A note on backward anaphora[1]

Valentina Bianchi
Università di Siena
(revised draft, April 2010)

Abstract: In this paper I present four empirical observations about backward anaphora; building on Schlenker (2005), I suggest that the phenomenon should be though of in terms of a left–to–right interpretation process, in which discourse referents can be dynamically stored (by NPs) and retrieved (by pronouns) at two distinct and interacting levels of organization of the grammar: the «local» level of semantic computation of constituents, and the global discourse level.

1. Introduction

In generative approaches to anaphora, a persistent view is the distinction between binding and coreference, dating back to Reinhart (1983, 1986).[2] True binding is syntactically constrained by a c–command requirement (Reinhart's Bound Anaphora Condition), and results in the binder's argument position and the bound pronoun being interpreted as two occurrences of the same variable bound by a lambda operator :

(1)　　a. John$_1$ loves his$_1$ wife
　　　　b. John [λx. loves (x, x's wife)]

On the most restrictive assumptions (see e.g. Heim & Kratzer 1998: 260–266), semantic binding as in (1b) strictly requires syntactic binding as in (1a), and vice versa, syntactic binding is invariably interpreted as semantic binding.

On the other hand, the kind of anaphora that obtains in non–c–command configurations is "accidental coreference": the unbound pronoun is a free variable which is assigned a value by the context; the assigned value may happen to be the referent of some other noun phrase (presumably made salient by the linguistic context).[3]

(2)　　*John* is unhappy. [*His* wife] has left *him*.
　　　　　　　　　↓　　　　　　　↓
　　　　　　　　　j　　　　　　　j

The main reason for maintaining this distinction is the (not uncontroversial)[4] empirical observation that quantificational antecedents are sensitive to the c–command requirement: a quantifier can semantically bind a pronoun in its c–command domain (3), but not one outside its c–command domain (4).

(3)　　No man$_1$ loves his$_1$ wife (as much as she deserves).
(4)　　* No man$_1$ is happy. His$_1$ wife has left him$_1$.

Reinhart's original view, also reiterated in much subsequent work (e.g. Grodzinsky & Reinhart 1993, Reinhart 2000, 2004) is that accidental coreference as in (2) is a pragmatic phenomenon: it is not syntactically constrained, and it falls outside the domain of semantics proper.[5]

Nowadays, various discourse–oriented approaches to anaphora have been developed, notably Grosz et al.'s (1983, 1995) Centering Theory and Ariel's (1990) Accessibility Theory, and also some pragmatic neo–gricean approaches (Levinson 1991, Huang 2000; see Huang 2007, 245–274 for a recent overview). The former are

generally more concerned with the problem of inter–sentential anaphora and discourse topic management; the neo–gricean approaches seem to be more concerned with binding configurations. All in all, there has been relatively little interest in intra–sentential anaphora without c–command, also generally known as "backward anaphora":[6]

(5) Those who know *her* respect *Zelda (*every girl)*.

In this paper I would like to bring together four simple empirical observations about backward anaphora, and to critically re–examine Reinhart's view of this phenomenon. I will suggest that backward anaphora is constrained by the interplay of two factors: the placement of focus, on the one hand, and the intrinsic content of the co–referring expressions, on the other; furthermore, I will argue that a proper understanding of the phenomenon requires us to assume an incremental left–to–right interpretation process.

2. Four empirical observations

The first empirical observation is actually fairly old and well known: backward anaphora between a pronoun and a following R–expression is blocked when the R–expression bears focus (cf., among others, Chomsky 1976: 343–45, Williams 1997: 587–91, Erteschik–Shir 1997: 78 ff.).
    Here I will only consider instances of new information focus (focus in question–answer pairs), leaving open the question of whether other kinds of focus (e.g. contrastive focus) have a similar or a different effect.[7] The relevance of new information focus is illustrated by the contrast between (6) and (7):[8] in (6), the R–expression is focussed and the preceding pronoun cannot corefer with it; in (7), on the contrary, the R–expression is destressed – due to contrastive focus on the preceding verb – and backward anaphora is possible.

(6)   (Context: As for *John*, who does *his* wife really love?)
        ?* *His* wife loves *JOHN*.

(7)   (Context: –As for *John*, I believe *his* wife hates him. –You're wrong:)
        His wife LOVES *John*.

Although this observation is generally accepted in the literature, its scope is actually very restricted. When the two co–referring terms are both pronominal, focus on the second one does not block coreference. This holds not only for first and

second person pronouns, but also for third person ones: (8c) forms a minimal pair with (6).

