

Interpreting Japanese Head-Internal Relative Clauses

Michael Yoshitaka Erlewine and Isaac Gould
Massachusetts Institute of Technology

1. Introduction

Japanese has both *head-external* and *head-internal* relative clauses (HERC & HIRC). Example (1a) shows a standard HERC. In contrast, the HIRC in (1b) has the form of an independent sentence, the bold-faced head occurring inside the relative, and it is followed by the complementizer *-no* and a case morpheme that relates the internal head to the matrix verb. A core question is what the interpretation of HIRCs is. How is their meaning computed compositionally, and how does this account for their differences with HERCs?¹

- (1) a. **Head-external relative clause (HERC; Shimoyama, 1999):**
Yoko-wa [[_{HERC} Taro-ga sara-no-ue-ni oita] **keeki**]-o tabeta
Yoko-TOP Taro-NOM plate-GEN-on-LOC put cake-ACC ate
'Yoko ate the cake that Taro put on a plate.'
- b. **Head-internal relative clause (HIRC; Shimoyama, 1999):**
Yoko-wa [_{HIRC} [Taro-ga sara-no-ue-ni **keeki-o** oita] -no]-o tabeta
Yoko-TOP Taro-NOM plate-GEN-on-LOC cake-ACC put -NO-ACC ate
Literally 'Yoko ate [Taro put **cake** on a plate].'

The contribution of this paper is to propose a new kind of semantic analysis for Japanese HIRCs, which is based on new empirical observations concerning the interpretation of HIRCs with quantificational heads. These findings are presented in Section 2. The analysis we propose in Section 3 involves movement of a null operator that is attached to the relative's head noun. After the null operator moves, a *maximally informative* definite operator is attached. We also present an alternative syntactic formulation using the Copy Theory of movement in Section 3.4.

Our analysis builds on earlier syntactic analyses of Itô (1986) and Watanabe (1992). Both analyses propose a movement relation for the HIRC's head: for Itô, the entire head moves covertly, and for Watanabe a quantificational null operator moves from the head's specifier. These approaches to HIRCs parallel HERCs in that the head is involved in a movement relation. However, as Hoshi (1995) and Shimoyama (1999) note, these analyses incorrectly predict the HIRC's head to be able to scope outside the relative like HERCs. In fact, unlike with HERCs, HIRC heads scope only within the relative. Our analysis is crucially different in that what moves out of the head is only non-quantificational material – specifically, a restriction on the head's domain of quantification – thus correctly predicting scope of the quantifier within the relative clause. We conclude that HIRCs are syntactically more similar to HERCs than has been assumed in recent analyses.

We provide a brief comparison of other semantic analyses of Japanese HIRCs in Section 4 and note that they have difficulty accounting for the new data discussed in Section 2. We conclude in Section 5 with some challenges and questions for our approach.

* We thank audiences at the Tokyo Semantics Reading Group, MIT LF Reading Group, LSA Annual Meeting 87, and WCCFL 31 for helpful comments and judgments. We especially thank Chris Davis, Danny Fox, Min-Joo Kim, Hadas Kotek, Junya Nomura, Yasutada Sudo, Ayaka Sugawara, and Satoshi Tomioka. Errors are ours.

¹ An important difference between HERCs and HIRCs, but one that we will not discuss, is that HIRCs (but not HERCs) are sensitive to what has been called the Relevancy Condition, which seems to place semantic-pragmatic constraints on what can be an HIRC head vis-à-vis properties of the surrounding relative clause and the matrix clause. See Kuroda (1976-77) for an early formulation of this condition.

2. New observations on the interpretation of HIRC

Our analysis is motivated by new empirical observations concerning the interpretation of HIRCs that have quantified heads. In this section, we discuss HIRCs that involve proportional and numeral quantifiers. We show a novel contrast in the interpretation of these HIRCs: for the speakers we consulted, a *domain reading* is possible, but a *witness set reading* is not. This contrast is crucially sensitive to the context, and we show how manipulations of the context result in different judgments.

2.1. Domain readings

In order to sharpen intuitions about the meanings of HIRCs, we did two things. First, we added the floating quantifier *zenbu* ‘all’ immediately following the HIRC (3). Second, we provided explicit contexts with our sentences. We will begin with the discussion of the context in (2), where there are two salient groups of six apples each: in the first group, half of the apples are peeled (the white apples); in the second group, all apples are unpeeled.

(2) **Context with two groups of apples:**



(3) **Head-internal relative clause with head quantifier *hanbun* ‘half’:**

Junya-wa [_{HIRC} [Ayaka-ga **hanbun-no-ringo**-o muita]-no]-o zenbu taberu
 Junya-TOP Ayaka-NOM half-GEN-apple-ACC peeled-NO-ACC all eat
 Literally ‘Junya will eat all of [Ayaka peeled **half of the apples**].’

