Mieko Ueno* and Andrew Kehler

Grammatical and pragmatic factors in the interpretation of Japanese null and overt pronouns

Abstract: Pronoun interpretation in English has been demonstrated to be sensitive to an interaction between grammatical and pragmatically driven factors. This study investigated the interpretation of pronouns in Japanese, which has both null and overt forms. Thirty-two native speakers of Japanese per experiment participated in passage completion studies with transfer-of-possession contexts (Experiment 1) or implicit causality contexts (Experiment 2), varying prompt type, aspect, and topic/nominative-marking of the previous subject. Two judges annotated reference and coherence relations in the completed passages. Japanese overt pronouns were revealed to pattern closely with English overt pronouns in their sensitivity to pragmatic factors, whereas null pronouns showed a mixed resilience to pragmatic factors. Topic-marking only showed marginal effects on reference in limited contexts. Despite different degrees of sensitivity to pragmatic factors, Japanese null and overt pronouns were both mostly subject-biased, casting doubt on the existence of a division of labor between the two forms. There was also an intrinsic link between reference and coherence relations throughout the experiments. We discuss the overall results in terms of language specificity and universality, the latter of which includes interactions between grammatical and pragmatic factors and the importance of discourse coherence in the interpretation of various pronouns across languages.

Keywords: Japanese pronoun interpretation, discourse processing, crosslinguistic language processing

1 Introduction

Understanding pronoun resolution is crucial for understanding language comprehension. Pronouns provide a classic example of an underspecified linguistic

---

*Corresponding author: Mieko Ueno, Department of Cognitive Science, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093–0515, USA, E-mail: miueno@ucsd.edu
Andrew Kehler, Department of Linguistics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093–0108, USA, E-mail: akehler@ucsd.edu

form that must be semantically interpreted within a context, and as such the study of their behavior offers a window into questions about the sentence and discourse interpretation processes that take place around them.

Unlike English, which only has overt pronouns,[1] Japanese has both null and overt pronouns. Null pronouns in Japanese occur most commonly in subject position, but also occasionally in object position (Ueno and Polinsky 2009). Unlike in European null pronoun languages such as Italian and Spanish, there is no verbal agreement marking that aids the identification of null pronouns in Japanese.[2] Nevertheless, Japanese overt pronouns, especially the third person pronouns (e. g., *kare* 'he', *kanojo* 'she'), occur much less commonly than null pronouns (Martin 1975). In fact, the Japanese third person overt pronouns are generally considered to be direct translations of their European counterparts (cf. Fiengo and Haruna 1987).[3] However, they are increasingly incorporated into daily Japanese. For instance, a corpus count of *Asahi Shimbun* (a popular Japanese newspaper) articles shows that out of 11,073,167 sentences, *kare* was used 28,795 times and *kanojo* 14,209 times (Amano and Kondo 2000).[4]

The fundamental question addressed in this paper is how pronouns are interpreted in typologically distinct and similar languages. The present study investigates the interpretation of Japanese null and overt pronouns mainly in comparison to English overt pronouns and Korean null and overt pronouns, in order to shed light on language-universal and language-specific aspects of sentence and discourse processing. More specifically, the present study investigates the interaction of various grammatical and pragmatic factors in Japanese pronoun interpretation by testing a hypothesis that Japanese pronouns are more influenced by grammatical factors than English pronouns. It also investigates how different referential forms interact within and across languages by testing

---

1 Except for pro-drop by children and a few restricted cases by adults (imperatives, "diary-drop"; cf. Haegeman 1990; Rizzi 1994).

2 The only exception is honorifics, but they are not obligatory and only show the relative hierarchy between the speaker and the referent in a given sentence.

3 Historically, *kare* originated as a speaker-distal demonstrative 'that' referring to speaker-distal persons and things dating back to the eighth century. However, *kare* gradually lost its deictic force by losing competition with *a*-series distal forms. As a result, exophoric usage of *kare* was gradually replaced with anaphoric usage around the eighteenth century, and around 1900 became established as a third person masculine form used anaphorically, presumably due to the influence from third person pronouns in European languages. Literally, *kanojo* is a combination of the distal demonstrative determiner *kano* 'that' and *jo* 'woman'. The use of *kanojo* dates to 1900 as a third person feminine form (cf. Ishiyama 2008).

4 In colloquial speech, *kare* and *kanojo* are also used as 'boyfriend' and 'girlfriend', as well as vocative expressions (cf. Ishiyama 2008). However, we consider that the formal writing styles in the newspaper would allow a very few (if any) instances of such usages.

claims associated with the influential theory of Gundel et al. (1993), which predicts that the distributions of reference associated with certain competing referential forms within a language form a complementary distribution and that referential forms across languages pattern together in their behavior. Further, this work examines discourse structure and its effects on pronoun interpretation by testing a hypothesis that pronoun comprehension is influenced by underlying discourse coherence across languages.

1.1 Underlying questions on pronoun interpretation

1.1.1 Grammatical and pragmatic factors that determine pronoun interpretation

A wide variety of principles have been proposed to influence the interpretation of pronouns in English; the earliest proposals relied mostly on grammatical factors. Consider (1).

(1)   a. *John saw Bob at the station.*
       b. *He…*

The context sentence in (1a) contains two possible referents for the pronoun, one in subject position (*John*) and one in object position (*Bob*). The principle of subject preference (e. g., Crawley et al. 1990; Walker et al. 1994) favors referents that occupy the subject position of the previous clause. Since the subject of (1a) is *John*, John becomes the preferred referent of the subject pronoun *He* in (1b). Similarly, the grammatical-role parallelism preference (Smyth 1994) predicts that an ambiguous subject pronoun resolves preferentially to a subject antecedent, as does the first-mention privilege (Gernsbacher and Hargreaves 1988) in such examples.

However, recent work suggests that pronoun interpretation in English is driven by the interaction of grammatical and pragmatic factors, as well as underlying discourse structure. For instance, Stevenson et al. (1994) reported a series of story completion experiments that included passages with a transfer-of-possession (hereafter ToP) context sentence followed by an ambiguous pronoun prompt, as in (2).

(2)   Transfer-of-possession (ToP) context
       a. *John$_{\text{SOURCE}}$ handed a book to Bob$_{\text{GOAL}}$.*
       b. *He* _____

In this case, the subject *John* fills the Source thematic role and the oblique *Bob* fills the Goal role. The results revealed that Goal continuations occurred about

as frequently as Source continuations, an unexpected result in light of the three aforementioned preferences. Stevenson et al. (1994) attributed this effect to a comprehender's bias toward the end state associated with the event structure of the previous event. Specifically, since the Goal is the recipient of an object of transfer in ToP events, the Goal is presumably more salient than the Source with respect to the end state.

Rohde et al. (2006) tested this claim further by running a similar experiment but adding a condition with imperfective aspect as in (3), in light of the fact that the imperfective focuses attention on the ongoing development of an event rather than its end state. Because the Source in a ToP event is presumably still salient to the ongoing development of the event, they hypothesized that the imperfective condition would yield significantly more interpretations of pronouns to the Source than the perfective condition. This prediction was confirmed.

(3)   a. *John*$_{SOURCE}$ *was handing a book to Bob*$_{GOAL}$.
      b. *He* _____

In addition, again following Stevenson et al. (1994), Rohde and Kehler (2008) ran passages with pronoun prompts as in (2)–(3) against those with "free" prompts as in (4), in which participants not only chose who to refer to from subject position but also the form of reference used to do it.

(4)   a. *John handed/was handing a book to Bob.*
      b. _____

Results showed more references to the Source in the pronoun condition than in the free condition, indicating that pronouns overlay a grammatical subject bias on top of the biases that were revealed by the aspect manipulation.[5]

---

5 The difference between the referential patterns witnessed for free prompts (4a)–(4b) and pronoun prompts (2a)–(2b) and (3a)–(3b) is a potential problem for purely expectation-based approaches to pronoun interpretation, such as Arnold's (2001) Expectancy Hypothesis. In contrast to analyses that rely predominantly on a single heuristic, the Expectancy Hypothesis is a multiple-constraints approach according to which the accessibility of a referent, and hence its ability to be pronominalized by a speaker, is driven by the hearer's contextually-generated expectation that it will be mentioned again at a particular point in the discourse. Because the free-prompt condition can be seen to provide estimates for such expectations, on this hypothesis it is not clear why, in referentially-ambiguous contexts, biases seen in the pronoun prompt condition would display a significantly different distribution. Our experiments will provide a test of the predictions of the Expectancy Hypothesis in Japanese, a topic to which we return in Section 4.1.

In addition to reference, Rohde et al. (2006) also reported that the influence of aspect on pronoun interpretation in ToP contexts was correlated with certain semantic relationships inferred to hold between the two clauses (henceforth *coherence relations*; cf. Hobbs 1979, 1990; Kehler 2002; Rohde 2008), suggesting that the shift in the distribution of pronoun interpretations mirrors a shift in the distribution of coherence relations induced (also see Arnold 2001 for a related finding). Specifically, following a context sentence with a ToP verb like (2) or (3), the continuation sentences that participants produced that referred to the Source participated predominantly in Elaboration and Explanation relations (5a and 5b), whereas those that referred to the Goal participated predominantly in Occasion and Result (5c and 5d).

(5)　Source-biased and Goal-biased coherence relations
　　a. Elaboration: continuations that provide additional details about the eventuality (i. e., event or state) described in the context sentence
　　　e. g., *John handed a book to Bob. He did so slowly and carefully*.
　　b. Explanation: continuations that describe the cause of the eventuality described in the context sentence
　　　e. g., *John handed a book to Bob. He no longer had a use for it*.
　　c. Occasion: continuations that describe an eventuality that initiates from the end state of the eventuality described in the context sentence
　　　e. g., *John handed a book to Bob. He began reading it*.
　　d. Result: continuations that provide the effect or result of the eventuality described in the context sentence
　　　e. g., *John handed a book to Bob. He thanked him for the gift*.

Crucial to this pattern is the aforementioned fact that ToP verbs have a special property, whereby the prominence of the event participants varies across event structure: Whereas the Source (which also acts as an Agent) is central to the initial state and ongoing development of the event, it is the Goal – in its role as the recipient of the object of transfer – that is prominent when the event has been completed. Because the second clause in Elaboration relations provides additional details of the event and in Explanation it describes the cause of the event, both focus on the event's initial state and are thus biased to the Source for reference. On the other hand, because the second clause in Occasion and Result relations both describe what happened next, these relations focus on the end state of the previous event and thus favor the Goal for reference. Rohde and Kehler (2008) also showed that referential biases affect the distribution of coherence relations: Pronoun prompts, in introducing an independent subject bias, yielded more continuations with Source-biased relations compared to the

free prompt condition. Based on these results, Rohde et al. (2006), Rohde and Kehler (2008), and Rohde (2008) concluded that the preferences in pronoun interpretation cannot be adequately captured by superficial heuristics such as subjecthood and parallelism, and that coherence and coreference exert bidirectional influences on each other (also see Kehler et al. 2008 for an overview).[6]

Another set of studies that establish a role for semantics and pragmatics in reference interpretation have utilized implicit causality (hereafter "IC") verbs. IC verbs are said to impute the cause of the events they denote to either their subject or object NP, and as a result, they have been reported to impose divergent biases on the interpretation of a subsequent pronoun in English (Caramazza et al. 1977; McKoon et al. 1993; Rohde 2008; see Ferstl et al. 2011 for a verb list and norming data). For instance, *surprise* in (6a) is called an "IC1 verb," as the subject *John* (1st sentential NP) is likely to be the cause of the surprising event (having done something to surprise *Bill*), biasing the subject of the following *because*-clause toward being John. On the other hand, *praise* in (6b) is an "IC2 verb", as the object *Bill* (second sentential NP) is likely to be the cause of the praising event (having done something worthy of praise), biasing the subject of the following *because*-clause toward being Bill. Obviously this difference in referential preferences must be a result of the verb, since (6a)–(6b) are otherwise identical.

(6)   Implicit causality (IC) contexts
    a. IC1 (IC verbs with NP1 bias): *John surprised Bill because he got an A + in his math class.*
    b. IC2 (IC verbs with NP2 bias): *John praised Bill because he got an A + in his math class.*

As in ToP contexts, Rohde (2008) reported a tight link between reference and coherence relations with IC contexts. First, they demonstrated that both types of

---

6 Hobbs (1979, 1990) famously argued that pronoun interpretation is determined as a by-product of the establishment of coherence relations. In his analysis, pronouns are modeled as free variables that will get unified with the representations of their antecedents during the process of reasoning about discourse meaning. Similar to the Expectancy Hypothesis, this analysis predicts that pronouns are effectively inert, without contributing form-specific biases of their own (beyond morphosyntactically-imposed constraints such as number and gender, of course). Further, whereas the analysis predicts that coherence relations should affect coreference, it should not work the other way around: an occurrence of a pronoun should not cause the shift in coherence relations that Rohde and Kehler (2008) found. We return to discuss this analysis in the context of our experimental results in Section 4.3.

IC verbs yield significantly more continuations that participate in Explanation relations than do non-IC verbs. Second, they demonstrated that IC1 and IC2 verbs are biased toward subject and object reference respectively when the coherence relation is Explanation, replicating the results of previous studies that relied on prompts that explicitly contained the subordinating conjunction *because*.

1.1.2 Different referential forms within and across languages

There have been proposals on how different referential forms interact with each other within and across languages. For instance, the Givenness Hierarchy (GH) of Gundel et al. (1993) argued that the referential systems of different languages are organized around a common system of six categories of cognitive status. According to the GH, English overt pronouns and Japanese null pronouns require referents of the highest status (in focus), meaning that the referent is the current center of attention. On the other hand, English demonstratives and Japanese overt pronouns are associated with the second highest status (activated), requiring only that the referent is in the hearer's short term memory. In addition, the six cognitive statuses that comprise the GH participate in an implicational hierarchy, and are thus expected to give rise to scalar implicatures. For instance, while Japanese overt pronouns are compatible with both activated and in focus referents, Grice's (1975) Maxim of Quantity ("say as much as you need to say") predicts that the informationally stronger null form should be used for in focus referents, in turn predicting that overt pronouns will be used only for referents that are activated but not also in focus. This creates a division-of-labor effect, whereby the preferred referents of the two forms are predicted to fall roughly into complementary distribution.[7] Gundel et al. also suggested that the referential forms associated with the same status across languages should pattern together in their usage.

In accordance with Gundel et al. (1993), previous research on Japanese pronouns has associated null pronouns in Japanese with overt pronouns in

---

7 With respect to Japanese overt pronouns, which Gundel et al. (1993: 299) classify alongside demonstratives, they say, "What is of interest here is that demonstratives not only don't require the referent to be in focus, but often implicate that the referent is not currently in focus. This is particularly striking in cases where there is potential ambiguity of reference." They reason that demonstratives are a mechanism for topic shift precisely as a result of their division of labor with null pronouns. Our experiments will provide a test of this prediction, as they employ referentially-ambiguous contexts.

English. Kuroda (1965) showed that the Japanese equivalent of (7a) in English is quite unusual when all the pronouns are expressed (7b), and it is necessary to delete all the overt pronouns to get a fully felicitous sentence (7c).

(7)   a. George does his work when he feels like doing it.
       b. ?ジョージは　彼が　　彼の　　仕事を　したい時に
          George-wa   kare-ga   kare-no shigoto-o shi-tai-toki-ni
          George-TOP   he-NOM   he-GEN   work-ACC do-want-when-at
          彼の　　仕事を　　する。
          kare-no shigoto-o suru.
          he-GEN   work-ACC do
          'George does his work when he feels like doing his work.'
       c. ジョージは　したい時に　　　仕事を　する。
          George-wa shi-tai-toki-ni   shigoto-o suru.
          George-TOP do-want-when-at work-ACC do
          'George does (his) work when (he) feels like doing (it).'

Kuroda (1965) argued that there is a distributional correspondence between the overt pronouns in English and the null pronoun ("ø/zero/PRO" in his terminology) in Japanese, in that the null pronoun in Japanese corresponds to the whole class of English personal pronouns. Kuroda (1965) further argued that forms that have been generally considered to be personal pronouns in Japanese actually pattern much more like nouns in their behavior, in that they involve specializations in meaning that go beyond the typical grammatical features (number, person, gender) associated with personal pronouns. Related to this, Hoji (1990) argued that Japanese overt pronouns cannot be construed as bound variables as in *everyone$_i$ thought his$_i$ paper was the best* while English overt pronouns can, because *ka* in *kare* is analogous to *a* in the *ko/so/a/do* 'this/that/that/which' demonstrative paradigm in Japanese. The distal *a* demonstrative system including *ano* 'that (determiner)', *ano* 'that (noun)', and *asoko* 'there' is too deictic for its members to be construed as bound variables, and *kare* is not a pronoun but rather a demonstrative in Japanese. Kuroda's (1965) comparison between the Japanese null pronoun and the English overt pronouns was limited to their grammatical properties; it was not his goal to address the behavior of those forms in context. However, subsequent theoretical and experimental studies on the interpretation of Japanese pronouns (e. g., Kameyama 1985; Walker et al. 1994) adopted Kuroda's (1965) analogy between the two forms and assumed that the interpretation of null pronouns in Japanese is analogous to the interpretation of overt pronouns in other languages without a null form.

