The Proceedings of the Seventh Tokyo Conference on Psycholinguistics

Edited by Yukio Otsu

Hituzi Syobo
Tokyo

Sloppy Readings of a 'Referential' Pronoun in Japanese

Misato Hiraga and Jon Nissenbaum
McGILL UNIVERSITY

1. Two Reflexives in Japanese

Japanese has two versions of the reflexive pronoun corresponding to English 'himself': *kare-zisin* ('him-self') and *zibun-zisin* ('self-self'). Both of these items require a local, c-commanding antecedent. In the examples in (1), both sentences can only mean: Taro criticized himself, Taro.

(1) a. Taro$_i$-ga **kare-zisin$_{i/*j}$**-o hihansita.
T -NOM him- self -ACC criticized.
"Taro criticized himself"

b. Taro$_i$-ga **zibun-zisin$_{i/*j}$**-o hihansita.
T -NOM self- self -ACC criticized.
"Taro criticized himself"

However, in some environments, they are famously not interchangeable. Following unpublished work of Takako Aikawa, Richards (1997) argues that *karezisin* and *zibunzisin* are unambiguously interpreted, respectively, as co-referential and bound reflexives, exemplifying the distinction introduced in Reinhart (1983). As evidence for this conclusion, Richards noted that in a

situation in which Taro *and* Ziro both criticized Taro, then sentence (2a) is false, while in the same situation, (2b) is true.

(2) a. Taro$_i$-dake-ga *karezisin*$_i$-o hihansita.
T -only-NOM himself-ACC criticized
"Only Taro criticized Taro."

b. Taro$_i$-dake-ga zibunzisin$_i$-o hihansita.
T -only-NOM selfself-ACC criticized
"Only Taro criticized himself."

(3) a. λx[x criticized Taro]

b. λx[x criticized x]

It seems that with *karezisin*, what gets attributed to Taro is the property of 'criticizing Taro' (3a). That being the case, (2a) is false because Taro is not the only one in the situation who holds this property. In contrast, with *zibunzisin* the property that gets attributed to Taro is that of 'self-criticizing' (3b). In the situation described by Aikawa/Richards, (2b) is true, since Taro is the only self-criticizer.

This way of distinguishing the two reflexives is corroborated, Richards/Aikawa note, by the fact that *karezisin* cannot be bound by quantifiers, as illustrated by (4):

(4) Daremo -ga {**karezisin* / zibunzisin}-o hihansita.
Everyone-NOM { him-self / self-self }-ACC criticize
"Everyone criticized himself."

So it looks as though whenever the two readings can be truth-conditionally distinguished, *karezisin* always gets the co-referential, or "Strict Identity" reading, while *zibunzisin* receives the bound-variable, or "Sloppy Identity" reading, using the terminology of Ross (1967).

However, in what follows, we aim to show that this picture is too simple. In particular, there are environments in which *kare-zisin*, too, can get a bound variable reading. After presenting some novel data illustrating these surprising bound variable uses of *kare-zisin*, we will argue that the distribution of variable and referential readings follows from a claim about semantic processing — namely that pronouns are optionally assigned their values at an early or late stage of computation — together with an inescapable fact about *kare-zisin* (it needs a referential antecedent), and the fact that discourse contexts (that is, sets of propositions held in the common ground) can have a filtering effect on the semantic outputs of otherwise ambiguous sentences.

2. A Puzzle: Bound Variable Uses of *Kare-zisin*

We will start by looking more closely at the sentences in (2) that led Richards to conclude *kare-zisin* is a co-referential anaphor while *zibun-zisin* is a bound variable. The situation that Richards considered in drawing his conclusion is what we can call a "Multiple-Taro-criticizing situation". This is illustrated in the top row of the table in (5) — the arrows are intended to indicate a situation in which all three participants are criticizing Taro. Richards' hypothesis correctly predicts (2a) to be false and (2b) to be true in the Multiple-Taro-criticizing situation.