Judgments in context (2):

- a. ⇒ Junya will eat all six of the apples in the first group. *domain reading*
- b. ⇏ Junya will eat the three peeled apples. *witness set reading*

There are two potential readings of (2) that we now consider. The first is one in which Junya eats all of the apples in the first group—the group that Ayaka peeled half of. We call this reading the *domain reading* (3a). The second reading is one where Junya eats all three of the peeled apples, but no other apples. This is the set of apples that constitute the half that Ayaka peeled, and we call this reading the *witness set reading* (3b). Our consultants report that only the domain reading (3a) holds in this context (2); the witness set reading (3b) is not a felicitous interpretation of (3).

We see similar judgments with numeral quantifiers. This is shown with *mittsu* ‘three’ in (4) below, with judgments based on the same context in (2) above. Our consultants again judge the domain reading (4a) to be possible and report that the witness set reading (4b) is unavailable. We conclude, then, that HIRCs denote the relevant *domain* of the head noun’s quantifier, not the witness set.

(4) **Head-internal relative clause with head quantifier *mittsu* ‘three’:**

Junya-wa [_{HIRC} [Ayaka-ga **ringo-o mit-tsu** muita]-no]-o zenbu taberu
 Junya-TOP Ayaka-NOM apple-ACC three-CL peeled-NO-ACC all eat
 Literally ‘Junya will eat all of [Ayaka peeled **three of the apples**].’

Judgments in context (2):

- a. ⇒ Junya will eat all six of the apples in the first group. *domain reading*
- b. ⇏ Junya will eat the three peeled apples. *witness set reading*

2.2. The effect of context

In the previous subsection we reported that Japanese HIRCs have a domain reading when interpreted with a context with a salient partition. In this section, we show that manipulating the context affects the interpretation of HIRCs. We illustrate this with the same quantified HIRCs from the previous subsection, but crucially the context that these HIRCs are to be evaluated against has no salient partition. This is illustrated graphically below: we present a context in which there are a total of twelve apples, three of them peeled, but without a salient partition.

(5) **Context with no salient subgroups:**



Let us now return to the examples in (3) and (4) in light of the context in **Error! Reference source not found.** In such a context, it is not possible to utter the HIRC in (3) with the head quantifier ‘half’: the sentence is simply infelicitous. Things are different for the HIRC in (4) with the head quantified by ‘three.’ When uttering (4) in the context of **Error! Reference source not found.**, the HIRC denotes the three peeled apples. Thus Junya eats only the three peeled apples, and the witness set reading holds. These judgments are different from the readings reported above where Junya ends up eating both peeled and non-peeled apples in the context in (2). The judgments are summarized in the table below. We note that there is some speaker variation with respect to the witness set reading and comment on this briefly in the conclusion.

(6) **Summary of interpretations of HIRC in different contexts:**

Context:		
HIRC with ‘half’ (2):		#
HIRC with ‘three’ (4):		

Thus, different quantifiers react differently to the (un)availability of contextually salient groups. The domain readings are possible when we have a contextually salient partition. In a context without such a partition, however, we get an entirely different set of judgments ranging from infelicity to the witness set reading. The task, then, is to provide an analysis that correctly models the contrast between domain and witness set readings and that is also sensitive to varying the context, modeling the resulting change in interpretation. We provide such an analysis in the following section.

3. Analysis

In the next section, we will introduce our syntactic analysis, sketch the proposed semantics, and describe some immediate consequences of the analysis. We will then introduce the notion of definiteness as *maximal informativeness*, following the work of von Stechow, Fox, and Iatridou (2012). We will then return to the semantics of HIRCs, formalizing our assumptions about the role of context and demonstrating how our compositional semantics yields the correct interpretation of HIRCs.

3.1. Proposal

We propose that HIRCs are definite descriptions, corresponding to the restriction of the HIRC’s head noun quantifier. Syntactically, this is accomplished by (a) the adjunction of a null operator of set type to the quantifier’s restriction, (b) movement of the null operator to the edge of the HIRC, and (c) the addition of a null definite head. Semantically, the extension of the HIRC will correspond to the *domain* of the head noun’s quantifier, rather than its witness set, as desired.

We illustrate this derivation using the HIRC in (2). The head noun of this HIRC is the quantifier “half apples.” A null operator of set type (*op*) is adjoined to the head’s NP, *apples*, and the modified structure is interpreted using Predicate Modification (7a). This operator moves to the edge of the HIRC, to yield a predicate by Predicate Abstracting over the variable *X* in the trace position of the operator (7b).² A null definite head (THE) is merged, yielding the definite description in (7c).