Competing with approaches that posit unidimensional salience scales are form-specific analyses (Brown-Schmidt et al. 2005; Kaiser and Trueswell 2008). On such analyses, there is no single notion of entity salience that governs referential form usage; instead different referential expressions exhibit varying degrees of sensitivity to different contextual factors. For instance, Kaiser and Trueswell (2008) reported on a study that compared the behavior of the Finnish pronoun *hän* 's/he' with the demonstrative *tämä* 'this' with respect to the grammatical role and word order position of possible antecedents. The results showed that *hän* is sensitive primarily to syntactic role (preferring the subject of the previous sentence) regardless of word order, whereas *tämä* is sensitive mainly to word order (preferring postverbal referents, with only a smaller secondary effect favoring objects over subjects). Crucially, this entails that the preferred antecedents of the two forms will sometimes overlap; for instance, for context sentences that display object-verb-subject word order, both forms exhibit a subject bias. This behavior is perfectly compatible with form-specific approaches as such accounts do not predict competing forms to generate division-of-labor effects. The results are surprising, however, for approaches based on a unidimensional salience hierarchy such as Gundel et al. (1993).

1.2 Previous research on null pronouns in Japanese and other languages

There have been a number of corpus studies that investigate the principles of argument ellipsis in null pronoun languages (e. g., Japanese: Fry 2003; Nariyama 2003; Ueno and Polinsky 2009; Chinese: Tao 1996; Spanish: Bentivoglio 1992; Sacapultec Maya: Du Bois 1987, among many others), and some studies also examine the linguistic mechanisms underlying null pronoun resolution to account for how referents of null pronouns are determined. For example, Nariyama (2003) proposed a "principle of direct alignment", which states that the subject always has to be higher on a person/animacy hierarchy than nonsubjects, as well as a "principle of argument ellipsis", which posits that a higher argument in terms of the person/animacy hierarchy and discourse salience (unless focused) tends to be elided, and formulated an algorithm for null pronoun resolution based on these principles.

In addition to corpus studies, a growing number of experimental studies have been performed on the interpretation of null pronouns in recent years (e. g., Japanese: Mazuka 1991; Walker et al. 1994; Ueno and Kehler 2010; Chinese: Yang et al. 1999, 2003; Spanish: Alonso-Ovalle et al. 2002; Fernandez-Salgueiro et al. 2007; Italian: Carminati 2002; Odawa: Christianson and Cho 2009). For instance, working within Centering Theory, Walker et al. (1994)

reported an influence of grammatical/information-structural factors found in a referent-choice experiment. Their results suggested a higher degree of salience for topic-marked (-*wa*; cf. Kuno 1973) than nominative-marked (-*ga*) referents, providing support for the notion that topics are by default the most salient discourse entity.[8] Similarly, a recent study by Christianson and Cho (2009) suggests that topical arguments in Odawa are more likely to be realized as null pronouns than nontopical arguments.

More recently, there have been several studies (Kim and Kaiser 2009; Kwon and Polinsky 2011; Kim et al. 2013; Kwon and Sturt 2013) that investigated the processing of null pronouns in Korean, which is syntactically similar to Japanese. For instance, Kwon and Sturt (2013) showed that the interpretation of Korean null pronouns is more dependent on discourse cues than morphosyntactic cues. When there is a possible antecedent given by the previous sentence in the discourse, a morphosyn-tactic mismatch between a null pronoun and a subsequent possible catacedent within the same sentence does not slow down reading times in Korean, while it does for overt pronouns in English (cf. Liversedge and Gompel n.d.).

Most relevant to the present study, Kim et al. (2013) investigated the inter-pretation of Korean null and overt pronouns. Their hypotheses and experi-mental design were modeled directly on Ueno and Kehler (2010), upon which contexts Experiment 1 of the current paper is based. Specifically, Kim et al. used a Korean version of our ToP stimuli and reported effects of a pragmatic bias for both null and overt pronouns. In their data, null and overt pronouns were equally subject-biased and there was no statistically supported effect (despite some numerical differences) of topic-marking. Kim et al. also reported a strong relationship between reference and coherence relations in Korean, following what Rohde and colleagues (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008) found for English and Ueno and Kehler (2010) found for Japanese. Despite the fact that, chronologically speaking, Kim et al. conducted their study based on (and hence following) Experiment 1 of the present study, for an ease of exposition we will treat the study as a parallel study to ours and incorporate the discussion of their results throughout the paper.

1.3 Research questions and predictions in the present study

Taken together, the foregoing work leaves a number of open questions for Japanese pronouns, which we address by adopting the passage completion

---

8 However, Walker et al. (1994) did not conduct any statistical analysis due to a small number of participants.

task used by Stevenson et al. (1994) and Rohde and colleagues (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008). These questions are grouped in terms of the following three categories:

(8)   Research questions in the present study
    a. What factors (grammatical, information-structural, pragmatic) influence the interpretation of Japanese null and overt pronouns?
    b. How do different referential forms interact within, and pattern across, languages?
    c. Do the factors in (a), by virtue of influencing the interpretation of Japanese pronouns, also influence how the discourse is likely to be continued?

Each of these major questions carries several subquestions, so we expand on them in turn. The first major question (8a) concerns the factors that influence the interpretation of Japanese null and overt pronouns. The previous work we have surveyed on English has imputed factors that are both grammatical (e. g., subjecthood) and semantic/pragmatic (e. g., context type per the ToP and IC manipulations discussed in Section 1.1). We ask whether the same factors affect both Japanese null and overt pronouns in a similar way. First, we will employ ToP contexts (Experiment 1) and IC contexts (Experiment 2) to test if Japanese pronouns are sensitive to pragmatic factors. If Japanese null and overt pronouns behave similarly to overt pronouns in English, we should see an effect of aspect (Experiment 1) and IC bias (Experiment 2) on the distribution of pronominal references. On the other hand, it could be that these forms are instead sensitive primarily to grammatical factors, in which case we would not find any effect of context type. Second, we will follow previous work (Stevenson et al. 1994; Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008). in employing a free prompt condition, from which participants can choose their own referring expressions. Because the studies of Stevenson et al. (1994) and Rohde and collegues have suggested that pragmatic manipulations to the context govern likelihood of subsequent mention generally rather than applying specifically to pronoun interpretation, we expect that any effect of context type on the pronominal forms will also be realized in the free prompt condition as well. Importantly, the referential patterns found in the free prompt condition will then provide a baseline with which to assess the existence of overlaid biases contributed by the use of a pronoun. If Japanese null and overt pronouns pattern similarly with English overt pronouns, we should see a greater number of references to the previous subject in the pronoun prompt conditions than in the free prompt condition. Further, if it is the case that the degree of bias to the subject is correlated with degree of reduction of linguistic form

(with null pronouns constituting more reduction than overt pronouns), we would further expect that null pronouns will be more biased to the subject than overt forms.

As noted above, previous work on Japanese suggests a role for information-structural factors as well, predicting that entities marked as topics are more salient, and hence will attract more pronominal references, than entities not so marked. We will also investigate this question by varying the morphological marking on the subject of a context sentence between -*ga* (nominative marker) and -*wa* (topic marker).[9] Following Walker et al. (1994) for Japanese and Christianson and Cho (2009) crosslinguistically, topic-marked antecedents are thus predicted to attract more pronominal references than nominative-marked antecedents.

We now consider our second question (8b), which is how different referential forms interact within, and pattern across, languages. Within Japanese, the major question is whether null and overt Japanese pronouns exhibit a division-of-labor effect. That is, if Gricean analyses like Gundel et al. (1993) are correct, then the use of an overt pronoun should implicate that the intended referent is not the one that would have been the preferred referent for a null pronoun; otherwise the speaker would have chosen the null form instead. This should in turn yield a complimentary distribution, whereby null pronouns display a strong subject bias and overt pronouns display a complementarily strong nonsubject bias. On the other hand, as mentioned above, no such division of labor is predicted on a form-specific approach (Kaiser and Trueswell 2008), according

---

9 An anonymous reviewer correctly points out that both topic and nominative markers in Japanese can be ambiguous. For instance, according to Kuno (1973), -*wa* can be "thematic" or "contrastive". The thematic -*wa* ('speaking of...') is suffixed to an NP that is either anaphoric (given in the discourse) or generic and is thus known to already be in the speaker and hearer's knowledge, while the contrastive ('X..., but...,') -*wa* does not have such a restriction and suffixed NPs receive prominent intonation in speech. In light of the fact that there is no preceding context or prominent intonation in our stimuli, however, we assume that *wa*-marking will not evoke a contrastive reading. Similarly, Kim et al. (2014) also show that topic-marked subjects in Korean yield contrastive reading only when prominent prosody is physically placed. We also note that the present study showed extremely low instances of contrastive coherence relations in the continuation sentences (Parallel – about 1% for both experiments), indicating that the context sentences were unlikely to have contrastive readings. Likewise, -*ga* can be "descriptive" to denote a neutral description of an action or temporary state, as well as have an "exhaustive listing" reading to denote an 'X (and only X)...' state. Again, in light of there being no preceding context in the stimuli, we assume that *ga*-marking will not induce an exhaustive-listing reading. There is also another usage of -*ga* for marking the object of stative transitive verbs, which is not relevant here.

to which different forms might exhibit overlapping patterns of reference. We will answer this question by analyzing the data collected from the experimental design described above.

We will also analyze this data to assess similarities and differences across the referential systems of different languages. Specifically, we will examine whether the behavior of Japanese null and/or overt pronouns patterns with that of English overt pronouns, or display a different character. If we follow the assumptions made in previous studies (e. g., Kameyama 1985; Walker et al. 1994) and the requirement for the IN FOCUS status proposed by Gundel et al. (1993), the distributions of reference associated with Japanese null pronouns should mirror those of English overt pronouns. On the other hand, there may be a more complex and nuanced crosslinguistic picture at work. On a form-specific approach like Kaiser and Trueswell's (2008), it would be unsurprising to find that different referential forms across languages, while exhibiting some common characteristics, nonetheless reveal a varying sensitivity to different factors. For example, if it turns out that Japanese null pronoun usage is governed more by grammatical and/or information-structural factors than English overt pronouns, then the character of the referential patterns will be different. In that case we then ask whether Japanese overt pronouns, which have been characterized as being demonstrative-like by the previous works, pattern closely with English overt pronouns. Alternatively, it could be that both null and overt pronouns in Japanese will behave differently from English pronouns, casting doubt on theories that posit unified crosslinguistic organizational principles for referential systems.

Finally, with respect to our third question (8c), we ask whether any factors revealed to be associated with these referential forms affect the distribution of ensuing coherence relations, as Rohde and colleagues found for English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008). This question will be answered by having judges annotate the completed passages with respect to coherence relations and examining the resulting distributions. We predict that factors that affect the interpretation of Japanese pronouns also influence how the discourse is likely to be continued. More specifically, following Rohde and colleagues, we predict that ToP contexts (Experiment 1) will reveal more Source-biased relations for imperfective than perfective contexts, so that the distribution of coherence relations will generally follow the pattern found for reference. For IC contexts (Experiment 2), following Rohde (2008), we predict more Explanation relations expressed in object-referenced IC2 continuations (continuations whose subject rementions the object of the context sentence) than in object-referenced IC1 continuations, reflecting the expectation that causal completions will typically mention the causally implicated entity. For the same reason, we predict more Explanation continuations in subject-referenced

IC1 than subject-referenced IC2 continuations. However, if Japanese pronouns (at least null pronouns) are not as pragmatics/discourse driven as English pronouns, it is possible that pronominal reference and coherence relations will not be as strongly associated as they are in English.

2 Experiment 1 – ToP contexts

In order to address our research questions, Experiment 1 first investigated the interpretation of Japanese pronouns with ToP contexts.

2.1 Methods

2.1.1 Participants

Thirty-two native speakers of Japanese recruited from the San Diego area participated in the study and were reimbursed for their time.

2.1.2 Materials

2.1.2.1 Stimuli

The experiment employed a $3 \times 2 \times 2$ design that varied prompt type (Null pronoun[10]/Overt pronoun[11]/Free), aspect (Perfective/Imperfective[12]), and topic/nominative-marking of the context sentence subject (*-wa/-ga*), as shown in (9).

---

10 The "subject-omission" prompt was used to indicate the presence of a null pronoun. A pilot study revealed that most participants were capable of continuing such prompts appropriately, which was confirmed in the actual study.

11 All overt pronoun prompts were topic-marked. This was done because the pilot study revealed that nominative-marked overt pronouns tend to be used to express an embedded subject of a complex sentence rather than a matrix subject. Topic-marking the pronoun resolved the issue.

12 Imperfectivity is not as straightforward to encode in Japanese as in English, since *-teita* 'was ~ing' is ambiguous between a perfective and imperfective reading depending on the verb (or VP) with which it co-occurs. Because ToP verbs typically express achievement events as a default, the more natural interpretation of *-teita* with these verbs is perfective. We therefore used *tokoro* 'was in the scene of' to "stretch out" instantaneous events and facilitate an imperfective reading of achievement events, in a manner similar to what the English progressive does to achievement events.

(9) Stimuli

太郎は/が　　次郎に　本を　　　　渡した/渡している　　ところだった。
Taro-wa/ga　　Jiro-ni　hon-o　　　watashita/watashi-te-iru　tokoro-datta.
Taro-TOP/NOM　Jiro-DAT　book-ACC　handed/hand-INF-ASP　　　scene-was
'Taro handed/was handing a book to Jiro.'
主語省略/彼は/自由＿＿＿＿＿＿＿＿＿＿＿＿＿＿
shugo-shoryaku/kare-wa/jiyu
subject-omission(Null)/he-TOP(Overt)/free(Free)

The 60 experimental stimuli each had context sentences with different ToP verbs. The Source referent (*Taro* in (9)) always appeared in subject position, and the Goal referent (*Jiro*) was the dative/'to'-marked indirect object of the sentence. All verbs described physical transfer events (e. g., 'hand', 'throw').

Fillers consisted of 40 sentences with transitive or intransitive nontransfer verbs in the perfective or imperfective aspect. The transitive verbs varied between active and passive voice, and adverbs, names, and gender-unambiguous overt pronouns served as prompts. The 100 sets of sentences instantiating the 12 ($3 \times 2 \times 2$) experimental conditions were placed in a Latin square design to create 12 parallel lists of 100 sentences, such that no one participant saw more than one sentence from each set.

2.1.2.2 Norming study
A norming study was conducted to test if stimulus sentences used in the experiment were natural. One particular concern was the naturalness of context sentences with topic-marked subjects in light of the absence of any previous discourse. Ten native speakers of Japanese (a different group from those who participated in the present Experiment) rated the context sentences distributed over four lists in a Latin square design and intermixed with fillers on a scale from "1" (very unnatural) to "5" (very natural).

The mean naturalness ratings on the four conditions and fillers ranged between 3.6 and 4.2 (Figure 1). Repeated measures analyses of variance (ANOVAs) were run with "aspect" and "topic/nominative-marking" as fixed within-group factors, and "participant" or "item" as a random factor. There was a significant effect of aspect [$F_1(1, 9) = 38.23$, $p < 0.0001$; $F_2(1, 59) = 29.88$, $p < 0.0001$], suggesting that Perfectives were rated higher than Imperfectives, but there was no significant main effect of topic-marking, nor any interactions involving aspect or topic-marking. Therefore, while Topic-marked context sentences were rated as natural as Nominative-marked context sentences, Imperfective
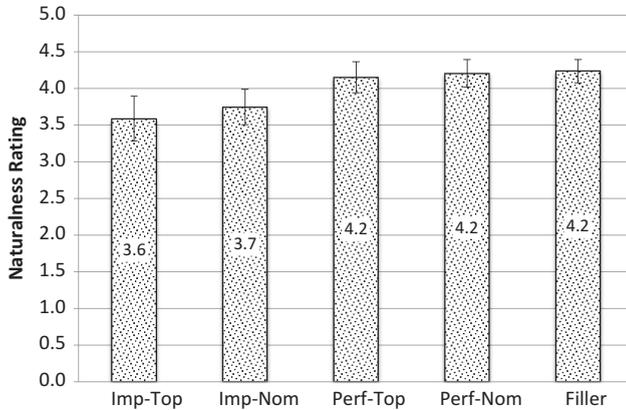
Figure 1: Naturalness ratings for stimulus and filler sentences (participant means with standard errors).

sentences were rated less natural than Perfective sentences. However, since Imperfective sentences were rated above the midpoint (3.7 out of 5 on average), we considered them natural enough and proceeded with the experiment.

2.1.3 Task

Using a web-based interface, participants were asked to write continuations for the 100 passages. They were instructed to imagine a natural continuation to the story, write the first continuation that came to mind and avoid humor.

2.1.4 Data annotation

Following previous studies on English, we focused our analysis on the interpretation of matrix subjects. However, identifying the matrix subject can be less straightforward in Japanese than in English. Since Japanese clauses may contain multiple null elements and are characterized by flexible and head-final word order, we first translated the continuations into English, thereby recovering the referents of null elements. For instance, if the original sentence in Japanese said 'felt happy because passed exam', detectable null pronouns were postulated, as in '(s/he) felt happy because (s/he) passed (her/his) exam' in the relevant English translation. The first author (who is a native speaker of Japanese) then

underlined the likely matrix subject of the given sentence for the subsequent annotation processes.