(8) (Context: Who does your/my/his wife really love?)
    a. *My* wife loves *ME*.
    b. *Your* wife loves *YOU*.
    c. *His* wife loves *HIM*.

Similarly, focus does not block coreference when the co–referring terms are both R–expressions.[9] (9) forms another minimal pair with (6).

(9) (Context: Who does John's wife really love?)
    John's wife loves *JOHN*.

The final observation is that focus or destressing are completely irrelevant in configurations in which the second term is an R–expression c–commanded by the first term: (10b) is as unacceptable as (10a).[10]

(10) a. * *John/He* claims that Mary hates *JOHN*.
     b. * *John/He* claims that Mary HATES *John*.

These empirical observations can be schematically summarized as follows (the square brackets represent the c–command domain of the first term):

(i) Focus effect in [...*pro*]... *R*... configurations
(ii) Focus irrelevant in [...*pro*]... *pro*... configurations
(iii) Focus irrelevant in [...*R*]... *R*... configurations
(iv) Focus vs. destressing irrelevant in *[*pro/R*... *R* ..] configurations.

3. Reconsidering Reinhart's account

Let us start from the focus effect in (6) vs. (7). A fairly common assumption is that the focussed R–expression in (6) undergoes covert LF movement to some leftward position, crossing over the pronoun: this results in a Weak Crossover configuration.

(11) * $JOHN_i$ [his_i wife] loves t_i

Whatever the exact nature of the Weak Crossover constraint, it will rule out (11).[11]

This much is unobjectionable: note that Weak Crossover is a constraint on *binding*, and in fact, even the grammatical examples in (8) above do not allow for a genuinely bound interpretation (i.e., one involving two occurrences of the same variable). For instance, (8a) does not mean that I am the (only) individual who is loved by his own wife (12a); rather, it means that I am the (only) individual who my wife loves (12b).[12]

(12)   a. * The x such that x's wife loves x = the speaker
        b.   The x such that the speaker's wife loves x = the speaker

However, invoking the Weak Crossover constraint gives only half of an explanation for (6). Even if binding is correctly ruled out, the question remains of what rules out coreference. It seems to be often tacitly assumed that a focussed R–expression is a "quasi–quantifier", and as such, it cannot (co)refer. But this is too strong an assumption, as shown by the ambiguity of (13):

(13)   Only Joel voted for his proposal.   (Büring 2005b: 262)
        a. [λx. x voted for x's proposal] is true of Joel and of noone else (sloppy)
        b. [λx. x voted for Joel's proposal] is true of Joel and of noone else (strict)

The focussed R–expression *only Joel* gives rise to a strict–sloppy ambiguity. In Reinhart's approach, the sloppy reading (13a) (whereby Joel is the only self–voter) arises from syntactic (and semantic) binding, but the strict reading (13b) (whereby Joel is Joel's only supporter) is an instance of coreference between *Joel* and *his*. A similar ambiguity arises in the question–answer pairs in (14):

(14)   a. Who$_i$ voted for his$_i$ proposal? JOEL$_i$ did (vote for his$_i$ proposal).   (sloppy)
        b. After the discussion of Joel's report, who voted for his (=J's) proposal?
        JOEL (HIMSELF) did (vote for *his* proposal).   (strict)

Since a focussed R–expression can corefer, we do need some constraint to rule out coreference in (6). One possible candidate would seem to be Reinhart's Rule I:

(15)   Rule I (Intrasentential Coreference)
        NP A cannot corefer with NP B if replacing A, at LF, with a variable bound by the trace of B, yields an undistinguishable interpretation.
                                        (Grodzinsky & Reinhart 1993: note 11).

However, Rule I only applies in configurations where the pronoun can be bound by the *trace* of the antecedent at LF, where the trace corresponds to the argument position. Since in (11) the trace of the focussed R–expression fails to c–command the pronoun, (11) is not a possible binding configuration by the Bound Anaphora Condition, and Rule I cannot even apply.[13] On the other hand, Rule I applies to the c–command configurations in (10), and correctly rules out coreference independently of any focus effect.

Reinhart's (1986: 138–140) view of the contrast between (6) and (7) involves a specific discourse constraint, which she (tentatively) formulated as follows:

(16)   Topic–antecedent hypotesis: Backward anaphora is possible only if the antecedent is in sentence–topic position.[14]

Evidence in support of (16) is provided by the pair in (17a–b). According to Reinhart, in (English) active clauses either the subject or the object can qualify as the sentence topic, whereas in passive clauses the promoted subject is the obligatory sentence topic, excluding the by–phrase. As predicted by (16), in (17a) either the matrix subject or the object can be the antecedent of the backward pronoun, whereas in (17b), the passive subject is the only available antecedent. Finally, (17c) is inappropriate because the only plausible antecedent for the backward feminine pronoun fails to satisfy the Topic–antecedent constraint.