(5)

	(2a) Taro-dake-ga *karezisin*-o hihansita. Only T criticized ***karezisin***		(2b) Taro-dake-ga zibunzisin-o hihansita. Only T criticized ***zibunzisin***	
	Predicted	Actual	Predicted	Actual
Situation (i): ***Multiple Taro-criticism*** 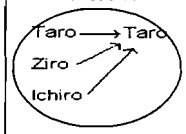	F	F	T	T
Situation (ii): ***Multiple SELF-criticism*** 	T	**F** *(for many speakers)*	F	F

However, Richards did not consider a different kind of situation, one in which each person criticizes only himself. This is shown in the bottom row of (5). In this situation, the sentence with *zibun-zisin* (2b) is false just as expected; Taro is not the only one criticizing himself. However, it turns out that many speakers also find (2a) false as well. This fact is quite unexpected if *kare-zisin* is always interpreted referentially. In the multiple self-criticizing situation, Taro is indeed the only person criticizing Taro. It seems, then, that not only are speakers able to get a bound-variable reading for *karezisin* in this situation — they're getting *only* the bound-variable reading for it. This suggests that Richards and Aikawa's conclusion was not fully correct. And now it becomes quite puzzling

why *karezisin* cannot get a sloppy reading in situation (i), or have a quantificational antecedent.

Example (6) provides another environment in which *karezisin* very clearly gets a bound-variable interpretation: "*Only Taro* is allowed, if Hanako is recommending him, to be a self-recommender". This sentence would be contradictory if it got the referential reading described in (7b). But (6) is sensible; it is understood as asserting that a situation like (8A) is not permitted, while (8B) is permitted. Specifically, (8B) is permitted because Taro is the only self-recommender.

(6) Taro$_i$-dake-ga [Moshi Hanako-ga kare-o suisensiterunara] t_i *karezisin*$_i$-o suisensite-yoi.
T-only-NOM [if H -NOM him-ACC recommend] himself-ACC recommend-may

a. "Only Taro is allowed, if Hanako is recommending him, to be a self-recommender."

b. # "Only Taro is allowed, if Hanako is recommending him, to recommend Taro."

(7) a. If Hanako is recommending him (Taro), only Taro [λx. x may recommend Taro].

b. If Hanako is recommending him (Taro), only Taro [λx. x may recommend x.]

(8) **Situation A**: Hanako → Taro ⟲ Ziro ⟲ Multiple-*SELF*-recommending

Situation B: Hanako → Taro ⟲ ← Ziro Multiple-*Taro*-recommending

Why can *karezisin* receive a sloppy reading just in (6), or in (2a) in a multiple-self-criticizing situation? If it can sometimes get a sloppy reading, why can't it always get sloppy reading like *zibunzisin*?

3. A Proposal

3.1. Three Theses about pronouns and semantic computation

We would like to suggest that this puzzle could be solved in a straightforward way, if the following three empirical theses hold:

THESIS A: What it means for *karezisin* to be referential (or co-referential, to be precise) is that it is (obligatorily) coindexed with an antecedent that is referential.

THESIS B: Free indices are interpreted by means of a **Variable Assignment Function**, i.e. a function that maps a **syntactic index** to some value in the domain of discourse.

THESIS C: Valuation by the assignment function is *optional* at a given stage of the computation.

The last of these really constitutes the novel part of our proposal. The first two play an important role but are not especially interesting or novel. Thesis A seems like a plausible and natural way of characterizing the notion of (co-)referentiality of a syntactic constituent. Thesis B is likewise plausible, in that it has a fair amount of independent motivation as a framework for interpreting natural language expressions. It is intended to embody the specific assumptions listed in (9), adopting the framework of Heim & Kratzer's (1998) theory of semantic computation:

(9) a. Denotations are computed by a procedure that maps (sub-)trees to truth conditions relative to an assignment function *A*. The assignment function is to be understood as specifying part of the utterance context.

b. What it means for an NP to be referential is that its index is in the domain of the assignment function (henceforth, in *DOM(A)*).

c. If an NP moves, targeting some constituent β, its index, *i*, gets interpreted as a λ-operator which composes with the assignment-relative denotation of β — and it yields the lambda expression that maps x to the denotation of β under the **modified** variable assignment, $A^{x/i}$.

This is just the mechanism of predicate abstraction that is described in Heim and Kratzer. The λ-function mentioned in (9c) is interpreted relative to the **modified assignment function** $A^{x/i}$. This is the function that is identical to *A*, except that it maps every occurrence of the index, *i*, to the value of the λ-bound variable x. In other words, when a structure is derived by movement, the trace and every constituent co-indexed with it become interpreted as variables bound under λ–abstraction. Since this is one of the now-standard approaches to movement and variable binding in work on syntax and semantics, neither Thesis A nor Thesis B is novel or especially exotic.