Two trained judges[13] – native speakers of Japanese but blind to the experimental hypotheses – annotated the referent of the matrix subject of each continuation sentence with respect to six categories: Source (*Taro* in (9)), Goal (*Jiro*), Theme (*book*), Both (*Taro* and *Jiro*), Other, and Unsure. The judges were instructed to do the annotation separately, without talking to each other. Then the first author compared their annotations and discarded the tokens the judges did not agree on, the tokens in which participants did not omit a subject even though they were given a null pronoun prompt, sentence fragments such as prepositional phrases only (which tended to occur with a null pronoun prompt), as well as overt pronoun prompt continuations which started with a null pronoun subject (thus making the overt pronoun an embedded subject). The tokens discarded in this way constituted about 17 % of the data.

The remaining tokens were then given back to the judges for annotating the coherence relations that held between each context sentence and continuation (Hobbs 1990; Kehler 2002; Rohde 2008). These relations include Elaboration, Explanation, Occasion, and Result as shown in (5) in Section 1.1.1, as well as Violated Expectation, Denial, Parallel, and Background as shown in (10) below.

(10) Additional coherence relations
   a. Violated Expectation: continuations that provide an unexpected result of the eventuality described in the context sentence
      e. g., *John handed a book to Bob. He showed no interest in reading it.*
   b. Denial: continuations that provide a reason why the eventuality described in the context sentence would be unexpected
      e. g., *John handed a book to Bob. He was not sure if Bob would like it.*
   c. Parallel: continuations that provide two events or states that share some common or contrasting property
      e. g., *John handed a book to Bob. He gave a magazine to him as well.*
   d. Background: continuations that provide the background to the eventuality described in the context sentence
      e. g., *John handed a book to Bob. He was Bob's friend.*

The judges resolved disagreements through discussion.

---

13 These two judges came from a group of Japanese research assistants in different combinations. A given sentence was always annotated by two judges.

2.1.5 Statistical analysis

Continuations with Source or Goal referents were sent to statistical analyses on reference. Mixed-model repeated measures ANOVAs were run on the percentage of Source referents as a function of the total number of Source and Goal referents. "Prompt type", "aspect", and "topic/nominative-marking" were used as fixed within-group factors, and "participant" or "item" was used as a random factor. Subsequent Tukey HSD comparisons were used to test for any significant differences among different prompt types, which will be the same for all the ANOVAs conducted for the present experiment. In addition, one-tailed paired t-tests were conducted on the percentage of Source referents over the total number of Source and Goal referents for each participant or item comparing Imperfectives vs. Perfectives or Topics vs. Nominatives within each of the three prompts, in order to test possible effects of aspect or topic-marking on each prompt type. In addition, the free prompt continuations were further divided according to their matrix subject type, namely, Null pronouns, Overt pronouns, and Names, and the same ANOVAs as above were conducted, in order to test if the pattern of results on reference would stay the same as the results with the predetermined prompts.

As for coherence relations, among the tokens referenced as Source or Goal, tokens that were judged as Elaboration, Explanation, Occasion, and Result were sent to further statistical analyses. Several other coherence relations which sometimes occurred – e. g., Violated-Expectation, Background, Parallel, and Denial (about 18 % of data in total) – were not included in the statistical analysis for the present experiment, as they were not strongly biased toward the Source or the Goal and were few in number for each category. We collapsed the proportion of Elaborations and Explanations (Source-biased relations) on the one hand and Occasions and Results (Goal-biased relations) on the other hand for each participant's/item's continuations, and conducted ANOVAs on the proportion of Source-biased relations over Source- and Goal-biased relations. As in reference, "prompt type", "aspect", and "topic/nominative-marking" were used as fixed within-group factors, and "participant" or "item" was used as a random factor. Furthermore, one-tailed paired t-tests were also conducted on the percentage of Source-biased relations over Source- and Goal-biased relations for each participant or item comparing Imperfectives vs. Perfectives or Topics vs. Nominatives within each of the three prompts, to test possible effects of aspect or topic-marking on each prompt type. In addition, ANOVAs and pairwise comparisons with the same factors as above were conducted on the proportion of Elaboration over Elaboration and Explanation relations, to test a post-hoc observation that the proportion of Elaborations within the Source-biased relations was consistently higher for Imperfectives than Perfectives for all prompt types.

An alpha level of 0.05 was used for all statistical tests, with a *p*-value up to 0.10 considered marginally significant.

2.2 Results and discussion

2.2.1 Reference

Figure 2 shows the proportion of Source and Goal referents for each prompt type and aspect (collapsed over topic/nominative-marking differences) averaged across participants. The Null pronoun conditions had about 80 % Source referents irrespective of aspect, while Overt and Free conditions varied by aspect.
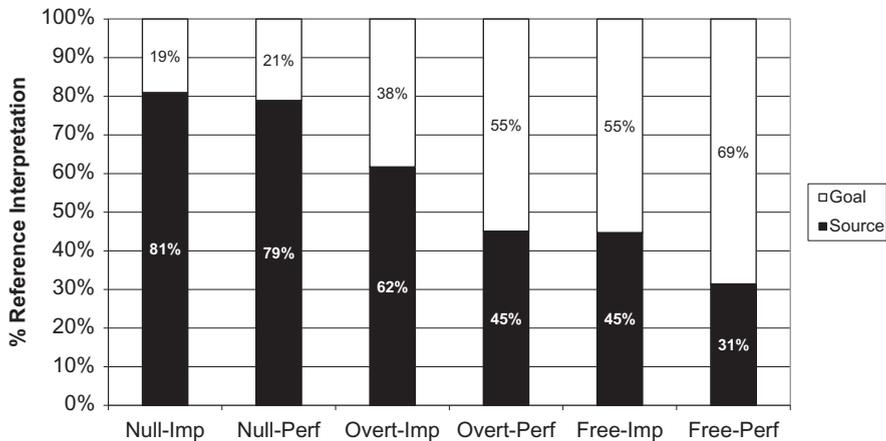

Figure 2: Proportion of source and goal referents for all conditions (collapsed over topic/nominative-marking differences).

ANOVAs run on the percentage of Source referents over the total number of Source and Goal referents revealed a significant main effect of prompt type $[F_1(2, 31) = 74.11, p < 0.0001; F_2(2, 59) = 64.10, p < 0.0001]$. Subsequent Tukey HSD comparisons found significant differences in the order of Null > Overt > Free by both participants and items, i. e., Null continuations were mostly Source biased, followed by Overt, followed by Free continuations (participant mean: Null 80 %, Overt 53 %, Free 36 %). Mirroring Rohde and colleagues for English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008), Japanese overt pronouns led to significantly more subject mentions of the Source than free prompts. In addition,

Japanese null pronouns led to even more subject mentions of the Source than overt pronouns. Our results are also partly consistent with Kim et al. (2013), who reported that Korean null and overt pronouns are about equally Source biased (72~78 % with no statistical difference), while free prompt continuations are the least Source biased (30 %). We further divided the free prompt continuations according to their matrix subject type, and performed the same ANOVA as above. The results yielded a significant main effect of subject type $[F_1(2, 31) = 67.24, p < 0.0001; F_2(2, 59) = 44.10, p < 0.0001$; Tukey HSD: Null > Overt > Name by both participants and items][14] which shows the highest proportion of Source referents for Null, followed by Overt, followed by Name continuations (participant mean: Null 80 %, Overt 40 %,[15] Name 15 %[16]). This is also mostly consistent with the findings of Rohde and colleagues in English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008) and Kim et al. (2013) in Korean,[17] in that Names were least Source biased.

There was also a significant main effect of aspect $[F_1(1, 31) = 15.81, p < 0.0001; F_2(1, 59) = 21.02, p < 0.0001]$, indicating that Imperfectives yielded more Source referents than Perfectives. Pairwise comparisons between Imperfectives and Perfectives within each prompt type revealed significant/marginal differences for Overt continuations $[t_1(31) = 3.72, p = 0.0004; t_2(59)^{18} = 3.95, p = 0.0001]$ and Free continuations $[t_1(31) = 1.59, p = 0.0611; t_2(52) = 3.47, p = 0.0005]$, but not for Null continuations. This suggests that Japanese overt pronouns show sensitivity to pragmatic factors in a way similar to English overt pronouns and Korean null and overt pronouns, but that is not the case for Japanese null pronouns.

Unlike prompt type or aspect, there were no significant main effects or interactions involving topic-marking. Figure 3 shows the proportion of Source

---

14 These ANOVAs on free prompt continuations only (with "subject type", "aspect", and "topic/nominative-marking" as factors) also revealed the same overall pattern of results as the ANOVAs on all prompt continuations as discussed below, including a significant main effect of aspect $[F_1(1, 31) = 6.36, p = 0.0122; F_2(1, 59) = 9.22, p = 0.0027]$ indicating that Imperfectives yielded more Source referents than Perfectives, as well as no significant effects involving topic-marking.

15 The difference between this percentage and the results of the overt pronoun condition (53 %) is probably due to many missing cells. Only 12 out of the 32 participants produced overt pronoun continuations with free prompts, and these tokens did not cover all the conditions with aspect and topic-marking variations covered by the predetermined overt prompt continuations.

16 Free prompt continuations were more Source-biased than Names as they also included pronoun continuations.

17 Kim et al. (2013) reported that their null and overt pronoun continuations within free prompts were equally Source-biased.

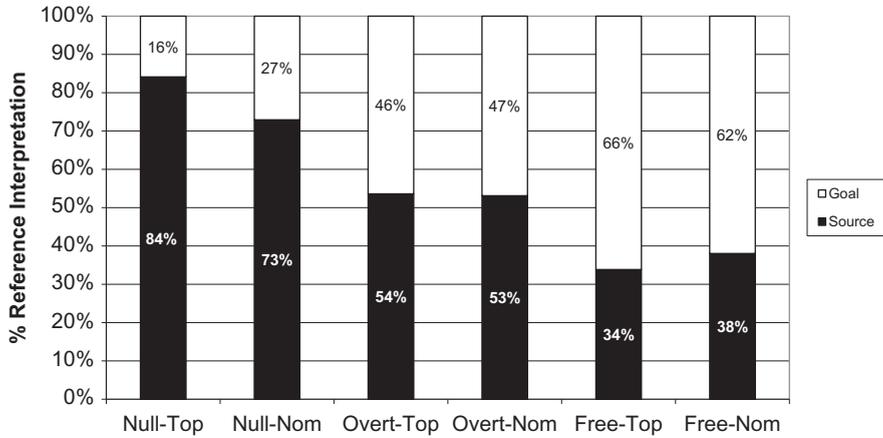**18** Some degrees of freedom vary due to missing cells.

Figure 3: Proportion of source and goal referents for all conditions (collapsed over aspectual differences).

and Goal referents for each prompt type and topic-marking (collapsed over aspectual differences) averaged across participants. Pairwise comparisons between Topics and Nominatives within each prompt type revealed no significant differences for any prompt type, although there was a marginal difference in Null continuations $[t_1(28) = 1.37, \ p = 0.0908; \ t_2(55) = 1.87, \ p = 0.0665]$ that favored Source referents in the Topic condition. Therefore, there was no consistent overall effect of topic-marking in the present data, unlike Walker et al. (1994) and like Kim et al. (2013).[19]

2.2.2 Coherence relations

Figure 4 shows the Source/Goal referent count for each coherence relation (collapsed over 12 experimental conditions) averaged across participants. As in English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008) and Korean (Kim et al. 2013), Elaboration and Explanation are highly Source biased whereas Occasion and Result are highly Goal biased.

We collapsed the proportion of Elaborations and Explanations (Source-biased relations) on the one hand and Occasions and Results (Goal-biased

---

19 However, it should be kept in mind that there were some numerical differences in Kim et al. (2013), and Walker et al. (1994) did not conduct any statistical analysis due to a small number of participants.
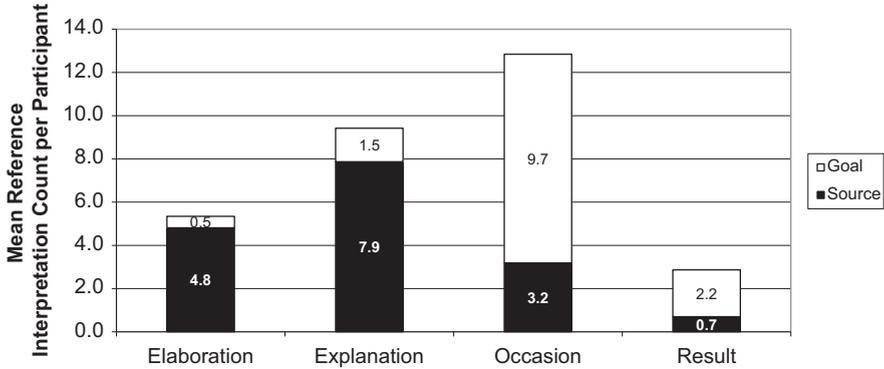
Figure 4: Mean source/goal referent count for each coherence relation (collapsed across different conditions).

relations) on the other hand for each participant's/item's continuations, and conducted repeated measures ANOVAs on the proportion of Source-biased relations over Source- and Goal-biased relations. There was a significant main effect of prompt type [$F_1(2, 31) = 25.94$, $p < 0.0001$; $F_2(2, 59) = 22.34$, $p < 0.0001$; Tukey HSD: Null > Overt > Free by participants, and Null > Overt, Free by items], indicating that Null prompt continuations led to the most Source-biased relations (Figures 5 and 6).
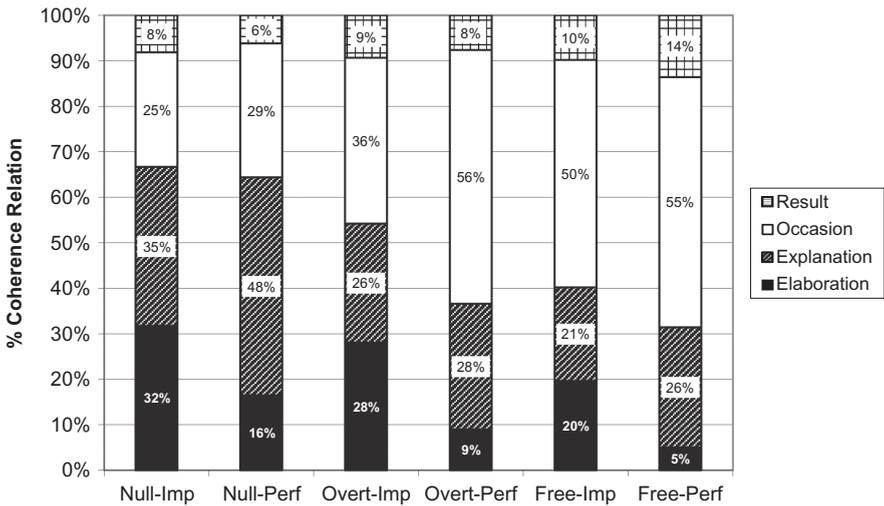

Figure 5: Proportion of coherence relations for all conditions (collapsed over topic/nominative-marking differences).
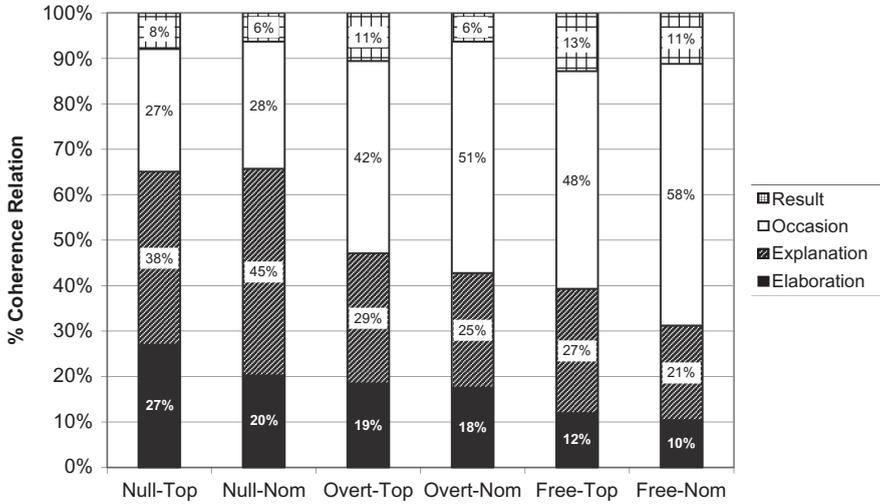
Figure 6: Proportion of coherence relations for all conditions (collapsed over aspectual differences).

There was also a significant main effect of aspect [$F_1(1, 31) = 9.75$, $p = 0.0018$; $F_2(1, 59) = 9.54$, $p = 0.0021$], suggesting that Imperfectives yielded more Source-biased relations than Perfectives (Figure 5). Pairwise comparisons between Imperfectives and Perfectives within each prompt type revealed significant/marginal differences for Overt [$t_1(31) = 3.68$, $p = 0.0004$; $t_2(59) = 3.82$, $p = 0.0002$] and Free [$t_1(29) = 1.64$, $p = 0.0559$; $t_2(44) = 1.90$, $p = 0.0319$] continuations. There was no effect of aspect for Null continuations, however, which is expected in light of the lack of effect on reference in that condition.