² See Heim & Kratzer (1998) for definitions of the terms *Predicate Modification* and *Predicate Abstraction*.

- (7) **Syntactic derivation for the HIRC (3):**
- a. Modification of the head noun’s restriction by a null operator:
Ayaka peeled half [apples in *op*]
 - b. Movement of the null operator to the edge of the HIRC:
 $\lambda X . \text{Ayaka peeled half [apples in } X]$
 \uparrow
 \downarrow
 - c. Merge null definite head THE:
THE ($\lambda X . \text{Ayaka peeled half [apples in } X]$)

Intuitively, the HIRC in (3) is best translated into English as “the apples that Ayaka peeled half of.” As this English translation shows, the extension of the HIRC corresponds to a salient group of apples, half of which are peeled, rather than those apples which are peeled. Details of the interpretation of (7c) are discussed shortly. First, we highlight some immediate consequences of this novel analysis.

Our proposal captures the intuition of Itô (1986) and Watanabe (1992) that the head noun in a HIRC is involved in a movement relation parallel to the movement of the head noun in a HERC. Under our analysis, the head noun of a HIRC is involved in a movement relation, but it is only *a part* of the head noun that moves, rather than the entire head noun.

The null operator movement proposed here predicts that HIRCs should be sensitive to syntactic islands. This is observed in Watanabe (1992, 2003). Example (8)(8) shows that the head noun in a HIRC can be within an embedded clause; but (9) shows that when the head noun is embedded within a relative clause island, the HIRC becomes ungrammatical.

- (8) **Grammatical HIRC with head embedded in a complement clause (Watanabe, 2003):**
 Mary-ga [HIRC [John-ga [CP jibun-no gakusei-ga **juuyouna kasetu**-o teianshita-to]
 Mary-NOM John-NOM self-GEN student-NOM important hypothesis-ACC proposed-that
 jimanshiteita -no]-no kekkan-o shitekishita
 boasted -NO-GEN defect-ACC pointed.out
 Literally ‘Mary pointed out a defect of [John boasted that [his student proposed an **important hypothesis**]].’
- (9) **Ungrammatical HIRC with head embedded in a relative clause island:**³
 * Mary-ga [HIRC [John-ga [_{island} [**atarashii kasetu**-o teianshita] gakusei]-o hometeita]
 Mary-NOM John-NOM new hypothesis-ACC proposed student-ACC praised
 -no]-no kekkan-o shitekishita
 -NO-GEN defect-ACC pointed.out
 Literally ‘Mary pointed out a defect of [John praised the student_i [that *t_i* proposed a **new hypothesis**]].’

Our analysis also predicts the correct scope for the quantifier on the head noun. Shimoyama (1999) shows that quantifiers on HERC heads scope over their embedding clause (10), but quantifiers on the head of HIRCs scope within the relative clause (11). This is predicted by our analysis: the quantifier of the head noun (‘half’ in example (3) above) stays in the HIRC and therefore takes scope only over the material in the HIRC.

- (10) **Quantifier on head of HERC scopes over the embedding clause (Shimoyama 1999):**
 Taro-wa [[Yoko-ga reezooko-ni ireteoita] **kukkii**]-o **hotondo** paatii-ni motteitta
 Taro-TOP Yoko-NOM refrigerator-LOC put cookie-ACC almost all party-to brought
 ‘Taro brought **almost all of [the cookies]** that Yoko had put in the refrigerator] to the party.’
 ⇒ ALMOSTALL(cookie $\cap \lambda x . \text{Yoko put } x \text{ in the fridge})(\lambda x . \text{Taro brought } x \text{ to the party})$
- (11) **Quantifier on head of HIRC scopes within the relative clause (Shimoyama 1999):**
 Taro-wa [[Yoko-ga reezooko-ni **kukkii**-o **hotondo** ireteoita]-no]-o paatii-ni motteitta
 Taro-TOP Yoko-NOM refrigerator-LOC cookie-ACC almost all put -NO-ACC party-to brought
 Literally ‘Taro brought [Yoko put **almost all of the cookies]** in the fridge] to the party.’
 ⇒ ALMOSTALL(cookie)($\lambda x . \text{Yoko put } x \text{ in the fridge}$)

³ This example comes from Grosu (2010), citing Akira Watanabe (p.c.). This point was first made in Watanabe (1992, 2003) with examples less parallel to (8).

Two additional ingredients are necessary in order to properly compute the denotation of a HIRC under our analysis. First is the interpretation of definiteness as maximal informativeness, as proposed by von Fintel, Fox, and Iatridou (2012). Second is an assumption regarding the role of context in the interpretation of referential expressions.