Suppose we understand semantic computation to involve a bottom-up procedure that computes the values of increasingly larger subtrees, as in (10). Note that this is just an expository move, and that nothing crucially hinges on taking this as a bottom-up computation.

(10) a. $[\![\alpha]\!]^A$ → *VALUE 1*

b. $[\![\beta]\!]^A$ → *VALUE 2*

c. $[\![\ [[\alpha][\beta]]\ ]\!]^A$ → *VALUE 3* = Combination of *VALUE 1* and *VALUE 2* given by composition rules

We will assume that the value computed for each subtree is stored in a short-term semantic buffer, to be accessed at the next step of computation. This is illustrated in (11):

(11)

	Subtree	**Stored in semantic buffer**
a.	$[\![[_{V'}$ criticize Mary$]]\!]^A$	→ λx[TRUE iff *x criticizes Mary*]
b.	$[\![[_{VP}$ [John $[_{V'}$ criticize Mary]]$]]\!]^A$	→ [*OUTPUT OF STEP a*](*John*) = TRUE iff *John criticizes Mary*

In the step shown in (11a), the value of the V' — that is, the lambda-expression shown on the right — is stored to the buffer. This value is then accessed at the step shown in (11b), where the V' composes with the VP-internal subject by Function Application. The output of this step, in turn, would be accessed at the subsequent step, and so on. So far, this is just standard, rule-by-rule bottom-up computation.

3.2. Optional Valuation at a given stage of computation

Now we turn to the case of sub-trees that contain pronouns. Remember that pronouns (and other constituents) whose indices are in the domain of the local assignment function are interpreted as referring to individuals in the context. Let's consider the context that's partially specified in (12). The assignment function *A* maps syntactic indices to individuals in this context. Here, *A* maps the index "3" to *'Professor Shimoyama'*.

(12) **Context^A**: *A(3)* = *'Professor Shimoyama'*

It is important to keep in mind here that we are taking assignment functions and indices to be mental objects that play a role in semantic computation. More specifically, (12) states an equivalence between two symbolic expressions in a mental representation of a context. Noun Phrases like [$_{NP}$ Professor Shimoyama] that describe or refer to individuals can be taken, in the mentalistic approach to semantic computation, to "denote" symbolic objects in the representation of a context. Given the equivalence stated in (12), it is just as correct to say that the NP [$_{NP}$ Professor Shimoyama] denotes 'A(3)' as it is to say that the NP denotes the mental object 'Professor Shimoyama'. The denotation of a referentially indexed pronoun, like [$_{NP}$ her]$_3$ or [$_{NP}$ kare-zisin]$_3$, is always 'A(3)'. In the context (12), since the objects 'A(3)' and 'Professor Shimoyama', are equivalent expressions, pronouns indexed *3* also denote 'Professor Shimoyama'.

Given the Assignment Function of (12), the value of a pronoun with the index 3 can be written to the semantic buffer in one of two ways, shown in (13):

(13)

	Subtree		**Stored in semantic buffer**
a.	$[\![pronoun_3]\!]^A$	→ *'A(3)'*	→ 'PROFESSOR SHIMOYAMA'
b.	$[\![pronoun_3]\!]^A$	→	*'A(3)'*

One option, shown in (13a), is to take the value of the pronoun, A(3), and rewrite it into the semantic processing buffer as *'Professor Shimoyama'* by undertaking the additional step of applying the function A to the index. Alternatively, the pronoun's value can be written into the buffer as simply 'A(3)', as in (13b). Crucially, under this second option what gets stored in the short term semantic buffer is an expression that contains the index. In other words, we propose, the computational step that is optional with pronouns is to rewrite the index as the equivalent symbol 'Professor Shimoyama' given by the assignment function. We can refer to Option (13a) as the ***Immediate Valuation*** of a pronoun, and to Option (13b) as the ***Indexed Interpretation*** for a pronoun.

The option that is selected at an early computational step has consequences for the options available at subsequent steps. For instance, consider the interpretation of the V' [criticize her] in (14). If the pronoun itself was originally interpreted by **Immediate Valuation** — that is, if it was written into the buffer as 'Professor Shimoyama' as in step (13a) — then the interpretation of the V' dominating it will access the output of (13a), and it will produce (14a). On the other hand, if the pronoun was initially written to the buffer with an **Indexed Interpretation** — "A(3)", as in step (13b) — then this will be the value accessed at step (14), and we'll end up with (14b).