Furthermore, as was the case for reference, ANOVAs revealed no significant main effect of topic-marking. While Figure 6 indicates a small numerical trend of more Source-biased relations for topic-marked than nominative-marked Overt and Free continuations, pairwise comparisons between Topics and Nominatives within each prompt type revealed no statistically supported differences.

Finally, posthoc observation revealed that the proportion of Elaborations within the Source-biased relations was consistently higher for Imperfectives than Perfectives for all prompt types (see Figure 5). Additional ANOVAs run on the proportion of Elaboration over Elaboration and Explanation relations revealed a significant main effect of aspect [$F_1(1, 31) = 28.29$, $p < 0.0001$; $F_2(1, 59) = 30.63$, $p < 0.0001$] with no other statistically supported main effects involving prompt type or topic-marking or interactions involving these factors. Likewise, pairwise comparisons between Imperfectives and Perfectives within

each prompt type revealed significant/marginal differences for all types [Null: $t_1(23) = 2.42$, $p = 0.0120$; $t_2(39) = 2.54$, $p = 0.0075$; Overt: $t_1(24) = 3.36$, $p = 0.0013$; $t_2(39) = 3.06$, $p = 0.0020$; Free: $t_1(17) = 1.70$, $p = 0.0535$; $t_2(28) = 2.36$, $p = 0.0128$], supporting our observation that Imperfective continuations had a uniformly higher proportion of Elaboration relations across prompt types than Perfective continuations. On the other hand, pairwise comparisons between Topics and Nominatives within each prompt type revealed no statistically supported differences. Participants were therefore more likely to elaborate an event described as ongoing (imperfective) than one described as completed (perfective), indicating an effect of aspect on coherence relations that is independent of prompt type or topic-marking. This is also consistent with Rohde (2008) who reported that there were more Elaborations following imperfective than perfective context sentences in her ToP experiment.

In summary, the distribution of coherence relations generally followed the pattern found for reference, being consistent with the studies in English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008) and in Korean (Kim et al. 2013).

2.2.3 Experiment 1 summary and research questions

We are now in a position to offer preliminary answers to the questions posed in Section 1.3. We begin by considering the questions pertaining to (8a), which concerned the factors that influence the interpretation of Japanese null and overt pronouns. First, we found that context type influenced reference in the overt pronoun and free prompt conditions in the predicted direction (i. e., a greater number of subject references in the imperfective condition), but that there was no significant effect in the null pronoun condition. This is a striking result. The fact that the difference was found in the free prompt condition tells us that the effect of context type on subsequent mention is there, and that it was also found for overt pronouns shows that pronouns can in principle be sensitive to it. Null pronouns, on the other hand, appear to be less affected by this pragmatic bias, instead being strongly and uniformly Source biased. Second, an effect of prompt type was found across the board. This result demonstrates that the pronominal forms overlay a bias toward the previous subject as compared to the baseline distribution found in the free prompt condition. Further, the results reflected a correspondence between subject reference and reduction of linguistic form, whereby null pronouns were more strongly subject-biased than overt pronouns. Third, the manipulation of topic vs. nominative marking did not reveal an advantage for topic-marked antecedents, with the exception of

a marginal significance for null pronouns only. This result runs counter to the predictions of Walker et al.'s theory, but consistent with what Kim et al. (2013) found for Korean.

The second set of questions we posed (8b) pertained to how different referential forms interact within, and pattern across, languages. We first ask whether the interactions among different referential forms in Japanese create a division-of-labor effect across their referential biases. Clearly the answer is no. While there was a difference in degrees, both null and overt pronouns displayed a subject bias, and hence their referents were not in complimentary distribution. Note that the pattern held not only when the data across the prompt types was analyzed, but also when free prompt continuations were further divided according to their matrix subject type. Thus we cannot say that the use of the predetermined prompts obscured a possible division-of labor effect. The use of an overt pronoun does not seem to implicate that the referent is an entity other than what the preferred referent would have been if a null pronoun had been used (here, the Source). Recall that Kim et al.'s (2013) results for Korean likewise did not show a division-of-labor effect, and indeed the Source biases for null and overt forms were actually more similar to each other than we found for Japanese.

Second, we ask whether Japanese null pronoun interpretation is analogous to English overt pronouns as previous researchers have suggested, in light of the fact that they are the most reduced forms in their respective languages. The results again suggest that the answer is no. Although it is hard to draw conclusions based on comparisons across experiments with different stimuli, we observe that Japanese null pronouns were more Source biased (80 %) than English overt pronouns (60 % reported by Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008), and showed no sensitivity to the aspect manipulation unlike English overt pronouns. This is also different from Korean null pronouns which demonstrated a sensitivity to an aspect manipulation (Kim et al. 2013). Instead of null pronouns, overt Japanese pronouns actually patterned with English in their degree of Source bias (53 %) and sensitivity to the aspect manipulation. Furthermore, overt pronouns led to significantly more mentions of the Source than free prompts, demonstrating that, like English overt pronouns, Japanese overt pronouns overlay a subject bias on top of the pragmatic bias. In Korean, on the other hand, both null and overt pronouns pattern with English in terms of their sensitivity to pragmatic factors, although they are both more Source biased than English overt pronouns.

The answer to our last major question (8c) – whether factors that influence the interpretation of Japanese pronouns also influence how the discourse is likely to be continued – is clearly yes. The appearance of the null or overt pronoun in a prompt biased the continuation toward mentioning the previous

subject referent first, which in turn biased the participants toward continuing the story using a Source-biased coherence relation. Further, while the aspect manipulation in Null continuations created differences in the distribution of Source-biased relations – Imperfectives resulted in a greater number of Elaborations at the expense of Explanations – it did not change the allocation between Source- and Goal-biased relations, in accord with the fact that the manipulation resulted in no difference in the distribution between Source and Goal referents for null pronouns.

Whereas the results went mostly according to predictions, several are worthy of follow-up. Foremost among these are the lack of effect of context type on null pronoun interpretation (particularly in light of Kim et al.'s (2013) finding of such an effect in Korean) and the absence of a clear uniform effect of topic-marking. We therefore look to another way in which we can confirm or refute the results of this experiment. One approach is to vary the contexts across different verb types; whereas we manipulated aspect on a single type (ToP) in Experiment 1, we could instead vary the contexts by choosing opposing verb types likely to yield substantially different subsequent-mention biases. Whereas pronouns in (perfective) ToP contexts in English tend toward exhibiting equal biases across the two possible referents (Stevenson et al. 1994; Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008), pronouns in IC1 and IC2 contexts, for instance, are known to have stronger biases toward the subject and object NPs respectively (cf. Ferstl et al. 2011).[20] Thus varying IC type would potentially offer a stronger test of the lack of effect of context on null pronoun interpretation. Similarly, IC2 contexts may also provide a stronger test of the lack of effect of topic-marking on subject referents, since IC2 biases point strongly away from the subject. We pursue this idea in Experiment 2.

3 Experiment 2 – IC contexts

Experiment 2 tested our observations from Experiment 1 by comparing contexts that contain IC1 and IC2 verbs, which have been documented to impose highly divergent biases toward the interpretation of a subsequent pronoun in English (Caramazza et al. 1977; McKoon et al. 1993; Rohde 2008; Ferstl et al. 2011; *inter*

---

20 This is particularly the case when the operative coherence relation is Explanation, which in many previous IC experiments is forced by the inclusion of the connective 'because' in the prompt. Since we will categorize the passage completions according to coherence relation, we can analyze not only the effect of context type on the completions as an ensemble, but on the subset in which Explanation relations are operative as well.

alia). The results of Experiment 1 lead us to the following expectations. First, in terms of factors, we expect that null and overt pronouns will both display a subject bias as compared to the free prompt condition, with the bias being greater for null pronouns. We also expect that topic-marking will not yield substantial differences in both reference and coherence relations. However, since the results of Experiment 1 showed a marginal effect on reference with null pronoun continuations, we will again examine what type of effect topic-marking has in IC contexts. The role of pragmatic influences is the major question in the present experiment. The results found with ToP contexts lead us to expect that Japanese null pronouns will be uniformly subject-biased for both reference and coherence relations, unlike English and Korean pronouns. On the other hand, Japanese overt pronouns are expected to pattern with English and Korean in demonstrating sensitivity to pragmatic biases, in that overt continuations following an IC1 context sentence will be more subject-biased than overt continuations following an IC2 context sentence. Aspect is not expected to affect pronoun interpretation as it did in Experiment 1, however, since there is no reason to expect that entity prominence varies across the event structure for IC verbs in the way that it does for transfer-of-possession verbs. Finally, we again predict that whereas Japanese null and overt pronouns will be subject-biased to different degrees, their referential patterns will not display a complementary distribution that is characteristic of a division-of-labor effect.

In terms of coherence relations, we continue to predict that factors that affect reference also influence how the discourse is going to be continued. For the present experiment, we predict that there will be more Explanations expressed in object-referenced IC2 continuations than in object-referenced IC1 continuations, and likewise more Explanations in subject-referenced IC1 than subject-referenced IC2 continuations. In addition, following Experiment 1, we predict more Elaborations for Imperfectives than Perfectives, reflecting a presumed tendency to elaborate on an ongoing (imperfective) rather than completed (perfective) event, despite the fact that aspect is not predicted to affect reference.

3.1 Methods

3.1.1 Participants

Thirty-two native speakers of Japanese recruited from the San Diego area (but different from those in Experiment 1) participated in the study and were reimbursed for their time.

3.1.2 Materials

3.1.2.1 Stimuli

The experiment employed a $3 \times 2 \times 2 \times 2$ design varying prompt type (Null pronoun/Overt pronoun/Free), IC context (IC1/IC2), aspect (Perfective/Imperfective), and topic/nominative-marking (-*wa/ga*) of the context sentence subject, as shown in (11).

(11)   Stimuli

> 太郎は/が   次郎を   驚かした/驚かしている   ところだった。
> *Taro-wa/ga* *Jiro-o*   *odorokashita/odorokashi-te-iru* *tokoro-datta.*
> Taro-TOP/NOM Jiro-ACC surprised/surprise-INF-ASP   scene-was
> 'Taro surprised/was surprising Jiro.'
> 主語省略/彼は/自由_____
> *shugo-shoryaku/kare-wa/jiyu*
> subject-omission(Null)/he-TOP(Overt)/free(Free)

The 48 experimental stimuli each had context sentences with 24 IC1 or 24 IC2 verbs, which were adapted from the lists of IC verbs used in Garvey and Caramazza (1974) and Rohde (2008) (which she in turn adapted from those in McKoon et al. 1993) and then modified to fit Japanese. They were all in the active voice, with a subject referent (*Taro* in (11)) and an object referent (*Jiro*).

We used a separate sentence for eliciting a continuation. Although IC continuations are typically framed in a 'because'-clause, a Japanese 'because'-clause has to precede its main clause, reversing the necessary order of a context sentence and a continuation sentence for the experiment.[21] Further, using a separate sentence for eliciting a continuation allows us to analyze referential patterns across different coherence relations as we did in Experiment 1 (using only 'because' prompts would give us nearly 100 % Explanations). Rohde (2008: Experiment IV) compared continuations with and without a 'because'-prompt, and found that an overall IC bias remained for IC1 and IC2 continuations when 'because' was omitted, albeit with the strength of the biases being reduced since not all continuations were Explanations, the relation for which the IC bias is relevant. (However, the referential biases between the 'because' and full-stop

---

21 There is a 'because' construction for a separate sentence (*Nazenara~karada* 'The reason why is because~') that can follow the main clause, but such a construction is quite marked and not used very often.

prompt conditions were statistically identical when only the Explanation continuations in the latter condition were analyzed.) Based on these results, we therefore expect a significant difference between IC1 and IC2 biases even when the subsequent clause is a separate sentence instead of a 'because'-clause. As in Experiment 1, -*wa* versus -*ga* marking was included as a manipulation to continue to test effects of topic-marking in IC contexts. The aspect manipulation was also kept in the present experiment primarily to test for effects on the distribution of coherence relations.

Fillers consisted of 52 context sentences, containing 24 ToP verbs with the indirect objects all scrambled to the beginning of the sentence and 28 transitive or intransitive nontransfer verbs in the perfective or imperfective aspect. The transitive verbs varied between active and passive voice, and adverbs, names, and gender-unambiguous overt pronouns served as prompts. The sets of sentences instantiating the 24 (2 verb types × 3 prompts × 2 aspects × 2 topic/nominative-markings) experimental conditions were placed in a Latin square design to create 12 parallel lists[22] of 100 sentences, such that no one participant saw more than one sentence from each set.

3.1.2.2 Norming study

As in Experiment 1, a norming study was conducted to test if stimulus sentences used in the experiment were natural. Ten native speakers of Japanese (a different group from those who participated in Experiment 1 or 2 but the same group as those who participated in the Experiment 1 norming study) rated the context sentences distributed over four lists in a Latin square design and intermixed with fillers on a scale from "1" (very unnatural) to "5" (very natural).

The mean naturalness ratings on the eight conditions and fillers ranged between 3.6 and 4.5 (Figure 7). ANOVAs were run with "IC context", "aspect", "topic/nominative-marking" as fixed within-group factors, and "participant" or "item" as a random factor. There was a significant effect of aspect [$F_1(1, 9) = 71.90$, $p < 0.0001$; $F_2(1, 59) = 62.86$, $p < 0.0001$], suggesting that Perfectives were rated higher than Imperfectives. There were no other significant main effects nor interactions involving the within-group factors. Therefore, while Topic-marked context sentences were rated as natural as Nominative-marked context sentences, Imperfective sentences were rated less natural than Perfective sentences as in Experiment 1. However, since Imperfective sentences were rated above the midpoint (3.8 out of 5 on average), we again considered them natural enough and proceeded with the experiment.

---

22 Since all the participants saw 24 IC1 verbs and 24 IC2 verbs, IC context did not increase the number of the lists needed.
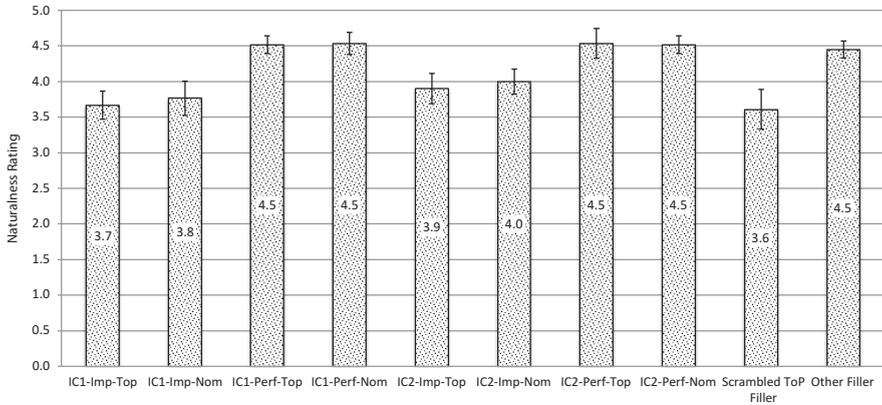
Figure 7: Naturalness ratings for stimulus and filler sentences (participant means with standard errors).

3.1.3 Task

As in Experiment 1, participants were asked to write continuations for the 100 passages in a web interface. They were instructed to imagine a natural continuation to the story, write the first continuation that came to mind, and avoid humor.

3.1.4 Data annotation

As in Experiment 1, participants' continuations were translated into English recovering null elements, then the matrix subject of each continuation was underlined for the subsequent annotation processes. Two trained judges[23] – native speakers of Japanese but blind to the experimental hypotheses – annotated the referent of the matrix subject of each continuation sentence with respect to six categories: Subject (*Taro* in (11)), Object (*Jiro*), Theme (*book* in ToP verbs as in (9)),[24] Both (*Taro* and *Jiro*), Other, and Unsure. The judges were instructed to do the annotation without talking to each other. As in Experiment 1, the first author compared their annotations and discarded the cases the judges did not agree on, the cases in which

---

23 As in Experiment 1, these two judges came from a group of Japanese research assistants in different combinations.

24 Although the "theme" category was irrelevant to IC context sentences, some filler sentences required it.

participants did not omit a subject with a null pronoun prompt, sentence fragments, and overt pronoun prompt continuations which started with a null pronoun subject (14.5 % of the data). The remaining tokens were then given back to the judges for annotating coherence relations as described in Section 2.1.4.