3.2. *Definiteness as maximal informativeness (von Fintel, Fox, and Iatridou 2012)*

We have proposed that Japanese HIRCs are interpreted as definite descriptions, using a null definite determiner. In this paper we will adopt a proposal by von Fintel, Fox, and Iatridou (2012) for modeling definiteness as *maximal informativeness*. As we will see in the next section, the adoption of a maximal informativeness semantics for the definite determiner is crucial for computing the correct interpretation for some HIRCs.

Traditionally the definite determiner THE has been described as a *maximality* operator which returns the maximal individual satisfying the restriction. Consider the interpretation of the definite description *the number of children that John has* in (12) below. We assume that the domain of the predicate $\varphi = \text{“}\lambda n. \text{John has } n\text{-many children”}$ is restricted to positive integers. The definite description is interpreted as the maximal value that satisfies this predicate in the interpretation world. In this case, the positive integers 1, 2, 3, and 4 satisfy the predicate φ . $\llbracket \text{THE} \rrbracket(\varphi)$ returns the maximal value, 4.

(12) **Maximality semantics for *the number of children that John has*:**

Context: John has exactly four children.

1 < 2 < 3 < **4** < ~~5~~ < ~~6~~ < ~~7~~ < ...

von Fintel, Fox, and Iatridou (2012) proposes a new, intensional semantics for THE which returns a *maximally informative* individual. The denotation of the definite determiner THE is defined as follows:

(13) **Definiteness as maximal informativeness (von Fintel, Fox, and Iatridou 2012):**

$\llbracket \text{THE} \rrbracket(\varphi_{\langle \alpha, t \rangle})$ is defined only if there is an object x such that $\varphi(x)$ is the maximally informative proposition among the true propositions of the form $\varphi(\dots)$. When defined, $\llbracket \text{THE} \rrbracket(\varphi)$ refers to that object x .

In many cases, maximal informativeness yields the same result as the traditional maximality semantics. For example, consider the interpretation of *the number of children that John has* using the maximal informativeness semantics for THE in (13). In this case we consider the intensional propositions of the form “ $\lambda w. \text{John has } n\text{-many children in } w$ ” with different values of w . Those φ -propositions that are not true in the real world are removed from consideration. The relation of entailment induces a partial ordering over propositions, where we define $p < q$ if and only if q is a stronger proposition than p : $p \Leftarrow q$. The φ -proposition “ $\lambda w. \text{John has 4 children in } w$ ” entails all other φ -propositions which are true in the context, and therefore it will be the *maximally informative* φ -proposition. $\llbracket \text{THE} \rrbracket(\varphi)$ returns the value corresponding to this maximally informative proposition: 4.

(14) **Maximal informativeness semantics for *the number of children that John has*:**

Context: John has exactly four children.

... < $\lambda w. \text{John has 3 children in } w$
 < **$\lambda w. \text{John has 4 children in } w$**
 < ~~$\lambda w. \text{John has 5 children in } w$~~ < ...

The difference between maximality and maximal informativeness lies in the ordering used over the values that satisfy the restriction. In the traditional maximality semantics, the natural ordering over numbers is used to choose the maximal element. In contrast, under the maximal informativeness proposal, values satisfying the restriction φ are ordered based on entailment between their corresponding φ -propositions. In the case of *the number of children that John has*, both approaches to the definite article yield the same results.

There are other cases, however, where the traditional maximality semantics cannot be used, but maximal informativeness yields the correct interpretation.⁴ For example, consider the definite description *the amount of flour sufficient to bake a cake*, beginning with the maximal informativeness interpretation in (13). We consider intensional propositions of the form “ $\lambda w. n$ -much flour is sufficient to bake a cake in w ” with different values of n , ordered by propositional strength. Note that if the proposition “ $\lambda w. n$ -much flour is sufficient to bake a cake in w ” is true, the corresponding proposition for any greater amount of flour, $m > n$, will also be true: “ $\lambda w. m$ -much flour is sufficient to bake a cake in w .” Therefore, φ -propositions corresponding to *smaller* values of n will be stronger. Of those φ -propositions which are true in the context, the maximally informative proposition is “ $\lambda w. 150\text{g}$ of flour is sufficient to bake a cake in w .” $\llbracket \text{THE} \rrbracket(\varphi)$ returns the corresponding value: 150g.

(15) **Maximal informativeness semantics for *the amount of flour sufficient to bake a cake*:**

Context: 150g of flour is sufficient to bake a cake.