(14)

	Subtree		**Stored in semantic buffer**
a.	$[\![[_{V'}$ criticize $her_3]]\!]^A$	→	$[\![$criticize$]\!]$([*OUTPUT OF STEP* (13*a*)])
			= λx[TRUE iff *x criticizes* PROFESSOR SHIMOYAMA]
b.	$[\![[_{V'}$ criticize $her_3]]\!]^A$	→	$[\![$criticize$]\!]$([*OUTPUT OF STEP* (13*b*)])
			= λx[TRUE iff *x criticizes A(3)*]

Again, these expressions are truth-conditionally equivalent, given the assignment function indicated in (12). But the empirical claim is that they are nevertheless symbolically distinct expressions, either of which may be written into the short term semantic processing buffer.

In the next section we will argue that selecting the **Indexed Interpretation** for a pronoun can yield a truth-conditional difference, in situations where *Predicate Abstraction* — the semantic rule for interpreting movement structures — applies on a subsequent step.

4. *Karezisin* and *Zibunzisin*

4.1. The Distribution of Strict and Variable Readings

If it is true that *karezisin* has a lexical requirement to have an index in the domain of the assignment function *A* — that is, if Thesis A of Section 3.1 is correct — then it follows that *karezisin* will never be able to have a quantificational antecedent. We therefore take Richards' (1997) observation about the lack of quantificational antecedents for *karezisin* as motivation for Thesis A. Specifically, we assume that *karezisin* has the following two lexically given requirements:

(15) Requirements for *karezisin*:

(i) Its index is in DOM(A), and

(ii) It must have a local antecedent which has the same index

That means that the two structures shown in (16), only (16a), on the left, is a legitimate structure for *karezisin*. In (16a) the local antecedent is referential, so its index is in the domain of the assignment function. In (16b), *karezisin* is

co-indexed with a local antecedent, but in this case the antecedent is quantificational. Therefore, its index cannot be in the Assignment Function's domain, and the requirement on *karezisin* is violated.

(16) a.

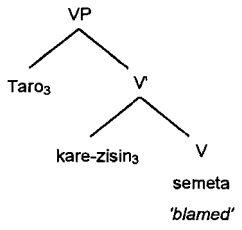

b. *

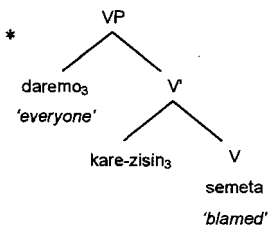

However, the fact that *karezisin* has to have a referential index does not stop it from receiving a bound variable reading. On the assumptions sketched in the previous section, it can indeed get a sloppy reading. Consider the structure shown in (17). This is a structure derived by movement of *Taro-Dake-ga*, "Only Taro", leaving a trace in the VP-internal position. We have represented this using the Heim & Kratzer (1998) method of adjoining the index to the target of movement, although this way of representing it is not crucial for our account. What is crucial is that movement is interpreted via Predicate Abstraction, as we described earlier. But the above assumptions allow it to receive a sloppy reading.

(17)

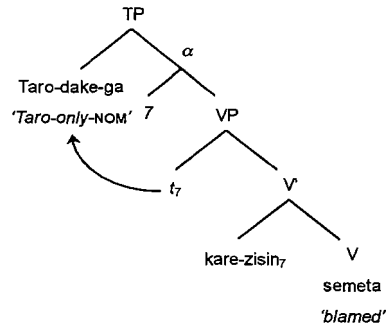

Because of the referentially indexed pronoun in (17), there are two different ways to compute its meaning. These are illustrated in (18).

(18) **Two different computations for the meaning of (17):**

	Immediate Valuation	*Indexed Interpretation*
a.	$[\![karezisin_7]\!]^A$	$[\![karezisin_7]\!]^A$
	→ $A(7)$ → Taro	→ $A(7)$
b.	$[\![V']\!]^A$	$[\![V']\!]^A$
	→ $[\![semeta]\!]^A$(Taro)	→ $[\![semeta]\!]^A(A(7))$
	→ λy.[y blamed Taro]	→ λy.[y blamed $A(7)$]
	...	...
c.	$[\![\alpha]\!]$	$[\![\alpha]\!]$
	→ λz $[\![VP]\!]^{A^{z/7}}$	→ λz $[\![VP]\!]^{A^{z/7}}$
	→ λz[z blamed Taro]	→ λz[z blamed z]