3.1.5 Statistical analysis

As in Experiment 1, tokens referenced as Subject or Object were sent to statistical analyses on reference. ANOVAs were run on the percentage of Subject referents as a function of the total number of Subject and Object referents. "Prompt type", "IC context" (i. e., IC1 or IC2), "aspect", and "topic/nominative-marking" were used as fixed within-group factors,[25] and "participant" or "item" was used as a random factor. Subsequent Tukey HSD comparisons were used to test any possible significant differences among different prompt types, which will be the same for all the ANOVAs conducted for the present experiment. In addition, one-tailed paired t-tests[26] were conducted again on the percentage of Subject referents over the total number of Subject and Object referents for each participant or item comparing IC1 vs. IC2 contexts, Imperfectives vs. Perfectives, or Topics vs. Nominatives within each of the three prompts to test possible effects of IC context, aspect, and topic-marking on each prompt type. In addition, the free prompt continuations were further divided according to their matrix subject type, namely, Null pronouns, Overt pronouns, and Names, and the same ANOVAs as above were performed, to test if the pattern of results will stay the same as the predetermined prompts. Likewise, tokens with Explanation relations (31 % of IC1 and 47 % of IC2 continuations in the present experiment)[27] were separated and the same ANOVAs and t-tests were run to test for the effect of IC contexts within Explanation continuations only.

As it is customary to analyze results for IC data separately by verb type, we also separated continuations with IC1 and IC2 contexts and investigated them individually to test for specific patterns in each category. ANOVAs were run on

---

25 "IC context" was nested in 'item' for items analyses. This was the same for all the ANOVAs involving "IC context" as a factor in the present experiment.

26 Grouped t-tests were conducted for items analyses comparing IC1 and IC2 contexts. Grouped t-tests were used for all the item analyses of pairwise comparisons between IC1 and IC2 contexts in the present experiment.

27 As expected, these numbers are higher than about 25 % Explanations with ToP continuations in Experiment 1, suggesting that the lexical semantics of IC verbs create a stronger-than-usual expectation for an explanation (Kehler et al. 2008).

the percentage of Subject referents as a function of the total number of Subject and Object referents with "prompt type", "aspect", and "topic/nominative-marking" used as within-group factors and "participant" or "item" used as a random factor. The same methods as those used for the IC1&2 data combined for Tukey HSD comparisons and pairwise comparisons (excluding IC1 vs. IC2 comparisons) were also used. Tokens with Explanation relations only were also sent to the same ANOVAs and t-tests to further investigate a post-hoc observation on continuations with all relations that IC2 continuations were influenced by topic-marking while IC1 continuations were not.

As for coherence relations, all the tokens referenced as Subject or Object were used for further statistical analyses in the present experiment (thus no tokens were discarded based on coherence relations unlike in Experiment 1). Explanations and Elaborations were analyzed in the following ways. For Explanations, ANOVAs were run on the proportion of Explanations over all the relations with "IC context" (IC1/IC2), "reference", "prompt type", "aspect", and "topic/nominative-marking" as within-group factors and "participant" or "item" as a random factor. One-tailed paired t-tests were also conducted on the percentage of Explanation relations over all the relations for each speaker or item comparing IC1 vs. IC2, Imperfectives vs. Perfectives, or Topics vs. Nominatives within Subject-referenced and Object-referenced continuations, to test for possible effects of IC context, aspect, or topic-marking on Subject-referenced and Object-referenced continuations. These comparisons also resulted in testing if there would be more Explanations with IC1-Subject than IC2-Subject continuations and with IC2-Object than IC1-Object continuations as predicted. As in reference, we also separated continuations in IC1 and IC2 contexts and investigated them on the distribution of Explanations individually to test for any specific pattern within each IC type. ANOVAs were run on the percentage of Explanations over all the relations with "reference", "prompt type", "aspect", and "topic/nominative-marking" as within-group factors and "participant" or "item" as a random factor. The same methods as IC1&2 combined above for Tukey HSD comparisons and pairwise comparisons (excluding IC1 vs. IC2 comparisons) were also used. In addition, after one of our predictions about the relationship between reference and Explanations in IC2 contexts was not borne out, an analysis of the referential patterns associated with three syntactic constructions that offered a potential explanation for why that was the case was also performed (see Section 3.2.2.2).

Furthermore, to test our prediction based on Experiment 1 that the proportion of Elaborations within the Source-biased relations would be higher for Imperfectives than Perfectives in IC contexts as well, we ran separate ANOVAs on the proportion of Elaborations over Elaborations and Explanations. "IC context", "prompt type",

"aspect", and "topic/nominative-marking" were used as within-group factors,[28] and "participant" or "item" was used as a random factor. One-tailed paired t-tests were also conducted on the percentage of Elaborations over Elaborations and Explanations for each speaker or item comparing Imperfectives vs. Perfectives within each of the three prompts to test a possible effect of aspect on each prompt type.

An alpha level of 0.05 was used for all statistical tests, with a *p*-value up to 0.10 considered marginally significant.

3.2 Results and discussion

3.2.1 Reference

Figure 8 shows the proportion of Subject and Object referents for each prompt type and IC context (collapsed over aspect and topic/nominative-marking differences) averaged across participants. As with ToP contexts in Experiment 1, Null continuations showed the highest proportion of Subject reference followed by Overt and then Free continuations. Furthermore, IC1 contexts yielded more Subject reference than IC2 contexts, yet unlike in Experiment 1, all prompt types including Null continuations appeared to vary by IC context.
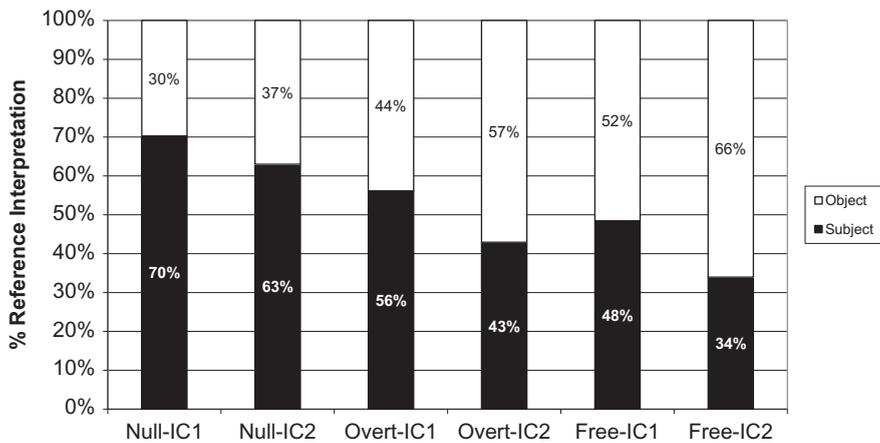

Figure 8: Proportion of Subject and Object referents for all conditions with all coherence relations (collapsed over aspectual and topic/nominative-marking differences).

---

28 "Reference" was not included as a factor, as the number of tokens was insufficient to accommodate the factor for repeated measures ANOVAs and the factor was not related to our research question here.

For IC1&2 continuations combined, ANOVAs run on the percentage of Subject referents over the total number of Subject and Object referents revealed a significant main effect of prompt type [$F_1(2, 31) = 28.77$, $p < 0.0001$; $F_2(2, 46) = 33.26$, $p < 0.0001$; Tukey HSD: Null > Overt, Free by participants, and Null > Overt > Free by items] suggesting that Null continuations were most Subject-biased, followed by Overt continuations, followed by Free continuations (participant mean: Null 66 %, Overt 49 %, Free 43 %). As in Experiment 1, we further divided the Free continuations according to their matrix subject type and performed the same ANOVAs as above. The results yielded a significant main effect of subject type [$F_1(2, 31) = 37.45$, $p < 0.0001$; $F_2(2, 46) = 27.46$, $p < 0.0001$; Tukey HSD: Null, Overt > Name by participants, and Null > Overt > Name by items][29] indicating the highest proportion of Subject referents for Null continuations at least by items, as well as the lowest proportion of Subject referents for Name continuations (participant mean: Null 65 %, Overt 53 %, Name 14 %). Therefore, as in Experiment 1 with ToP contexts, Null continuations showed the highest proportion of Subject referents followed by Overt continuations and then Free continuations, which can be further broken up to show the least Subject bias for Names.

As predicted, there was also a significant main effect of IC context [$F_1(1, 31) = 26.47$, $p < 0.0001$; $F_2(1, 46) = 10.38$, $p < 0.0001$], indicating that IC1 contexts yielded significantly more Subject referents than IC2 contexts. Pairwise comparisons between IC1 and IC2 contexts within each prompt type also revealed marginal/ significant differences for Null [$t_1(31) = 1.55$, $p = 0.0659$; $t_2(46) = 1.88$, $p = 0.0329$], Overt [$t_1(31) = 3.48$, $p = 0.0008$; $t_2(46) = 2.46$, $p = 0.0089$], and Free [$t_1(31) = 1.82$, $p = 0.0389$; $t_2(46) = 3.63$, $p = 0.0004$] continuations. However, unlike the case of ToP contexts in Experiment 1, here all prompt types varied by IC context, i. e., IC contexts affected the proportion of Subject reference even for Null continuations (albeit marginally by subjects). Furthermore, ANOVAs conducted on only the continuations with Explanation relations (Figure 9) revealed a significant main effect of IC context [$F_1(1, 31) = 89.46$, $p < 0.0001$; $F_2(1, 46) = 43.48$, $p < 0.0001$],[30] indicating

---

29 These ANOVAs on the free prompt continuations only also revealed the same overall pattern of results as the ANOVAs on all prompt continuations as discussed below, including a significant main effect of IC context [$F_1(1, 31) = 10.76$, $p = 0.0012$; $F_2(1, 46) = 15.99$, $p = 0.0002$] indicating that IC1 contexts yielded more Subject referents than IC2 contexts, as well as no significant effects involving topic-marking.
30 These ANOVAs on Explanation continuations only also revealed a significant main effect of prompt type [$F_1(2, 31) = 4.04$, $p = 0.0184$; $F_2(2, 46) = 7.75$, $p = 0.0005$; Tukey HSD: Null > Free by
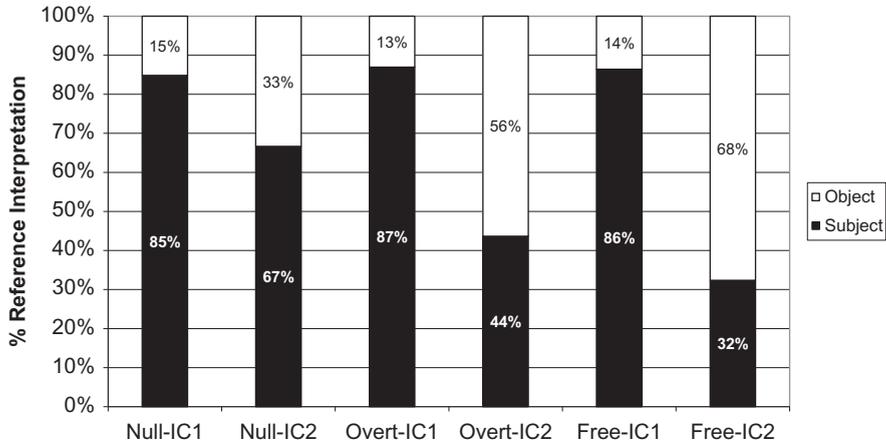
Figure 9: Proportion of Subject and Object referents for all conditions with explanation continuations only (collapsed over aspectual and topic/nominative-marking differences).

again that IC1 contexts yielded significantly more Subject referents than IC2 contexts. Pairwise comparisons between IC1 and IC2 contexts within each prompt type revealed significant differences for all prompt types [Null: $t_1(23) = 2.82$, $p = 0.0048$; $t_2(43) = 4.47$, $p < 0.0001$; Overt: $t_1(22) = 5.45$, $p < 0.0001$; $t_2(43) = 3.14$, $p = 0.0015$; Free: $t_1(31) = 1.82$, $p = 0.0389$; $t_2(36) = 6.18$, $p < 0.0001$]. These comparisons showed larger numerical differences and stronger statistical differences as indicated by smaller p-values than comparisons with all the relations combined, indicating a stronger effect of IC contexts in Explanation continuations only as in English (Rohde 2008). Perhaps most importantly, the distribution of references with null pronouns was significantly affected by the strongly opposing biases for our two verb classes in Explanation contexts, suggesting that the less definitive effect revealed by the analysis over all coherence relations above, and the lack of a parallel effect in Experiment 1, were due in part to the biases across contexts being less strongly opposed. In addition, notably, these ANOVAs on Explanations also showed a significant IC context x prompt type interaction [$F_1(2, 31) = 6.38$, $p = 0.0019$; $F_2(2, 46) = 3.63$, $p = 0.0276$], indicating that while IC1 continuations were consistently Subject-biased across prompt type (above 85%), IC2 continuations were affected by prompt type.

---

participants and items], suggesting that Null continuations were most Subject-biased while Free continuations were least Subject-biased. (Overt continuations were numerically in between Null and Free continuations and did not differ statistically from either condition.) There were no statistically supported effects involving aspect or topic-marking.

Returning to IC1&2 continuations with all coherence relations, unlike in ToP contexts in Experiment 1, there was no significant main effect of aspect. Pairwise comparisons between Imperfectives and Perfectives within each prompt type only revealed significant/marginal differences in Overt continuations $[t_1(31) = 3.17, \ p = 0.0017; \ t_2(47) = 1.64, \ p = 0.0535]$ that favored Subject referents in the Imperfective condition. Overall, therefore, an end-state focus with perfective contexts shown in ToP contexts does not seem to apply to IC contexts, as predicted by the fact that only in ToP contexts do we have reason to expect that entity salience will vary across event structure. In addition, there was no significant main effect of topic-marking (Figure 10). Pairwise comparisons between Topics and Nominatives within each prompt type revealed no significant differences for any prompt type, although there was a marginal difference in Overt continuations $[t_1(31) = 1.39, \ p = 0.0876; \ t_2(47) = 1.64, \ p = 0.0539]$ that favored Subject referents in the Topic condition. Overall, therefore, there was no significant overall effect of topic-marking on reference, consistent with the results of Experiment 1.
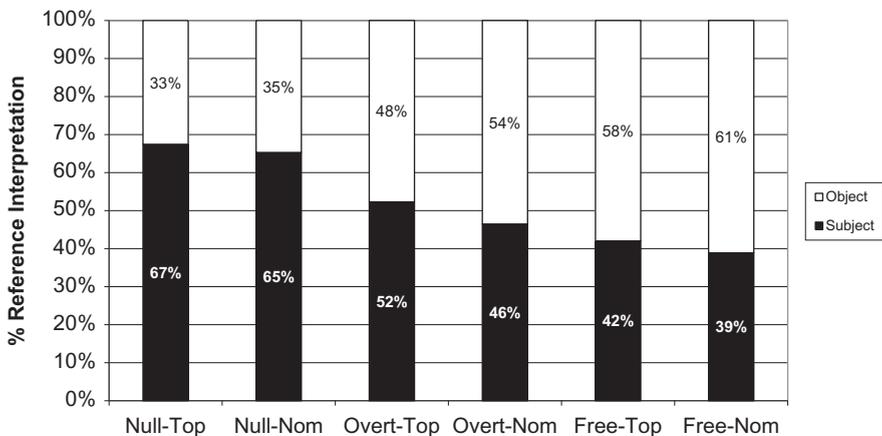

Figure 10: Proportion of Subject and Object referents for all conditions with all coherence relations (collapsed over IC1/2 and aspectual differences).

We then separated continuations in IC1 and IC2 contexts and investigated them individually. For IC1 continuations only, ANOVAs revealed a significant main effect of prompt type $[F_1(2, 31) = 11.00, \ p < 0.0001; \ F_2(2, 23) = 12.63, \ p < 0.0001;$ Tukey HSD: Null > Overt, Free by both participants and items], showing that Null continuations were most Subject-biased, followed by Overt and Free continuations. There were no significant main effects of aspect or topic-making

nor any significant interactions involving them. Pairwise comparisons between Imperfectives and Perfectives within each prompt type revealed a significant/marginal difference only in Overt continuations [$t_1(30) = 2.43$, $p = 0.0106$; $t_2(23) = 1.54$, $p = 0.0690$] that favored Subject referents in the Imperfective conditions. Pairwise comparisons between Topics and Nominatives within each prompt type revealed no significant differences for any prompt type. Therefore, there was no reliable overall effect of aspect or topic-marking in IC1 continuations.