(17)   a. When *he* entered the room, $^{\sqrt{}}$*Max* greeted $^{\sqrt{}}$*Bill*.
       b. When *he* entered the room, $^{\sqrt{}}$*Max* was greeted by **Bill*.
       c. # When *she* entered the room, Max was greeted by *Kora*.

The proposed constraint (16) can also account for the contrast between (6) and (7), repeated here:

(6)   *His* wife loves JOHN.
(7)   *His* wife LOVES John.

A phrase bearing new information focus, being part of the comment, cannot be the sentence topic (cf. Zubizarreta 1998: 9–10): thus, the focussed R–expression in (6) cannot be the antecedent of the backward pronoun, because it fails to satisfy (16). In (7), on the contrary, the destressed R–expression can qualify as the sentence topic, hence as a possible antecedent for the backward pronoun.[15]

The constraint in (16) is fairly natural, for it is well known that anaphoric pronouns (tend to) take topics as antecedents. However, it cannot be a general constraint on intra–sentential coreference, in view of the data in (8) and (9) above. Two important qualifications are hidden in the term «backward anaphora»:

condition (16) only applies when (a) the coreference relation involves a pronoun and an R–expression, and (b) the pronoun is to the left of the R–expression.

Concerning the first qualification (a), it may fall out naturally if we conceive of (16) as a constraint on the *assignment function*, i.e. the context–determined function from numerical indices to individuals which assigns a value to all free pronouns in an LF.[16] A possible reformulation might be as follows:

(18)   Topic–antecedent hypothesis (revised): An (index on an) unbound pronoun can be assigned the referent of a following R–expression only if the latter is the sentence topic.

This reformulation is compatible with the possibility of coreference in (8) and (9). In (8c), for instance, the two coreferring third person pronouns will bear the same index, and the assignment function will simply assign the same value to their index (whether it corresponds to the sentence topic or not).[17] In (9), the two R–expressions do not depend on the assignment function for their interpretation; hence the condition (18) is irrelevant.

However, note that the reformulation in (18) still embodies the second qualification (b), concerning the liner position of the pronoun and the R–expression: the Topic Antecedent constraint is sensitive to directionality in an apparently irreducible way. This does not seem to follow from the view of coreference as a «static» relation. (18) is even more puzzling from the viewpoint of incremental left–to–right processing, in that a constraint on the assignment of a value to a free pronoun depends on the information structure status of an R–expression *following* it.

Let me summarize the preceding discussion in the following points.

We have observed that the possibility of backward anaphora depends on the interplay between two factors. The minimal pair (6) vs. (7) shows that the information structure status of the coreferring terms is relevant (new information focus or, alternatively, sentence topic status); on the other hand, the minimal pairs (6) vs. (8c) and (6) vs. (9) show the relevance of the intrinsic "referential content" of the coreferring terms (i.e. their pronominal vs. nonpronominal status). The first factor characteristically plays a central role in discourse–oriented approaches to anaphora, whereas the second factor is central to all the competition approaches to anaphora.[18] The data from backward anaphora presented above strongly suggest that both these factors are relevant, and that they interact in a very specific way.

At a purely observational level, it appears that coreference is unconstrained when the two co–referring terms are on the same level of the referentiality hierarchy, as in (8) and (9); it is only in case of an asymmetry that some constraint becomes operative. However, this constraint is subject to another asymmetry concerning the respective linear position of the less referential and the more referential term.

Anaphora is subject to the focus effect only when it is backward, namely, when the pronoun precedes the more referential R–expression.

I believe that this linear dimension of the phenomenon is worth emphasizing. We already know from work on inter–sentential anaphora that the linear – or rather, temporal – order of presentation of the elements is absolutely crucial.[19] Even an independent phenomenon like presupposition projection reveals a left–to–right asymmetry which is plausibly due to an incremental interpretation process. The preceding discussion suggests that this linear/temporal dimension is also crucial to intra–sentential anaphoric relations, which are more commonly conceived of in a static way. This point has been independently made by Shan & Barker (2006), Barker & Shan (2007) in their analysis of quantificational binding.[20]

4. Semantic vs. discourse anaphoricity

A different account of the contrast between (6) and (7) (repeated here) is found in Williams (1997: 588), who emphasizes the role of focus vs. destressing.