On the first step of the computation in (18a), the semantic value of the pronoun is written into the processing buffer. This can be done in one of the two ways described in the previous section, according to Thesis C. Whichever option is selected at this step has an effect on the subsequent steps, because the output of step (a) becomes one of the inputs for step (b). If *karezisin* is **Immediately Valued** by the assignment function, as shown in the computation on the left side of (18a), then the argument for the verb in step (b) will be the expression 'Taro', given that the Assignment Function maps the index 7 to this (mental) individual. The value of the V' in step (18b) will then be the *Taro-blaming* function. When predicate abstraction applies in step (c), only the trace of the moved subject will be interpreted as a bound variable. The result of the computation on the left, therefore, will be the strict reading for the pronoun: "Only Taro has the property of being a Taro-blamer".

On the other hand, if the pronoun is valued on the first step by **Indexed Interpretation**, as shown on the right-hand side, then the input for the second step will be the expression 'A(7)'. Then, the value computed for the V' will be the function that describes the property of being "*a blamer of 'A(7)'*", that is, a blamer of the individual that the context assigns to the index 7. When Predicate Abstraction applies on the right-hand side of step (18c), however, the index will be re-valued as a bound variable, giving the *self*-blaming function. Consequently, *karezisin* receives a bound-variable interpretation (19).

(19) "Only Taro has the property of being a self-blamer"

(20)

(19) is False!

Therefore, as a description of the situation where there are *multiple* self-blamers, this proposition is false (20). As we pointed out at the beginning, this is in fact the judgment that many speakers have about the sentence, indicating the availability of this interpretation.

4.2. A Problem: Why does *karezisin* sometimes fail to get a sloppy reading?

Now that we have shown a set of assumptions that would allow *karezisin* to get a sloppy interpretation, we face a different problem. Namely, why does it sometimes *fail* to get that reading? That is, what accounts for the facts that have led to the conclusion that it is always strict? Remember that we argued part of the answer — the lack of quantificational antecedents — involves the rather bland assumption that it must have a referential index. But *karezisin* sometimes gets only a strict reading with referential antecedents as well. Recall situation (i) from *Table (5)*: When *Ziro* and *Ichiro* are also blaming Taro —what we called the multiple-Taro-blaming situation — then the sentence (21) is judged false.

(21) Taro-dake-ga *karezisin*$_{7}$-o semeta.
T. only-NOM him-self-ACC blamed
= "Only Taro blamed Taro"
≠ "Only Taro blamed himself."

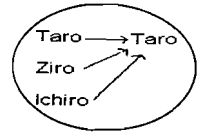

Evidently, in this kind of a situation, the sentence can not readily be understood as meaning, "Only Taro is a self-blamer". Rather, it seems to have just the meaning "Only Taro is a Taro-blamer". Given that *karezisin* can get a sloppy reading, as we have shown, this fact is now surprising. Why don't speakers judge (21) true in the multiple-Taro-blaming situation? That is, why don't

speakers interpret (21) as *"Only Taro blamed himself"* when they judge it against this situation?

We offer a speculative answer to this question. We think it is plausible that this is a result of an incongruity between the focus structure of the sloppy version of the sentence, on the one hand, and the set of alternative propositions that are made salient by the Multiple-Taro-blaming situation, on the other. Specifically, we suspect the following three factors interact to filter out one of the two possible LF outputs (the sloppy one) of (21):

(22) a. *Dake* ('only') is a focus-sensitive operator. It requires, for its interpretation, a set of alternative propositions that is dependent on the context and that relates appropriately to the focus structure of the sentence (Rooth 1992).

b. The salient alternatives evoked by the Multiple-Taro-blaming situation is the set of propositions of the form [*x* blamed Taro].

c. The sloppy LF of (21) has a focus structure that evokes alternative propositions of the form [*x* blamed *x*].

If these three assumptions are correct, the sloppy LF — while derivable — is inaccessible due to a mismatch between its focus structure and the discourse context. The alternative propositions mentioned in (22a) are part of the common ground of the discourse, and correspond to the covert discourse anaphor that Rooth (1992) proposed is actually the innermost syntactic complement of focus-sensitive operators like *only*. Rooth argued that this discourse-dependent set of propositions needs to be congruent with the focus structure of the overt

material in the scope of *only*, similar to the congruity that must hold between a constituent question and the focus structure of felicitous answers:

(23) Question: Who did John invite?