For IC2 continuations only, ANOVAs revealed a significant main effect of prompt type [$F_1(2, 31) = 19.62$, $p < 0.0001$; $F_2(2, 23) = 21.28$, $p < 0.0001$; Tukey HSD: Null > Overt, Free by both participants and items], showing again that Null continuations were most Subject-biased, followed by Overt and Free continuations. As in the cases of IC1&2 contexts combined and IC1 contexts only, there was no significant main effect of aspect. Pairwise comparisons between Imperfectives and Perfectives within each prompt type revealed no statistically supported differences for any prompt type, either. However, unlike in IC1 continuations, there was a marginal/significant main effect of topic-marking [$F_1(1, 31) = 3.32$, $p = 0.0692$; $F_2(1, 23) = 4.33$, $p = 0.0380$], indicating that Topics were more Subject-biased than Nominatives overall. Pairwise comparisons between Topic and Nominative contexts within each prompt type revealed marginal/significant differences for Null [$t_1(29) = 1.38$, $p = 0.0884$; $t_2(23) = 2.19$, $p = 0.0193$] and Overt [$t_1(31) = 1.13$, ns; $t_2(23) = 1.43$, $p = 0.0830$] continuations, but not for Free continuations (Figure 11).
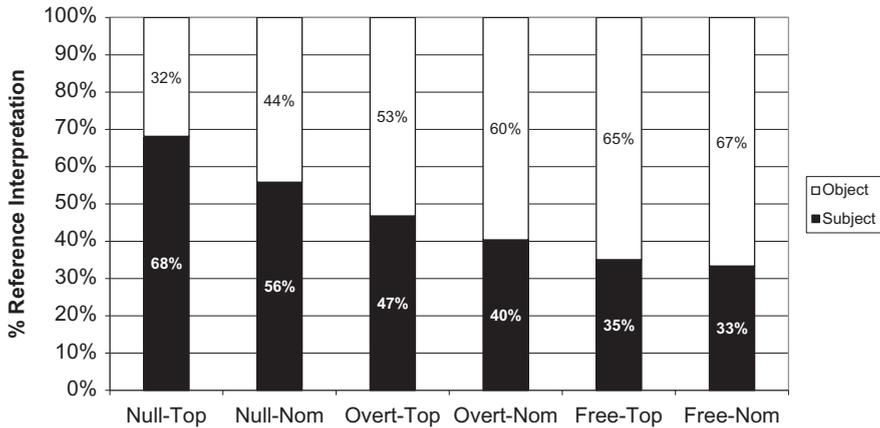

Figure 11: Proportion of Subject and Object referents with all coherence relations in IC2 contexts (collapsed over aspectual differences).

3.2.2 Coherence relations

Figures 12 and 13 show the proportion of coherence relations for all conditions collapsed over prompt type and topic-marking for IC1 continuations and IC2 continuations. In the following sections, we will discuss the pattern of results found for Explanations and Elaborations.
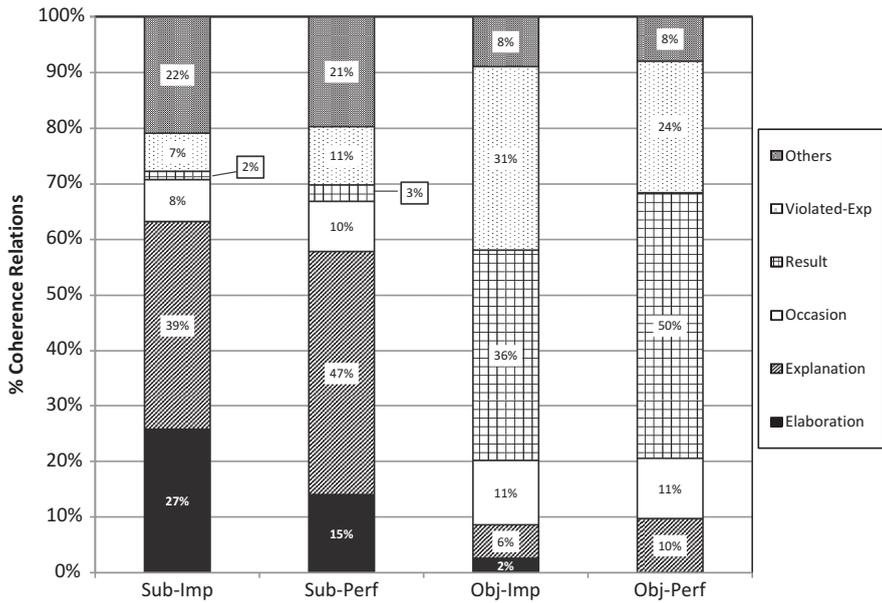

Figure 12: Proportion of coherence relations for IC1 continuations (collapsed over prompt type and topic-marking differences).

3.2.2.1 Explanations

For IC1&2 continuations combined, ANOVAs on the proportion of Explanations over all relations revealed a significant main effect of reference [$F_1(1, 31) = 16.60$, $p < 0.0001$; $F_2(1, 46) = 20.20$, $p < 0.0001$], whereby Subject-referenced continuations were more Explanation-biased than Object-referenced continuations. There was also a significant main effect of IC context [$F_1(1, 31) = 38.05$, $p < 0.0001$; $F_2(1, 46) = 25.16$, $p < 0.0001$], whereby IC2 continuations were more Explanation-biased than IC1 continuations. Pairwise comparisons between IC1 and IC2 contexts within Subject-referenced and Object-referenced continuations revealed a significant difference for Object-referenced continuations [$t_1(29) = 7.15$, $p < 0.0001$; $t_2(46) = 5.77$, $p < 0.0001$], in that Object-referenced IC2 continuations were more Explanation-biased than Object-referenced IC1
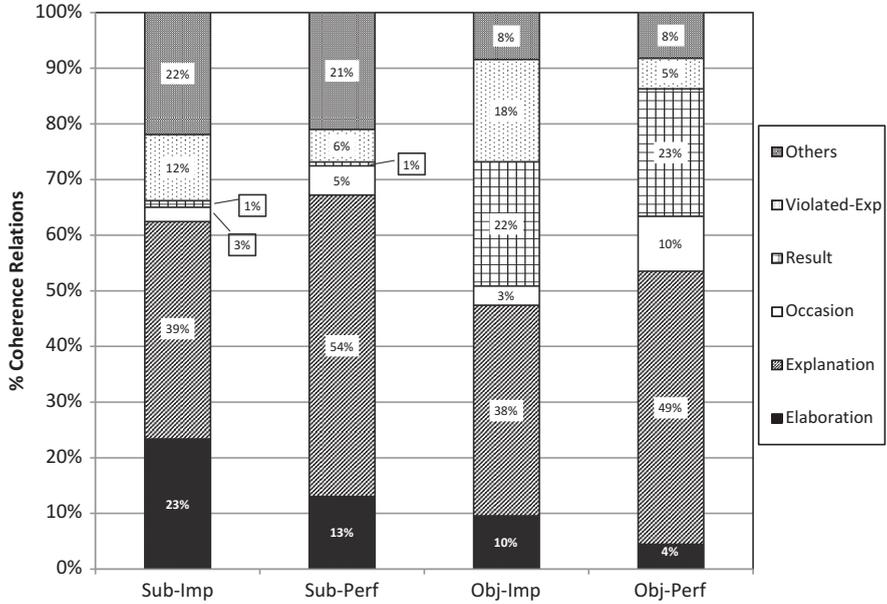
Figure 13: Proportion of coherence relations for IC2 continuations (collapsed over prompt type and topic-marking differences).

continuations. However, contrary to our predictions, there was no statistically supported IC1 vs. IC2 difference for Subject-referenced continuations. Consistent with these pairwise comparisons, there was also a significant IC context x reference interaction [F_1(1, 31) = 26.52, $p < 0.0001$; F_2(1, 46) = 27.14, $p < 0.0001$]. This finding differs from what Rohde (2008) found for English, who reported that the majority of Explanation continuations in IC1 contexts in English are subject-referenced, whereas the majority of Explanations in IC2 contexts are object-referenced. We will investigate this point further in Section 3.2.2.2 below.

There was a significant main effect of prompt type [F_1(2, 31) = 9.82, $p < 0.0001$; F_2(2, 59) = 9.98, $p < 0.0001$; Tukey HSD: Null > Overt, Free by both participants and items], suggesting that Null continuations were most Explanation-biased. There was also a significant main effect of aspect [F_1(1, 31) = 8.41, $p = 0.0038$; F_2(1, 46) = 8.99, $p = 0.0028$], indicating that Perfective continuations were more Explanation-biased than Imperfective continuations. Pairwise comparisons between Imperfective and Perfective contexts within Subject-referenced and Object-referenced continuations revealed significant differences both by participant and item for Subject-referenced continuations [t_1(31) = 2.31, $p = 0.0139$; t_2(47) = 3.40, $p = 0.0007$], and only by item for Object-referenced continuations [t_1(29) = 1.00,

p = ns; $t_2(45)$ = 2.55, p = 0.0072]. These pairwise comparisons overall indicate that Perfectives are more Explanation-biased than Imperfectives, which is probably due to Elaborations trading off with Explanations in Imperfectives (see Section 3.2.2.3). On the other hand, there was no significant main effect of topic-marking nor any statistically supported pairwise difference due to topic-marking within Subject and Object-referenced continuations.

For IC1 continuations only, ANOVAs revealed a significant main effect of reference [$F_1(1, 31)$ = 46.68, p < 0.0001; $F_2(1, 23)$ = 52.73, p < 0.0001], indicating again that Subject-referenced continuations were more Explanation-biased than Object-referenced continuations. There was a significant main effect of prompt type [$F_1(2, 31)$ = 4.71, p = 0.0094; $F_2(2, 23)$ = 4.38, p = 0.0131; Tukey HSD: Null > Overt by both participants and items], showing that Null continuations were more Explanation-biased than Overt continuations. On the other hand, there was no significant main effect of aspect or topic-marking, significant pairwise difference due to aspect or topic-marking, nor significant interaction involving any of the within-group factors. Therefore, IC1 continuations were basically influenced by reference and prompt type only.

In terms of IC2 continuations, unlike in IC1&2 continuations combined or in IC1 only continuations, ANOVAs revealed no significant main effect of reference, suggesting that both Subject and Object-referenced continuations were equally Explanation-biased. It thus seems more likely that the effect of reference found in IC1&2 continuations combined was mostly driven by IC1 continuations. As in IC1 contexts, there was a significant main effect of prompt type [$F_1(2, 31)$ = 5.27, p = 0.0054; $F_2(2, 23)$ = 4.38, p = 0.0131; Tukey HSD: Null > Overt, Free by both participants and items], showing that Null continuations were more Explanation-biased than Overt and Free continuations. Unlike in IC1 continuations, however, there was a significant main effect of aspect [$F_1(1, 31)$ = 6.17, p = 0.0134; $F_2(1, 23)$ = 6.14, p = 0.0135], indicating that Perfectives were more Explanation-biased than Imperfectives. Therefore, it seems more likely that the effect of aspect found in the IC1&2 combined data was mostly driven by IC2 continuations. Pairwise comparisons between Imperfective and Perfective contexts within Subject-referenced and Object-referenced continuations revealed a significant difference for Subject-referenced continuations [$t_1(29)$ = 2.55, p = 0.0081; $t_2(23)$ = 2.97, p = 0.0034], in that Subject-referenced Perfective continuations were more Explanation-biased than Subject-referenced Imperfective continuations, but there was no statistically supported difference for Object-referenced continuations. As in Experiment 1, the effect of aspect overall seems to be due to more Elaborations trading off with Explanations in Imperfectives, which will be discussed further in Section 3.2.2.3. Again there was no significant main effect of topic-marking nor any significant pairwise difference involving topic-marking.

3.2.2.2 Object-referenced embedded subject (OES) and Subject-referenced Experiencer + Object-referenced Stimulus (SE-OS) constructions

The preceding section revealed a much higher proportion of Explanations with Subject-referenced IC2 continuations than we had predicted. To solve the mystery, we further investigated the Subject-referenced continuations with Explanation relations and found that IC2 continuations in this group tend to fall into three categories, exemplified in (12)–(14). In (12), although the matrix subject of the continuation sentence refers to the subject of the previous sentence, the embedded subject refers to the object of the context sentence, which actually is the causal agent of the event denoted by the context sentence (OES construction). In (13), a passive sentence is used in which the object of the context sentence is an explicit or implicit causal agent. In (14), the subject of the continuation sentence is referenced as the subject of the context sentence and takes the Experiencer role (cf. Dowty 1991), while the object is referenced as the object (or something related to the object) of the context sentence and takes the Stimulus role, which is the causally-implicated referent (SE-OS construction).

(12) Object-referenced embedded subject (OES) explaining the cause of the context sentence (5% (7/139) for IC1 and 13% (17/126) for IC2 continuations)

    a. 康人が　　達郎を　　正している　ところだった。
       Yasuto-ga　Tasturo-o　tadashi-te-iru　tokoro-datta
       Yasuto-NOM Tasturo-ACC correct-INF-ASP scene-was
       主語省略
       shugo-shoryaku
       subject-omission
       'Yasuto was correcting Tasturo. subject omission'

    b. 達郎が　　　このままでは　　道を　　誤ると　　考えていた。
       Tatsuro-ga　kono-mama-de-wa　michi-o　ayamaru-to　kangae-te-i-ta
       Tatsuro-NOM this-state-at-TOP　path-ACC miss-that　think-INF-ASP-PST
       '(He) was thinking that Tatsuro would take the wrong path with the given situation.'

(13) Passive with the context object as an explicit/implicit agent (3% (4/135) for IC1 and 6% (7/126) for IC2 continuations)

    a. リサヨは　知世に　　感謝している　ところだった。
       Risayo-wa　Tomoyo-ni　kanshashi-te-iru　tokoro-datta
       Risayo-TOP Tomoyo-DAT thank-INF-ASP　scene-was

主語省略
shugo-shoryaku
subject-omission
'Risayo was thanking Tomoyo. subject omission'

b. 危ない　　ところを　助けて　　もらった　からだ
   Abunai　　*tokoro-o*　*tasuke-te*　*morat-ta*　*kara-da*
   dangerous　scene-ACC　save-INF　get-PST　　because-is
   'Because (she) got saved in a dangerous situation.'

(14)   Subject-referenced Experiencer subject + Object-referenced Stimulus object (including things related to the preceding object) (SE-OS) (9% (12/139) for IC1 and 33% (41/126) for IC2 continuations)

     a. 千秋は　　　民子を　　　　ほめた。　主語省略
       Chiaki-wa　*Tamiko-o*　　*home-ta*.　*shugo-shoryaku*
       Chiaki-TOP　Tamiko-ACC　praise-PST　subject-omission
       'Chiaki praised Tamiko. subject omission'

     b. 自分の　　生徒の　　　　成長を　　　　　　見届けるのが
       Jibun-no　*seito-no*　　*seicho-o*　　　　*mitodokeru-no-ga*
       Self-GEN　student-GEN　development-ACC　see-that-NOM
       うれしかった。
       ureshi-katta.
       glad-was
       '(She) was glad to see the development of (her) own student.'

Overall, only 17 % (23/139) of IC1 continuations met the above criteria while 50 % (63/126) of IC2 continuations did. These sentences were coded as 'Subject-referenced' according to our coding rule, but the object of the context sentence actually functions somewhat subject/agent-like. To investigate these observations, we conducted statistical analyses on OES and SE-OS construction continuations as below.

    With the Subject-referenced continuations with Explanation relations only, ANOVAs were run on the proportion of OES sentences over all continuations with 'IC context' and 'prompt type' as within-group factors,[31] and 'participant' or 'item' as a random factor. The results revealed a marginal/significant main effect

---

31 Unlike other ANOVAs, "reference" was not used as a factor, as every sentence examined in this section was Subject-referenced. Likewise, neither 'aspect' nor 'topic-marking' was included as a factor, as the number of tokens was insufficient to accommodate these factors for repeated measures ANOVAs, and these factors were not related to our research question.

of IC context [F_1(1, 31) = 2.98, p = 0.0857; F_2(1, 46) = 6.44, p = 0.0127], indicating that IC2 continuations had a higher proportion of OES sentences than IC1 continuations. There was no significant main effect of prompt type or IC context x prompt type interaction. One-tailed grouped t-tests[32] were also conducted on the percentage of OESs over Subject-referenced Explanation continuations for each speaker or item comparing IC1 vs. IC2 contexts within each prompt type, which revealed a significant/marginal difference for Null continuations only [t_1(18)[33] = 2.30, p = 0.0169; t_2(42) = 1.50, p = 0.0699], suggesting that Null continuations in IC2 contexts were the major cause of the difference between IC1 and IC2 continuations. This fact suggests that this complex embedded structure was used as a way to reconcile the strong bias toward subject referents for null pronouns with the IC2 biases toward the object. In the case of the passive construction in (13), there were not enough examples to run statistical analyses on. Nonetheless, the data seem to follow the same pattern: There were more instances of passive tokens in IC2 (7 cases) than in IC1 (4 cases) continuations, and the majority of passive sentences (7 cases) were Null continuations. In all of the 11 passive continuations, the subject referred to the subject of the context sentence, but actually the object of the context sentence functions as the implicit/explicit agent of a given causal event. Finally, again with Subject-referenced continuations and Explanation relations only, ANOVAs were run on the proportion of SE-OS sentences over all the continuations with 'IC context' and 'prompt type' as within-group factors. The results revealed a significant main effect of IC context [F_1(1, 31) = 23.51, p < 0.0001; F_2(1, 46) = 7.95, p = 0.0061], indicating that IC2 continuations had a higher proportion of SE-OS sentences than IC1 continuations. There was no significant main effect of prompt type or IC context x prompt type interaction. However, pairwise comparisons between IC1 and IC2 continuations within each prompt type revealed a significant difference for Overt continuations only [t_1(38)[34] = 3.81, p = 0.0002; t_2(42) = 3.81, p = 0.0003], suggesting that Overt continuations in IC2 contexts were the major cause of the difference between IC1 and IC2 continuations. Whereas such cases were catalogued as IC2 Explanations with a subject first-mention, they are not actually counterexamples to the bias since the matrix subject did not play the role of the causer in the (embedded) IC event.