... $< \lambda w. 160\text{g}$ of flour is sufficient to bake a cake in w

$< \lambda w. 150\text{g}$ of flour is sufficient to bake a cake in w

$< \lambda w. 140\text{g}$ of flour is sufficient to bake a cake in $w < \dots$

(16) **Maximality semantics for *the amount of flour sufficient to bake a cake*:**

Context: 150g of flour is sufficient to bake a cake.

... $< 130\text{g} < 140\text{g} < 150\text{g} < 160\text{g} < 170\text{g} < \dots$

there is no maximal value!

In contrast, we are unable to model this definite description accurately using the traditional maximality semantics. In the context, 150g of flour is sufficient to bake a cake, but so is *any greater amount of flour*. Using the natural ordering over these values, there will be no unique maximal value.

In our analysis of Japanese HIRC, we will be adopting von Stechow, Fox, and Iatridou’s maximal informativeness semantics for the definite determiner THE. One consequence will be that the denotation of a HIRC headed by ‘apple,’ for example, will return only a set of apples.

Recall that under our proposal, the interpretation of the HIRC with the head quantifier *hanbun* ‘half’ (3) will be $\llbracket \text{THE} \rrbracket(\varphi)$, where $\varphi = “\lambda X. \text{Ayaka peeled half [apples in X].”$ Notice that the predicate φ can be true of sets that contain both apples and non-apples.⁵ For example, $\varphi(\{\bullet, \circ, x\})$, where x is not an apple, will be true, because “the apples in $\{\bullet, \circ, x\}$ ” is the set $\{\bullet, \circ\}$, and half of $\{\bullet, \circ\}$ are indeed peeled (we assume by Ayaka). However, the denotation of our HIRCs headed by ‘apple’ is always a set of apples, not some apples together with other entities.⁶

Our adoption of a maximal informativeness semantics for the definite determiner resolves this potential issue. Under this semantics for the definite determiner, our goal is to identify the logically strongest intensional φ -proposition. Note that the intensional proposition “ $\lambda w. \text{Ayaka peeled half of the apples in } \{\bullet, \circ\}$ ” entails “ $\lambda w. \text{Ayaka peeled half of the apples in } \{\bullet, \circ, x\}$ ”. In fact, in general for any set of entities X , the intensional φ -proposition generated by the apple-part of X will entail the φ -proposition generated by X : $\varphi(X \cap \llbracket \text{apple} \rrbracket^{w*}) \Rightarrow \varphi(X)$.⁷ Therefore the result of $\llbracket \text{THE} \rrbracket(\varphi)$ must necessarily be a set of apples. This argument extends to any predicate that is the explicit domain of a HIRC’s head noun: if NP is the overt restriction of the head of a HIRC, the denotation of the HIRC will necessarily be a subset of $\llbracket \text{NP} \rrbracket$.

With this background and motivation in place, we now return to the compositional semantic interpretation of Japanese HIRC.

⁴ See von Stechow et al (2012) for discussion of the formal properties of those cases where the maximality and maximal informativeness semantics diverge.

⁵ Here we will use *set* terminology for the representation of plural individuals.

⁶ It has been claimed that the HIRC can sometimes refer to an entity that could not be ascribed the property denoted by the HIRC’s head NP. See Hoshi (1995: §3.2) for examples. One example is an HIRC whose head NP is ‘orange’, but which appears to refer to the juice resulting from squeezing said orange. The set of entities that can be ascribed the property ‘orange juice’ is not necessarily the same as that which can be ascribed the property ‘orange’. Nevertheless, it seems to us that some modification of the relation between the null operator and the head NP could account for these examples. We leave this as an issue for future research.

⁷ $\llbracket \text{apple} \rrbracket^{w*}$ is the set of apples *de re*—that is, the set of entities which are apples in the real world. We thank Danny Fox (p.c.) for pointing out this important technical detail.

3.3. Interpreting HIRCs and modeling the effect of context


As we have seen in section 2, the felicity and interpretation of HIRCs are highly context-dependent. In order to accurately model the interpretations of HIRCs and their context-sensitive behavior, we propose the following:

(17) **Salient Sets Restriction:**

The existence of salient sets in the context requires that only those salient sets be considered as candidates for the output of $\llbracket \text{THE} \rrbracket$.

We will begin with the interpretation of the HIRC with the head quantifier *hanbun* ‘half’ (3) in our context with two salient groups of apples (2), repeated below as (18) with apples numbered. Following our syntactic proposal, the interpretation of the HIRC will be $\llbracket \text{THE} \rrbracket(\varphi)$, where $\varphi = “\lambda X. \text{Ayaka peeled half [apples in X].}”$ In this context the Salient Sets Restriction will restrict our consideration to just two sets of apples: $A = \{1, 2, 3, 4, 5, 6\}$ and $B = \{7, 8, 9, 10, 11, 12\}$. A satisfies φ (half of the apples in A have been peeled by Ayaka) but B does not, so $\llbracket \text{THE} \rrbracket(\varphi) = A$, the six apples in the first group.