Possible Answer: John invited (only) $MARY_F$

#(Only) $JOHN_F$ invited Mary

In effect, Rooth argued that sentences with focus, even in the absence of explicit questions, are always evaluated in relation to an implicit question (taken to denote a set of propositions, in this case propositions of the form *John invited x*; see Hamblin 1975.) We think it is plausible that in the truth-value judgment task that elicits judgments of 'false' for (21), the Multiple-Taro-blaming situation is playing the same role as the Question in (23), namely, making salient the set of propositions mentioned in (22b). That set of propositions does not match the set determined by the focus structure of the sentence in its sloppy version, however (22c).[1]

Thus, we think that the discourse context is having a filtering effect on the two possible outputs for the semantic computation. These assumptions, together with the proposals made in the previous subsection, would explain the distribution of readings for *karezisin*.

4.3. Zibun-zisin

The other type of reflexive, *zibun-zisin*, is different. *Zibun-zisin* can receive only a sloppy reading. We assume that what is special about *zibun-zisin* is that it is required, as a lexical property, to have its value be computed by **Indexed Interpretation**. Thus, its value will always be written to the computational

buffer as just "A(7)", for example, in a structure where 7 is its index. It will not get re-written as "Taro" even in a context where that is the value assigned to 7 by the assignment function.

(24) $[\![zibunzisin_7]\!]^A$ → $A(7)$ *(* → Taro)*

Consequently, Predicate Abstraction applying on a subsequent step would always value *zibunzisin* as a bound variable. Alternative propositions evoked by the discourse context, then, could not have a filtering effect, simply because only one interpretation can be derived to begin with.

5. Extension: the Sloppy Indexical Puzzle

Consider the English examples (25).

(25) a. Only I understood my question.

b. Only I understood myself.

Heim (1994) noticed that, despite the fact that the first-person pronoun is indexical, these examples have a bound-variable reading. Thus, (25a) is understood as meaning "I was the only *x* who understood *x*'s question."

We think our suggestions about *karezisin* can explain Heim's puzzle as well, if we make a simple assumption about what exactly makes the first-person pronoun indexical. Specifically, we will assume that the first-person pronoun has a dedicated index that it always receives (e.g. "1"), and that *every*

assignment function maps this dedicated index to the "speaker" coordinate of the context.

If that is the method for assigning an interpretation to the first-person pronoun, then as a consequence its index is always in the domain of the local assignment function. Hence there will always be two ways to write its value into the semantic processing buffer, just like *kare-zisin*.

The two computations for the meaning of (25a) are illustrated in (26). If the pronoun *my* is subject to **Immediate Valuation** — that is, if it gets re-written as "the speaker" as on the left — then (25a) will get the strict reading.

(26)

	Immediate Valuation	**Indexed Interpretation**
a.	$[\![my_1]\!]^A$	$[\![my_1]\!]^A$
	→ $A(1)$ → "the speaker"	→ $A(1)$
b.	$[\![my_1\ question]\!]^A$	$[\![my_1\ question]\!]^A$
	→ "the speaker's question"	→ "$A(1)$'s question"
c.	$[\![VP]\!]^A$	$[\![VP]\!]^A$
	→ $[\![understood]\!]^A(\underline{output\ [b]})(t_1)$	→ $[\![understood]\!]^A(\underline{output\ [b]})(t_1)$
	→ [$A(1)$ understood the speaker's question]	→ [$A(1)$ understood $A(1)$'s question]
	...	...
d.	$[\![\alpha]\!]$	$[\![\alpha]\!]$
	→ $\lambda z\,[\![VP]\!]^{Az/1}$	→ $\lambda z\,[\![VP]\!]^{Az/1}$
	→ λz[z understood the speaker's question]	→ λz[z understood z's question]

Alternatively, as shown on the right, *my* can get the indexed interpretation. And, just as with *karezisin*, the result will be the sloppy reading, expressed in the output of step (d).

The first-person reflexive *myself* in (25b) receives only a sloppy reading. This is a general property of English reflexives. That suggests that English reflexives, just like *zibun-zisin*, have the special property of requiring the indexed interpretation.