To summarize, there were more OES, passive, and SE-OS constructions in IC2 than IC1 continuations, explaining why there was a much higher

---

32 Grouped t-tests were conducted for both participant and items analyses due to too many missing cells to conduct paired t-tests.

33 Some degrees of freedom vary due to missing cells.

34 Some degrees of freedom vary due to missing cells.

proportion of Explanations with Subject-referenced IC2 continuations than we had predicted. Faced with a referential bias toward the subject but a conflicting pragmatic bias toward the object in IC2 contexts, Japanese speakers often resolved the conflict by using a construction that allowed both biases to be satisfied at once, specifically one which puts the causally-implicated event participant in a grammatical position other than the matrix subject. The well-known association between IC2 verbs and object reference in Explanation relations assumes that the causality-implicated referent will be mentioned from subject position. As such, whereas on the surface these sentence types appear to violate this association, their prevalence in IC2 contexts is actually evidence that both the referential and coherence-driven biases are in fact driving the speaker's choice not only of referential form but also of syntactic construction.

3.2.2.3 Elaborations

As in Experiment 1, the proportion of Elaborations within the Source/Subject-biased relations (Explanation and Elaboration) appeared higher for Imperfectives than Perfectives in the present experiment. With IC1&2 continuations combined, ANOVAs run on the percentage of Elaborations over Elaborations and Explanations revealed a significant main effect of aspect [$F_1(1, 31) = 19.00$, $p < 0.0001$; $F_2(1, 46) = 13.19$, $p = 0.0003$], indicating that Imperfective continuations were more Elaboration-biased than Perfective continuations, as predicted. Pairwise comparisons between Imperfectives and Perfectives within each prompt type revealed significant/marginal differences for Null [$t_1(26) = 2.40$, $p = 0.0120$; $t_2(34) = 1.20$, ns], Overt [$t_1(28) = 2.80$, $p = 0.0046$; $t_2(34) = 4.01$, $p = 0.0002$] and Free [$t_1(21) = 1.43$, $p = 0.0836$; $t_2(21) = 1.02$, ns] continuations by participants. Although some item analyses did not reach significance, the proportion of Elaborations within the Source/Subject-biased relations had an overall trend of being higher for Imperfectives than Perfectives for all prompt types, being consistent with Experiment 1. In addition, there was no significant main effect of topic-marking or significant interactions involving the within-group factors like in Experiment 1, nor was there any significant main effect of IC context, indicating that both IC1 and IC2 continuations were equally Elaboration-biased. However, there was a significant main effect of prompt type [$F_1(2, 31) = 3.85$, $p = 0.0220$; $F_2(2, 46) = 2.77$, $p = 0.0059$; Tukey HSD: Overt > Null by participants and Overt > Null, Free by items], suggesting that Overt continuations were more Elaboration-biased than Null or Free continuations. This is different from Imperfectives having a uniformly higher proportion of Elaborations across prompt types in Experiment 1. The reasons for the effect of prompt type are unclear.

3.2.3 Experiment 2 summary and our research questions

We now return to the research questions posed in Section 1.3 beginning with (8a), which asked what factors influence the interpretation of Japanese null and overt pronouns. Analysis of these factors revealed that, as in Experiment 1, Null continuations showed the highest proportion of Subject reference, followed by Overt continuations and then Free continuations. Further, analysis of the Free continuations showed that Name continuations were the least Subject-biased. This further confirms the correspondence between the degree of linguistic reduction and the bias toward the subject found in Experiment 1. Also as predicted, there was a main effect of IC context whereby IC1 contexts yielded significantly more Subject reference than IC2 contexts. However, unlike in Experiment 1, all prompt types including Null continuations varied by IC context, with the most significant effect seen in the subset of the data that contained Explanation continuations. One of the primary motivations for this experiment was to see whether the lack of a similar effect in Experiment 1 was due to a true insensitivity of null pronouns to pragmatic biases, or whether the aspectual manipulation on ToP verbs was simply too weak to bring out an effect. The results of the current experiment support the latter explanation.

A particularly interesting pattern was revealed by the examination of the Explanation-only subset of the data (Figure 9). IC1 continuations were highly (above 85 %) Subject-biased across prompt type. This pattern is different from ToP contexts, in which both context types (Imperfective and Perfective) were affected by prompt type. We suspect that this is a result of the bias toward subjects being already near ceiling compared to that of objects, and thus there was no detectable effect of using either type of pronoun prompt. On the other hand, IC2 continuations were affected by prompt type, yielding the expected pattern whereby Japanese null and overt pronouns add a subject bias working in concert with the opposing pragmatic bias. Therefore, despite the lack of a significant effect for null pronouns in Experiment 1, we hypothesize that grammatical and pragmatic factors interact in generating the interpretation biases for both pronoun types, bringing the results in line with what Kim et al. (2013) found for Korean null and overt pronouns and Rohde and colleagues found for English overt pronouns (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008).

Following the results of Experiment 1, there was no reliable overall effect of topic-marking nor one in IC1 continuations only. However, there was a trend (significant by items but marginal by participants) in favor of topic-marking increasing subject reference in IC2 contexts. Thus, it could be that topic-marking only begins to have a measurable effect in those cases in which the pragmatic biases point strongly away from the subject. Walker et al. (1994) reported a

similar result with their Japanese passage 'Taro hit Jiro in front of all the other people. ø ignored ø all day.' They reported that when the subject 'Taro' is nominative-marked, the object 'Jiro' is the preferred referent of the following null subject pronoun, presumably due to the IC2 bias associated with the verb 'hit' and knowledge about the situation being described. However, when 'Taro' is topic-marked, topic-marking seems to counteract this pragmatic preference for 'Jiro' such that 'Taro' becomes the preferred referent. This could further explain the lack of statistically-supported effect of topic-marking in Experiment 1 as well as in Kim et al.'s (2013) Korean ToP experiment. That is, only when context gives rise to a pronounced subsequent mention bias toward the object, such as we see in IC2 contexts, is the speaker's decision to topic-mark the previous subject potentially understood by the comprehender to signal a different intended referent for a pronoun.

We now move on to our second set of questions in (8b), specifically on how different referential forms interact within and pattern across languages. Within Japanese, the first question was whether the opposition between null and overt pronouns would display a division-of-labor effect. The answer again is no. Again, while null pronouns were more subject-biased than overt pronouns, overt pronouns were also subject-biased. Furthermore, as in Experiment 1, the pattern of results was consistent in showing no division-of-labor effect both with the predetermined prompt continuations and when free prompt continuations were further divided by their matrix subject type. Thus we cannot say that using the predetermined prompts obscured a possible division-of-labor effect.

Comparing across languages, the present experiment revealed that Japanese null pronouns differed from English overt pronouns in showing a more limited sensitivity to pragmatic biases. Specifically, whereas Rohde and colleagues found that English overt pronoun biases were affected by both the ToP-aspect and the IC-type manipulations (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008), Japanese null pronoun biases were not affected by the aspect manipulation, and only to a limited extent (exhibiting the smallest difference among the three prompt types) by the IC manipulation. Like the results of Experiment 1, the results of this experiment therefore indicate that Japanese null pronoun interpretation is not completely analogous to English overt pronoun interpretation. Instead, the results continue to indicate that overt pronouns in Japanese pattern more closely with English overt pronouns in demonstrating a more consistent sensitivity to pragmatic biases.

Finally, with respect to question (8c), which asked whether factors that influence the interpretation of Japanese pronouns also influence how the discourse is likely to be continued, the close link between reference and coherence

relations found in Experiment 1 was also manifest in the present experiment. Just like the aspect manipulation in ToP contexts changed reference and the proportion of Source- and Goal-biased coherence relations in Experiment 1, IC contexts influenced both reference and coherence relations. The pattern of results mostly followed the pattern found in English reported by Rohde (2008), whereby there were more Explanations expressed in Object-referenced IC2 continuations than in Object-referenced IC1 continuations. In addition, as Rohde (2008) reported in her ToP study (and we reported in Experiment 1 of the present study), Imperfective continuations in the present experiment showed more Elaborations for describing ongoing events than Perfective continuations in IC contexts. These again came at the expense of another subject-biased relation, Explanation. One difference with Experiment 1 and previous work on English, however, was the unexpectedly high proportion of Explanations with Subject-referenced IC2 continuations; reference to the non-causally-implicated referent did not yield the predicted drop in Explanation relations. A detailed analysis of the continuations revealed that this was largely due to more OES, passive, and SE-OS constructions in Subject-referenced IC2 continuations. These constructions appeared to be used as a way to simultaneously satisfy the operative grammatical and pragmatic constraints, namely, 1) for pronouns to be subject-biased and 2) for IC2 contexts to have the preceding object as the cause of a given event.

4 General discussion

In this section, we discuss the overall pattern of results in the present study and other relevant studies concerning language universality and specificity in pronoun interpretation, comparing mainly English, Japanese, and Korean.

4.1 Factors that affect Japanese pronoun interpretation

The initial set of questions we posed concerns what factors (grammatical, pragmatic, information-structural) influence the interpretation of Japanese null and overt pronouns. First, we asked whether the interpretation of each type of pronoun is affected by both pragmatic and grammatical factors in a manner similar to what we see in English and Korean. The answer is mostly yes. As in previous studies, pragmatics was manipulated by varying context types that were known to yield different biases towards the re-mention of event participants. Overt pronouns were affected by the aspect manipulation in ToP

contexts (Experiment 1) in a manner similar to what has been found for English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008) and Korean (Kim et al. 2013). However, null pronouns did not appear to be similarly affected, in contrast to what Kim et al. (2013) found for Korean. To put the question to a stronger test, Experiment 2 manipulated context by pitting IC1 verbs against IC2 verbs. The results indicated that null pronouns are in fact sensitive to contextual biases of this sort, especially in those cases in which an Explanation relation was operative.

Second, we asked whether null and/or overt pronouns contribute a subject bias over and above the biases toward entity mention generated by the discourse context. Again consistent with both English and Korean, our results show that they do. Across both experiments, overt pronouns shifted the distribution of references toward the subject as compared to the biases measured in the free prompt condition. Similarly, null pronouns did as well, and in fact to a greater degree than overt pronouns. The results of both experiments are therefore consistent with the idea that degree of linguistic reduction (in which null > overt > name) correlates with the strength of bias toward the subject.

Finally, we asked whether topic-marking (-wa) on subject referents would attract a greater number of references than subject-marking (-ga). The answer is mostly no. There was no significant effect of topic-marking in Experiment 1, which is consistent with Kim et al.'s (2013) ToP experiment in Korean. However, there was a small effect in IC2 contexts (but not in IC1 contexts) in Experiment 2 of the present study. We tentatively hypothesized that this was because IC2 verbs create a strong lexical semantic bias that favors re-mention of the non-subject referent, which clashes with the signal sent by marking the subject as the topic. Only such contexts, therefore, is the effect of topic-marking detectable, since only here does the explicit indicator of what is being talked about contradict the direction of the discourse set up by the lexical semantics of the verb.

The picture that emerges for Japanese, as well as for English and Korean, is that different types of referring expressions contribute their own form-specific biases towards reference, but that these biases get integrated with top-down, pragmatically-driven expectations about which entities are likely to be mentioned next. The results are therefore mostly incompatible with any approach that relies primarily on grammatical 'heuristics' or 'preferences', such as a straight application of the subject preference (Crawley et al. 1990; Walker et al. 1994), the parallel grammatical role preference (Smyth 1994), or the first-mention privilege (Gernsbacher and Hargreaves 1988), since such approaches predict that the distribution of reference will be relatively uniform across different

context types.[35] Interestingly, the predictions of such approaches are in line with the pattern we saw for null pronouns in Experiment 1, whereby the distinction between perfective and imperfective aspect failed to lead to different patterns of reference. This occurred despite the fact that such differences were witnessed in the overt pronoun and free prompt conditions, hence demonstrating that the aspect manipulation did affect pragmatic biases. An effect of context on null pronouns was realized in Experiment 2, however, one that was strengthened even further when only Explanation continuations were analyzed. Whereas this result patterns more consistently with the behavior of null and overt pronouns in Korean and overt pronouns in English, it nonetheless seems that null pronouns in Japanese may be less sensitive to pragmatic biases than these other forms, in that more strongly biased contexts were needed before effects were seen. In contrast, overt Japanese pronouns were more robustly affected by the contextual manipulations in both experiments, on par with patterns previously witnessed for overt pronouns in English.

The patterns of reference found in all three languages are also inconsistent with the Expectancy Hypothesis of Arnold (2001). Recall from Section 1.1 that according to this account, biases toward pronominal reference are equated with the comprehender's expectations about who will be mentioned next. As such, no difference would be expected between the biases in the pronoun conditions and those in the free prompt condition, since the latter condition measures expectations about next-mention independently of the referential form used. Whereas such expectations are an important (and often overlooked) contributor to the interpretation biases for both null and overt pronouns, our results reveal that such influences are not the whole story.

The results instead reveal a more complex interplay between grammatical and pragmatic factors. In this regard, it is worth highlighting the unique role that the passage completion methodology has played in uncovering this pattern. Unlike other methodologies, passage completions not only reveal high-level

---

35 Historically, proponents of heuristic analyses (e. g., Crawley et al. 1990) allow for "pragmatic bias" to override the referent chosen by the heuristic, although what it means for a context to have pragmatic bias is typically left unspecified. On our analysis, the pragmatic bias is precisely the bias toward entity next-mention measured in our free-prompt conditions which, when quantified this way, makes it clear that no context is truly devoid of pragmatic bias (Kehler et al. 2008). Importantly, passage completions are a good test methodology in this regard, as participants are free to assign a pronoun prompt to a referent using a heuristic, and then complete the passage in a way that pragmatically supports that assignment. That is not what participants typically did in our experiments, however; their decisions about how to assign the pronoun were intertwined with their expectations about the direction the ensuing discourse would take.

referential behaviors, but also produce fine-grained numerical measurements of referential biases. The measured biases for different forms can then be compared against appropriate baselines, particularly those provided by the free prompt conditions. As argued by Kehler and Rohde (2013), there has historically been a problem in the literature whereby interpretation biases towards particular referents have been characterized using the assumption that a uniform distribution across available referents (e. g., a 50/50 distribution over two possible referents) is an appropriate baseline. On this assumption, for instance, a finding whereby a form yields more subject interpretations than non-subject interpretations will be used to argue that the form is associated with a subject bias. We submit, again following Kehler and Rohde (2013), that the contribution of a pronoun toward referential biases cannot appropriately be captured by analyzing only the overall reference pattern seen; it instead requires that the difference between referential biases with and without the pronoun be analyzed. That is, only by comparing reference statistics in completions with pronoun prompts with those with free prompts in otherwise identical contexts can the contribution of the pronoun itself be isolated.

This view helps to clear up confusion that might arise when analyzing data such as that we have presented here. For instance, we have claimed that overt pronouns carry a subject bias, albeit not one as strong as the one associated with null pronouns. An attentive reader might object, noting that the overt pronoun data for perfective contexts in Experiment 1 show only a 45 % bias toward the subject (Figure 2), and that the overt pronoun data for IC2 contexts in Experiment 2 show only a 43 % bias toward the subject (Figure 8). Whereas this reasoning might make sense on a heuristic approach that assumes a 50/50 distribution across the possible referents as a baseline, it does not make sense on the analysis provided here. For in each case, the percentage of subject references was greater than the corresponding next-mention biases measured in the free prompt condition (34 % and 31 % respectively). The fact that a greater percentage of subject mentions occurs with overt pronoun prompts than in free prompts demonstrates that overt pronoun prompts carry a subject bias, *even if the ultimate interpretation bias is toward a non-subject referent*. In every experiment that we are aware of that compares pronoun prompt and free prompt completions in all three languages under consideration here, the occurrence of a pronoun always contributes a subject bias, in that it pulls the distribution of referents toward the subject compared to when participants are free to select their own referring expressions in the free prompt condition. So to say that Japanese overt pronouns are only sometimes subject-biased would be an error, despite the fact that in some of our conditions the interpretation bias favored a referent other than the subject.

4.2 Different referential forms within and across languages

The second set of questions that we set out to explore concerns how different referential forms interact within and across languages. For example, we asked whether we would find a division-of-labor effect in the Japanese pronominal system of the sort predicted by an appeal to Gricean maxims. On this reasoning, the occurrence of an overt pronoun implicates that the intended interpretation is an entity other than the preferred referent for a null pronoun, for if that referent was the one intended by the speaker, she would have chosen the more informative and less prolix null form. This prediction was not borne out. While null pronouns were always more subject-biased, overt pronouns were also subject-biased, and hence their reference distribution was overlapping rather than complementary. Furthermore, both experiments showed that the pattern of results was consistent in showing no division-of-labor effect when the predetermined prompts were used and when free prompt continuations were further divided based on their matrix subject type, making it unlikely that the predetermined prompts skewed a possible division-of-labor effect.