(18) **Context with two groups of apples:**



This same procedure yields the same extension for the HIRC with the head quantifier *mittsu* ‘three’ (4) in this context (18). In this case, we compute $\llbracket \text{THE} \rrbracket(\varphi)$, where $\varphi = “\lambda X. \text{Ayaka peeled three [apples in X].}”$ The Salient Sets Restriction again limits our attention to the sets $A = \{1, 2, 3, 4, 5, 6\}$ and $B = \{7, 8, 9, 10, 11, 12\}$. A satisfies φ (three of the apples in A have been peeled by Ayaka) but B does not, so $\llbracket \text{THE} \rrbracket(\varphi) = A$, the six apples in the first group.

When the context no longer supplies salient sets for consideration, however, the Salient Sets Restriction does not apply. As we saw in section 2, this affects different HIRCs in different ways. Consider first the HIRC with head quantifier *hanbun* ‘half’ (3) in the context with no contextually salient groups (5), repeated below as (19) with apples numbered. This HIRC is judged as infelicitous in this context.

(19) **Context with no salient subgroups:**



Under our analysis, this HIRC is computed as $\llbracket \text{THE} \rrbracket(\varphi)$, where $\varphi = “\lambda X. \text{Ayaka peeled half [apples in X].}”$ We note that any set containing an equal number of peeled and unpeeled apples satisfies φ in the context. The problem is that the φ -propositions corresponding to these sets do not entail one another: for example, “ $\lambda w. \text{Ayaka peeled half of the apples in } \{1, 4\} \text{ in } w$ ” does not entail and is not entailed by “ $\lambda w. \text{Ayaka peeled half of the apples in } \{1, 2, 4, 6\} \text{ in } w$.” Therefore no maximally informative φ -proposition can be identified and $\llbracket \text{THE} \rrbracket(\varphi)$ will be undefined (13). This explains the infelicitousness of the HIRC with *hanbun* ‘half’ (3) in this context with no salient groups (19), even though it was felicitous in the context with two salient sets (18) above.

On the other hand, the HIRC with the head quantifier *mittsu* ‘three’ (4) is felicitous in the context with no salient groups (19), but now denotes the three peeled apples, $\{1, 2, 3\}$, instead of $\{1, 2, 3, 4, 5, 6\}$ which was its denotation when evaluated in the context with two salient sets (18). This too is accurately modeled by our analysis. We interpret this HIRC as $\llbracket \text{THE} \rrbracket(\varphi)$, where $\varphi = “\lambda X. \text{Ayaka peeled three [apples in X].}”$ φ is true of $\{1, 2, 3\}$ and any superset thereof, such as $\{1, 2, 3, 4\}$, $\{1, 2, 3, 5\}$, etc. We note that this φ is upward-entailing; that is, “ $\lambda w. \text{Ayaka peeled half of the apples in } \{1, 2, 3\} \text{ in } w$ ” entails “ $\lambda w. \text{Ayaka peeled half of the apples in } X \text{ in } w$ ” for any $X \supseteq \{1, 2, 3\}$. Therefore “ $\lambda w.$

Ayaka peeled half of the apples in $\{1, 2, 3\}$ in w ” is the strongest φ -proposition that is true in the context. $\llbracket \text{THE} \rrbracket(\varphi)$ returns the corresponding object, $\{1, 2, 3\}$.

Here is a case where we can clearly see the importance of using a maximal informativeness semantics for THE. In context (19), the predicate $\varphi = “\lambda X. \text{Ayaka peeled three [apples in X]}”$ is also satisfied by the set of all apples in the context, $\{1 \dots 12\}$. Under a traditional maximality semantics for the definite determiner, $\llbracket \text{THE} \rrbracket(\varphi)$ would thus incorrectly return this set of all apples, $\{1 \dots 12\}$.

3.4. Summary

In this section we presented a new proposal for the syntax and semantics of Japanese HIRC. We proposed that the derivation of a Japanese HIRC involves modification of the restriction of the HIRC’s head quantifier with a null operator. We abstract over this position by moving the null operator to the edge of the HIRC, and interpret the HIRC as a definite description with a null definite head THE. The null operator movement explains the island-sensitivity of HIRC documented in Watanabe (1992, 2003). It also predicts the correct scope for the quantifier of the HIRC’s head (see Shimoyama, 1999).