6. Conclusions

The prevailing view about Japanese reflexives has been that *kare-zisin* always receives a strict interpretation while *zibun-zisin* is always sloppy. In this paper we have identified some environments in which *kare-zisin* receives a bound variable reading, indicating that the prevailing view is too simplistic. We suggested that the explanation for the pattern of sloppy readings for *kare-zisin* boils down to two facts:

(i) *kare-zisin* is lexically restricted to have an antecedent in the domain of a contextually determined assignment function.

(ii) In the course of semantic computation, a pronoun can be immediately valued by the assignment function — i.e. assigned its contextually determined reference — or it can be written to a short term semantic buffer as simply an indexed expression.

If the latter option is chosen, even a referential pronoun can get a bound variable interpretation on a subsequent step of computation. We made what we think is a plausible speculation about why *kare-zisin* sometimes fails to get a sloppy reading even when its antecedent is referential. Namely, in contexts in which the sloppy reading would be true and the strict reading false — e.g. the Multiple-

Taro-blaming scenario — the salient alternative propositions are of the form *"x blamed Taro"*. That set of propositions is incongruent with the focus alternatives that are computed for the sloppy version of the sentence.

Finally, we suggested a way that this approach to *kare-zisin* can be extended to Heim's sloppy indexical puzzle. We proposed that indexical pronouns are always in the domain of the local assignment function, just like *kare-zisin*, so sloppy readings should emerge in exactly the same way in the two cases.

Notes

* We would like to thank the audience and organizers of the 2006 TCP, and we are also grateful to the TCP proceedings editors for their kindness and patience. This work was supported by a grant from the Fonds Québécois de la Recherche sur la Société et la Culture, award number 99440.

[1] Notice that a very similar situation seems to obtain for English. Consider the sentence (i), in relation to each of the two situations described in (ii). (In English, one cannot use the exact equivalent of a sentence like (21) due to the Principle B violation that would result; here we avoid that problem by embedding the pronoun in a possessive NP.)

(i) Only John likes his dog.

(ii) a. John likes John's dog, Bill likes Bill's dog, and Fred likes Fred's dog.
b. John, Bill and Fred all like John's dog.

Unlike *kare-zisin*, English *him* unquestionably is able to receive bound variable readings. Not surprisingly, judged in relation to the situation described in (ii)a,

(i) is false; in that situation John is *not* the only person who likes his own dog.

What is interesting is that the sentence (i) also seems false as a description of situation (ii)b. Indeed, speakers whom we consulted tended overwhelmingly to report that (i) is false in both situations. Thus, in the same kind of task that Richards found Japanese speakers get only the referential reading for *kare-zisin*, English speakers get just the referential reading for *his*. We take this as strong support for the ability of contexts/situations to disambiguate toward the referential reading of a pronoun in precisely the way we suggest for *kare-zisin*.

References

Aikawa, T. (1995) Unpublished ms. cited in Richards (1997).

Futagi, Y. (2004) *Japanese focus particles at the syntax-semantics interface.* Unpublished Ph.D. dissertation, Rutgers University.

Hamblin, C. (1973) Questions in Montague English. *Foundations of Language*, 10, 41–53.

Heim, I. (1994) Puzzling reflexive pronouns in *de se* reports: http://semanticsarchive.net/Archive/mVkNjZlY/Bielefeld%20handout%2094.pdf

Heim, I. and A. Kratzer. (1998) *Semantics in generative grammar*. Blackwell.

Miyagawa, S. and T. Tsujioka. (1994) Argument structure and ditransitive verbs in Japanese. *Journal of East Asian Linguistics*, 13, 1–38.

Nakayama, S. (1982) *On English and Japanese pronouns.* Unpublished M.A. thesis, University of Tokyo.

Reinhart, T. (1983) Anaphora and semantic interpretation. London: Croom Helm.

Richards, N. (1997) Competition and disjoint reference. *Linguistic Inquiry*, 28.1, 178–187.

Rooth, M. (1992) A theory of focus interpretation. *Natural Language Semantics*, 1, 75–116.

Ross, J. R. (1967) *Constraints on variables in syntax.* Ph.D. thesis, MIT.

Sag, I. (1975) *Deletion and Logical Form.* Ph.D. thesis, MIT. Published by MIT Working Papers in Linguistics.

Saito, M. and H. Hoji (1983) Weak crossover and move-α in Japanese. *Natural Language and Linguistic Theory*, 1, 245–259.

(mhiraga@linguist.umass.edu; jon.nissenbaum@mcgill.ca)