Gundel et al.'s (1993) GH assimilates Japanese overt pronouns to demonstrative pronouns in other languages (e. g., 'that' in English), in only requiring that their referents hold the second highest cognitive status (ACTIVATED). However, the referents of all of the four Japanese overt pronouns in their corpus data (p. 291, Table 4) are actually at the highest status (IN FOCUS). Thus, Japanese overt pronouns might be better classified with Japanese null pronouns in requiring their referents to be IN FOCUS, as Gundel et al. (1993) did for null and overt pronouns in Chinese, Russian, and Spanish. However, note that the overlap between Japanese null and overt pronouns was partial in our experiments, which indicates gradient differences of a character that do not fit well within a clear-cut implicational hierarchy. Therefore, it seems more likely that different types of referring expressions independently carry their own referential biases that might crosscut with one another, as a form-specific multiple constraints approach argues (Kaiser and Trueswell 2008). In addition, recall that Korean null and overt pronouns were equally subject-biased in Kim et al.'s (2013) ToP experiment, casting doubt on the existence of a division of labor between null and overt pronouns in Korean as well. Gundel et al. (1993) did not discuss Korean, but Korean null and overt pronouns would seem to associate with the same cognitive status as they both show about an equal degree of subject bias – although an examination of how both forms of pronouns differ from the other referential forms in the language would still be required to fully evaluate the predictions of the theory.

Across languages, it seems feasible to conclude that null and overt pronouns in Japanese and Korean, as well as overt pronouns in English, do not completely align in how they are interpreted, again being consistent with a form-specific approach, crosslinguistically construed. That is, these referential forms tend to show different patterns in their sensitivity to grammatical and pragmatic factors. Compared to English overt pronouns, Japanese null pronouns are more subject-biased, while not completely insensitive to pragmatic factors. Null and overt pronouns in Korean are both more subject-biased than English and Japanese overt pronouns, while also being sensitive to pragmatic factors. Furthermore, although Japanese overt pronouns patterned quite closely with English overt pronouns in terms of their sensitivity to grammatical and pragmatic factors in the present study, they still cannot take a bound variable interpretation as English overt pronouns can (cf. Hoji 1990). Thus they are not equivalent to each other in a number of respects. As such, there seems to be a quite complex and nuanced picture at work for pronoun interpretation across languages, as we speculated in the introduction.

This view can help explain why typologically similar null and overt pronouns in Korean do not behave exactly like those in Japanese. For instance, as mentioned above, Korean overt pronouns are equally subject-biased as Korean null pronouns and much more subject-biased than Japanese overt pronouns. Hoji (1990) argued that the reason why Korean *ku* 'he' can be construed as a bound variable while Japanese *kare* 'he' cannot is because Korean *ku* is less deictic/demonstrative (and thus more pronominal) than Japanese *kare*, in that the former belongs to the medial *ku* (Korean)/*so* (Japanese) demonstrative systems that are not as strongly deictic as the distal *ce/a* demonstrative systems the latter belongs to. This may be why Korean overt pronouns are more sensitive to a grammatical factor (the subject position in this case) than Japanese overt pronouns. Furthermore, while strongly subject-biased, Korean null (in addition to overt) pronouns are also sensitive to pragmatic factors even in ToP contexts. Recall from Section 1.2 that Kwon and Sturt (2013) showed that the interpretation of Korean null pronouns is more dependent on discourse cues than morphosyntactic cues. In light of the idea that multiple different constraints affect the interpretation of referential forms in varying degrees and that individual languages differ in the relative weight they assign to these different factors, it may be the case that Korean puts more weight on pragmatics/discourse-oriented factors in the interpretation of null pronouns than Japanese.[36]

---

36 Kwon and Sturt (2013) predicted that all null pronoun languages without rich verbal agreement, such as Japanese and Chinese, should be all discourse-oriented, but no study to our best knowledge has tested such a claim. Additionally, Japanese can also be discourse-

4.3 Coherence relations

Finally, we asked whether the various factors that were shown to affect Japanese pronoun biases also affect the distribution of coherence relations found in participants' continuations. An effect is expected on a model in which the relationship between coherence and reference is bidirectional: The use of pronominal forms that overlay a subject bias on top of next-mention expectations are predicted to lead to more subject-biased coherence relations, as previously shown by Rohde and colleagues for English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008). This prediction was confirmed. For instance, the aspect manipulation in the ToP contexts used in Experiment 1 changed both reference and the proportion of Source- and Goal-biased coherence relations, as also seen in English (Rohde et al. 2006; Rohde and Kehler 2008; Rohde 2008) and Korean (Kim et al. 2013). The IC contexts used in Experiment 2 also influenced both reference and coherence relations in the predicted direction, also as previously shown for English (Rohde et al. 2006; Rohde and Kehler 2008; and Rohde 2008). There were some differences between English and Japanese in how reference and coherence relations interact, but they are largely explainable; for instance, when pronoun biases and coherence biases clashed, Japanese participants sometimes used specific grammatical forms provided by the language that, by mentioning the causally-implicated referent in a position other than the matrix subject, allowed them to satisfy both biases at once. This study has therefore added crosslinguistic evidence to the claim that the dependency between coherence and reference is bidirectional.

Again, it is worth highlighting the crucial role that the passage completion methodology plays in examining these questions. Despite the fact that no fully predictive theory of coherence establishment exists, we can nonetheless analyze completions to estimate interlocutors' coherence-driven biases. These can then be used alongside the measured biases for reference to test predictions about the interaction between the two. Note that our results are inconsistent with classic analyses like that of Hobbs (1979, 1990) in which, as described in Section 1.1, pronouns are modeled as free variables that get bound to referents during the course of coherence establishment. Whereas Hobbs' (1979, 1990) analysis predicts that coherence relations should affect pronoun biases (which we have

oriented to a certain degree, but it may not place as much priority on discourse/pragmatic cues as Korean does, and as a result Korean rather than Japanese null pronouns may be relatively more sensitive to different types of pragmatic factors.

shown to be the case), it also predicts that it should not go the other way, since the contribution of a pronoun should be inert. Our results demonstrate that Japanese pronouns are not inert, but instead contribute their own referential biases that in turn cascade to influence coherence biases.

5 Conclusions

This study investigated the interpretation of null and overt pronouns in Japanese in ToP and IC contexts. Japanese overt pronouns were revealed to pattern closely with English overt pronouns in their sensitivity to pragmatic factors, whereas Japanese null pronouns showed a mixed resilience to pragmatic factors. Topic-marking showed marginal effects on reference only in limited contexts. Despite different degrees of sensitivity to pragmatic factors, Japanese null and overt pronouns were both mostly subject-biased, casting doubt on the existence of a division of labor between the two forms. There was also an intrinsic link between reference and coherence relations throughout the experiments. We have discussed the overall pattern of results in terms of language-specificity and universality, the latter of which includes interactions between grammatical and pragmatic factors and the importance of discourse coherence in the interpretation of various pronouns across languages.

These experiments compared reference behavior in a null pronoun language (Japanese) with an overt pronoun language (English) and another null pronoun language (Korean). As Japanese and Korean have no verbal agreement marking, a remaining question is how these reference patterns compare with null pronoun languages with rich verbal agreement (so-called 'pro-drop' languages) such as Spanish and Italian. A recent study by Filiaci et al. (2014), for instance, compared the processing of null and overt pronouns in these two languages. By way of two reading time studies, their results revealed that the processing of null subjects in both languages incurred a penalty when forced to refer to a non-subject antecedent. However, there was a difference between the two languages with respect to overt pronouns: Whereas Italian overt pronouns incurred a penalty when referring to a subject (suggesting a division-of-labor effect), Spanish overt pronouns did not. As such, their results reinforce the idea that the referential systems of languages, again even ones that are typologically close, do not always behave in a way that can be captured by a universal hierarchy. In line with our arguments above, it would be useful to carry out passage completion studies across the two languages to better understand the specific biases that each form contributes over and above the biases associated with next-mention.

References

Alonso-Ovalle, Luis, Susana Fernández-Solera, Lyn Frazier & Charles Clifton. 2002. Null vs. overt pronouns and the topic-focus articulation in Spanish. *Rivista di Linguistica* 14(2). 1–19.

Amano, Nariaki & Masahisa Kondo. 2000. *NTT database series nihongo-no goikokusei* [NTT database series lexical properties of Japanese], vol. 7. Tokyo: Sanseido.

Arnold, Jennifer E. 2001. The effects of thematic roles on pronoun use and frequency of reference. *Discourse Processes* 31. 137–162.

Bentivoglio, Paola. 1992. Linguistic correlations between subjects of one-argument verbs and subjects of more-than-one-argument verbs in spoken Spanish. In Paul Hirschbühler & Konrad Koerner (eds.), *Romance languages and modern linguistic theory: 20th Linguistic Symposium on Romance Languages* (LSRL XX), Ottawa, 10–14 April 1990, 11–24. Amsterdam & Philadelphia: John Benjamins.

Brown-Schmidt, Sarah, Donna. K. Byron & Michael Tanenhaus. 2005. Beyond salience: Interpretation of personal and demonstrative pronouns. *Journal of Memory and Language* 53. 292–313.

Caramazza, Alfonso, Ellen Grober & Catherine Garvey. 1977. Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behavior* 16. 601–609.

Carminati, Maria N. 2002. *The processing of Italian subject pronouns*. Amherst, MA: University of Massachusetts at Amherst dissertation.

Christianson, Kiel & Hee Youn Cho. 2009. Interpreting null pronouns (*pro*) in isolated sentences. *Lingua* 119. 989–1008.

Crawley, Rosalind A., Rosemary J. Stevenson & David Kleinman. 1990. The use of heuristic strategies in the interpretation of pronoun. *Journal of Psycholinguistic Research* 4. 245–264.

Dowty, David. 1991. Thematic proto roles and argument selection. *Language* 67. 547–619.

Du Bois, John W. 1987. The discourse basis of ergativity. *Language* 63. 805–855.

Fernandez-Salgueiro, Gerardo, Lillian Chen, Richard Lewis & Samuel Epstein. 2007. The when and how of positing null subjects (*pro*): The timing and surprising interaction of syntactic and semantic constraints. Poster presented at the 20th Annual CUNY Conference on Human Sentence Processing. San Diego, CA.

Ferstl, Evelyn C., Alan Garnham & Christina Manouilidou. 2011. Implicit causality bias in English: A corpus of 300 verbs. *Behavior Research Methods* 43. 124–135.

Fiengo, Robert & Makiko Haruna. 1987. Parameters in binding theory–Some suggestions based on an analysis of Japanese. In Takashi Imai & Mamoru Saito (eds.), *Issues in Japanese linguistics*, 107–128. Dordrecht: Foris Publication.

Filiaci, Francesca, Antonella Sorace & Manuel Carreiras. 2014. Anaphoric biases of null and overt subjects in Italian and Spanish: A cross-linguistic comparison. *Language, Cognition, and Neuroscience* 29. 825–843.

Fry, John. 2003. *Ellipsis and wa-marking in Japanese conversation*. New York, NY: Routledge.

Garvey, Catherine & Alfonso Caramazza. 1974. Implicit causality in verbs. *Linguistic Inquiry* 5. 459–464.

Gernsbacher, Morton A. & David J. Hargreaves. 1988. Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language* 27. 699–717.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Syntax and semantics 3: Speech acts*, 41–58, New York, NY: Academic Press.

Gundel, Jeanette K., Nancy Hedberg & Ron Zacharski. 1993. Cognitive status and the form of referring expressions in discourse. *Language* 69. 274–307.

Haegeman, Liliane. 1990. Understood subjects in English diaries: On the relevance of theoretical syntax for the study of register variation. *Multilingua* 9. 157–199.

Hobbs, Jerry R. 1979. Coherence and coreference. *Cognitive Science* 3. 67–90.

Hobbs, Jerry R. 1990. *Literature and cognition: CSLI lecture notes 21*. Stanford, CA: CSLI Publications.

Hoji, Hajime. 1990. On the so-called overt pronouns in Japanese and Korean. In Eung-Jin Baek (ed.), *Papers from the seventh international conference on Korean linguistics*, 61–78. Osaka, Japan: International Circle of Korean Linguistics & Osaka University of Economics and Law.

Ishiyama, Osamu. 2008. *Diachronic perspectives on personal pronouns in Japanese*. Buffalo, NY: State University of New York at Buffalo doctoral dissertation.

Kaiser, Elsi & John Trueswell. 2008. Interpreting pronouns and demonstratives in Finnish: Evidence for a form-specific approach to reference resolution. *Language and Cognitive Processes* 23. 708–748.

Kameyama, Megumi. 1985. *Zero anaphora: The case of Japanese*. Stanford, CA: Stanford University doctoral dissertation.

Kehler, Andrew. 2002. *Coherence, reference, and the theory of grammar*. Stanford, CA: CSLI Publications.

Kehler, Andrew, Laura Kertz, Hannah Rohde & Jeff Elman. 2008. Coherence and coreference revisited. *Journal of Semantics* 25. 1–44.

Kehler, Andrew & Hannah Rohde. 2013. A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics* 39. 1–37.

Kim, Kitaek, Theres Grüter & Amy J. Schafer. 2014. Effects of morphological and prosodic focus cues on topic maintenance in Korean. Poster presented at the 27th Annual CUNY Conference on Human Sentence Processing, Columbus, OH.

Kim, Lucy Kyoungsook & Elsi Kaiser. 2009. Effects of honorific agreement on null subject interpretation in Korean. In *Proceedings of 2009 Seoul International Conference on Linguistic Interfaces*. Seoul: Yonsei University.

Kim, Kitaek, Grüter Theres & Amy J. Schafer. 2013. Effects of event-structure and topic/focus-marking on pronoun reference in Korean. Poster presented at the 26th Annual CUNY Conference on Human Sentence Processing, Columbia, SC.

Kuno, Susumu. 1973. *The structure of the Japanese language*. Cambridge, MA: MIT Press.

Kuroda, Sige-Yuki. 1965. *Generative grammatical studies in the Japanese language*. Cambridge, MA: MIT doctoral dissertation.

Kwon, Nayoung & Maria Polinsky. 2011. Anaphoric inventories and bound variable interpretation: Evidence from Korean. Poster presented at the 24th Annual CUNY Conference on Human Sentence Processing, Stanford University, CA.

Kwon, Nayoung & Patrick Sturt. 2013. Null pronominal (*pro*) resolution in Korean, a discourse-oriented language. *Language and Cognitive Processes* 28. 377–387.

Liversedge, Simon P. & Roger P. G. van Gompel. n.d. *Resolving anaphoric and cataphoric pronouns*. Unpublished ms.

Martin, Samuel E. 1975. *A reference grammar of Japanese*. New Haven, CT: Yale University Press.

Mazuka, Reiko. 1991. Processing of empty categories in Japanese. *Journal of Psycholinguistic Research* 20. 215–232.

McKoon, Gail, Steven B. Greene & Roger Raticliff. 1993. Discourse models, pronoun resolution, and implicit causality of verbs. *Journal of Experimental Psychology* 19. 1040–1052.

Nariyama, Shigeko. 2003. *Ellipsis and reference tracking in Japanese*. Amsterdam & Philadelphia: John Benjamins.

Rizzi, Luigi. 1994. Root null subjects and early null subjects. In Teun Hoekstra & Bonnie Schwartz (eds.), *Language acquisition studies in generative grammar*, 151–176. Amsterdam & Philadelphia: John Benjamins.

Rohde, Hannah. 2008. *Coherence driven effects in sentence and discourse processing*. San Diego, CA: University of California, San Diego doctoral dissertation.

Rohde, Hannah & Andrew Kehler. 2008. The bidirectional influence between coherence establishment and pronoun interpretation. Poster presented at the 21st Annual CUNY Conference on Human Sentence Processing, Chapel Hill, NC.

Rohde, Hannah, Andrew Kehler & Jeff Elman. 2006. Event structure and discourse coherence biases in pronoun interpretation. In Ron Sun & Noami Miyake (eds.), *Proceedings of the 28th annual conference of the cognitive science society*, 697–702. London: Psychology Press.

Smyth, Ron. 1994. Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research* 23. 197–229.

Stevenson, Rosemary J., Rosalind A. Crawley & David Kleinman. 1994. Thematic roles, focusing and the representation of events. *Language and Cognitive Processes* 9. 519–548.

Tao, Hongyin. 1996. *Units in Mandarin conversation: Prosody, discourse, and grammar*. Amsterdam: John Benjamins.

Ueno, Mieko & Andrew Kehler. 2010. The interpretation of null and overt pronouns in Japanese: Grammatical and pragmatic factors. In Stellan Ohlsson & Richard Catrambone (eds.), *Proceedings of the 32nd annual meeting of the cognitive science society*, 2057–2062. Austin, TX: Cognitive Science Society.

Ueno, Mieko & Maria Polinsky. 2009. Does headedness affect processing? A new look at the VO-OV contrast. *Journal of Linguistics* 45. 675–710.

Walker, Marilyn, Masayo Iida & Sharon Cote. 1994. Japanese discourse and the process of centering. *Computational Linguistics* 20. 193–232.

Yang, Chin Lung, Peter C. Gordon, Randall Hendrick & Chih Wei Hue. 2003. Constraining the comprehension of pronominal expressions in Chinese. *Cognition* 86. 283–315.

Yang, Chin Lung, Peter C. Gordon, Randall Hendrick & Jei Tun Wu. 1999. Comprehension of referring expressions in Chinese. *Language and Cognitive Processes* 14. 715–743.