We showed that this proposal is able to accurately model the interpretation of HIRC with different head quantifiers in different contexts. Crucial to these computations was the use of a maximal informativeness semantics for the definite determiner THE (von Stechow, Fox, and Iatridou 2012) and a straightforward assumption about the role of context in the interpretation of definite descriptions.

3.4. Theoretical note: the syntax of HIRC in a Copy Theory of movement

In this section we present a variant of our syntactic proposal following recent work on the Copy Theory of movement and the interpretation of chains. This alternative formulation makes the syntactic derivation of Japanese HIRC much more parallel to the head-raising analysis of HERCs and obviates the generation of a null operator. This further reinforces the early insight of Itô (1986) and Watanabe (1992), of a deep syntactic similarity in the derivations of HIRC and HERCs. The compositional semantics for HIRC, however, will remain unchanged under this formulation.

Chomsky (1993) proposed that movement does not leave traces, but instead “copies” entire syntactic objects. How each copy in a chain is interpreted or realized is determined at the interfaces. In particular, the correct semantics for movement chains does not obtain if both copies of a movement chain are interpreted as is. Fox (2002) proposes the operation of *Trace Conversion*, which converts the tail copies in DP movement chains into bound definite descriptions. Trace Conversion itself is made up of two parts: *Variable Insertion* and *Determiner Replacement*. Under this view, the process of constructing and interpreting a movement chain can be broken down into the following operations:

(20) **Decomposing the syntax/semantics of DP movement:**

- a. *Agree and Attract*: a higher probe Agrees with the DP and Attracts it.
- b. *Merge*: the DP is remerged (copied).
- c. *Predicate Abstraction* (Heim & Kratzer, 1998): a λ -binder is inserted below the higher copy.
- d. *Variable Insertion* (Fox, 2002): the domain of quantification on the lower DP is modified to be equal to the variable bound by the λ -binder introduced in (c).
- e. *Determiner Replacement* (Fox, 2002): the determiner in the lower copy is replaced with THE.

For example, consider the interpretation of the sentence *John read every book* as in (21) below. We assume that the quantifier “every book” in object position must QR for interpretation (see, among others, Heim & Kratzer, 1998). This covert movement operation will follow the procedure in (20) above: the DP is copied to a higher position and a corresponding λ -binder is inserted (*QR* below). Then, the lower copy is converted into a definite description “the book x ” (*TC* below).

(21) **The interpretation of *John read every book* via QR and Trace Conversion:**

$$\begin{array}{l} \text{John read every book} \xrightarrow{QR} [\text{every book}] \lambda_i \text{John read } [\text{every book}]_i \\ \xrightarrow{TC} [\text{every book}] \lambda x \text{John read } [\text{the book } x] \end{array}$$

With this background in place, we propose that Japanese HIRC's can be derived using a subset of the fundamental operations in (20), followed by the addition of the null definite determiner THE. Such a derivation for the HIRC in (3) is given below.

(22) **Syntactic derivation for the HIRC (3) using the independently attested operations in (20):**

- a. *Agree* and *Attract*: Ayaka peeled [_{DP} half apples]
- b. *Merge*: (the head DP is **not** copied)
- c. *Predicate Abstraction*: λ_i Ayaka peeled [_{DP} half apples]
- d. *Variable Insertion*: λX . Ayaka peeled [_{DP} half [apples in X]]
- e. *Determiner Replacement*: (the head DP's quantifier is **not** replaced)
- f. Merge a null definite head THE: THE [λX . Ayaka peeled [_{DP} half [apples in X]]]

The resulting structure in (22f) is logically identical to the result of our proposal above in (7) and yields the same interpretation.

4. Previous approaches

We have proposed a new semantic analysis for Japanese HIRC's that posits a movement relation between the HIRC's head and the edge of the relative clause. Two prominent analyses for the semantics of Japanese HIRC's have been proposed in recent years: an E-type anaphora analysis (Hoshi 1995; Shimoyama 1999), and an event semantics analysis (Grosu 2010; Grosu and Landman 2012). Here we briefly review these analyses in light of the new data from Section 2.

In the E-type pronoun analysis, a HIRC is interpreted similarly to cross-sentential anaphora and refers to an entity made salient in the relative clause. Thus, it is assumed that there is a (possibly null) pronoun *outside* the HIRC, whose referent is determined by a salient entity *within* the relative in conformity with Kuroda's (1975) Relevancy Condition.

In Section 2, we discussed what we called domain and witness set readings of HIRC's with quantificational heads. The E-type analysis predicts both to be robustly available. To illustrate this more clearly, we will give an example of a standard E-type pronoun in English:

- (23) Only [**half of [the students]_i**]_j turned in papers...
- a. ✓ ...and they_j will all pass. = witness set of *half of the students*
 - b. ✓ ...but they_i will still all pass. = domain of *half of the students*

In (23), the E-type pronoun *they* can refer to either the domain or witness set of *half of the students*. If HIRC's involve anaphora that are analogous to the pronoun in (23), then we would expect both kinds of readings to be available for Japanese HIRC's, contrary to what we have reported for our consultants in Section 2. In fact, a core property of an E-type analysis is the flexibility it allows in what the E-type pronoun refers to. However, we know of no way that a theory of E-type anaphora could be restricted so as to systematically exclude the witness set reading, and thus the judgments reported here constitute a challenge for such an analysis.

Similarly, the event semantics analysis is not restrictive enough in that it predicts both domain and witness set readings to always be available. In such an analysis, Japanese HIRC's contain the functional category ChooseRoleP (ChRP) that introduces a free variable of type *e* that is in some salient relation with the event of the HIRC. The reference of the HIRC is then interpreted as equal to this free variable. However, the ChRP is free to choose any relation to the event description (so long as the Relevancy Condition is not violated, cf. fn. 1), and thus there is no principled way to exclude the witness set reading.

In summary, previous approaches to the semantics of Japanese HIRC's compute the reference of HIRC's in a way that overgenerates possible interpretations for HIRC's, and does not predict the systematic sensitivity to context presented in Section 2 and explained by our analysis.

5. Conclusion

We have presented new data on the interpretation of Japanese HIRC with quantificational heads, showing the availability of domain readings but not of witness set. This led to a new analysis involving abstraction over a restriction on the domain of the quantificational head and the use of a null definite determiner with maximal informativeness semantics. This analysis accounts for the new data, the island sensitivity of HIRCs, and the fact that quantificational heads take scope within the HIRC.

In closing, we note two difficulties for our approach. First, verbs of creation appear to allow the witness set reading: in (25), half a paper is written, and this is what the professor reads. (We thank Chris Davis, p.c., for this observation.) A domain reading where the entire dissertation is read is simply not possible because it has not been written; however it is not clear how our approach would exclude the domain reading and allow the witness set reading in just such contexts.

- (24) [_{HIRC} [tetsuya-shite **ronbun-o hanbun** kaita] -no]-o sensei-ga yonde-kureta
all.nighter-did paper-ACC half wrote -NO-ACC teacher-NOM read-BENEFACTIVE
Literally ‘The teacher read (for me) [(I) pulled an all-nighter and wrote **half paper**].’

Finally, there is speaker variation regarding HIRC interpretation: we have found some speakers who allow both domain and witness set readings in a variety of contexts. We leave this as an issue for future research, including whether a modification of our proposal could account for these speakers.

References

- Chomsky, Noam. A minimalist program for linguistic inquiry. *The view from Building 20*, 1993.
- von Stechow, Kai, Danny Fox, and Sabine Iatridou. 2012. Definiteness as maximal informativeness. Ms, MIT.
- Fox, Danny. 2002. Antecedent-contained deletion and the copy theory of movement. *Linguistic Inquiry* 33.
- Grosu, Alexander. 2010. The status of the internally-headed relatives of Japanese/Korean within the typology of definite relatives. *Journal of East Asian Linguistics* 19: 231-274.
- Grosu, Alexander, and Fred Landman. 2012. A quantificational disclosure approach to Japanese and Korean internally headed relatives. *Journal of East Asian Linguistics* 21: 159-196.
- Heim, Irene and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Blackwell.
- Hoshi, Koji. 1995. Structural and interpretive aspects of head-internal and head-external relative clauses. Doctoral Dissertation, University of Rochester.
- Itô, Junko. 1986. Head-movement at LF and PF: the syntax of head-internal relatives in Japanese. *University of Massachusetts Occasional Papers in Linguistics* 11: 109-138.
- Kim, Min-Joo. 2007. Formal linking in internally headed relatives. *Natural Language Semantics* 15: 297-315.
- Kuroda, Sige-Yuki. 1976-77. Pivot-independent relativization in Japanese III. *Papers in Japanese Linguistics* 5: 157-179.
- Rullmann, Hotze, and Sigrid Beck. 1998. Reconstruction and the interpretation of *which*-phrases. In *Reconstruction: Proceedings of the 1997 Tübingen workshop*.
- Shimoyama, Junko. 1999. Internally headed relative clauses in Japanese and E-type anaphora. *Journal of East Asian Linguistics* 8.
- Watanabe, Akira. 1992. *Wh*-in-situ, subjacency, and chain formation. *MIT Occasional Papers in Linguistics* 2.
- Watanabe, Akira. 2003. *Wh* and operator constructions in Japanese. *Lingua* 113:519-558.