(6)   *His* wife loves JOHN.
(7)   *His* wife LOVES *John*.

In his words, "if the putative antecedent does not have main sentence stress, then the antecedent itself is anaphoric… really, *John* itself is dependent on something preceding, and *his* is dependent on it as well." [21] This means that there is no real backward anaphora: in cases like (7), the pronoun is actually linked to an antecedent in the previous linguistic context, rather than to the following R–expression.[22]

(7')   *John*…   *His* wife LOVES *John*.

In slightly different terms, this amounts to the following hypotheses. First, a "backward" pronoun implies an antecedent in the previous discourse. Second, the focussed R–expression in (6) must be incompatible with the presence of such an antecedent. The first hypothesis is fully consistent with the idea, defended above, that anaphora must be thought of in terms of a left–to–right incremental interpretation process. From this perspective, the leftward pronoun *his* could not inherit its referent from the following R–expression *John*, for the simple reason that the interpretation of the pronoun *precedes* the interpretation of the R–expression.

But what about the second hypothesis? Williams (1997) derives it from his view of anaphoric destressing. The destressed R–expression in (7) is dependent on a previous discourse antecedent; in turn, when the R–expression is focussed, as in (6),

its discourse referent cannot have been already introduced in the previous discourse; therefore, the "backward" pronoun has no possible (preceding) antecedent.

However, Williams's claim that the destressed proper name in (7) is anaphoric requires some qualification. It is widely assumed that proper names are directly referential: so how could they possibily be anaphoric? Anaphora is generally conceived of as the property of *inheriting* a referent from an antecedent.

I think that it is necessary here to clearly distinguish two notions of anaphoricity: strict semantic anaphoricity and discourse anaphoricity. The former corresponds to the way in which a nominal expression comes to denote: pronouns are strictly anaphoric in that they must either undergo semantic binding or be assigned a value that they lack inrtrinsically (let's say, via a context–determined assignment function); on the contrary, R–expressions are either inherently endowed with a fixed referent (proper names)[23] or come to denote an individual through semantic composition (definite descriptions). From this perspective, the proper name *John* in (7) cannot be strictly anaphoric in the same way that a pronoun is:[24] it is intrinsically "more referential" than a pronoun.[25]

On the other hand, the notion of discourse anaphoricity is much broader: any noun phrase that is preceded by a coreferential term is discourse anaphoric. Unbound pronouns (apart from deictic uses, to which I return in § 6.1 below) are discourse anaphoric,[26] but definite R–expressions can be as well: in fact, a definite R–expression can alternate with a pronoun when its referent has already been linguistically introduced (this corresponds to Roberts's (2003) strong familiarity).[27]

(19) There's *a doctor* in our little town. *He/The doctor* is Welsh.
<div align="center">(Roberts (2003), (8))</div>

The notion of discourse anaphoricity is thus relational, and does not exclusively depend on the intrinsic content of the noun phrase in question.

There is an apparent tension between semantic and discourse anaphoricity: definite R–expressions are semantically non–anaphoric, but they can be discourse anaphoric, as in (7) or (19): they independently denote a referent, but this referent may be already available in the previous linguistic context. This tension dissolves once it is recognized that the two notions of anaphoricity pertain to different levels of organization of language. An R–expression, being strictly non–anaphoric, always «brings with it» a referent in the semantic computation of the clause containing it. At the same time, an R–expression may or may not introduce its referent at the discourse level, as a potential antecedent for future discourse anaphora (cf. the forward–looking centers of Centering Theory). In the following section I will try to make these ideas a bit more precise, and I will show how they can derive the above observations about backward anaphora, essentially implementing a version of Williams's insights.

One more clarification is in order. In the following dicussion, I will continue to describe anaphoric relations by the usual terms of introducing and retrieving/picking up a referent. However, Safir (2004b, 2007) has argued that the type of anaphoric relations that may link one noun phrase to another are much more varied and complex that this. The common core of all anaphoric relations is the existence of a dependency between an antecedent and an expression whose interpretation somehow depends on the interpretation of that antecedent. The mechanism informally sketched below might be conceived of as a way of implementing dependent identity in Safir's broader sense.

5. Sketch of a solution

5.1. Local vs. global memory registers

The preceding discussion suggested that anaphora is constrained by two levels of organization: the local level of semantic computation of truth conditions, and a global discourse level. I contend that the two levels, and the two relevant notions of anaphoricity, must remain distinct: thus, an R–expression can never be strictly anaphoric, even though it can be discourse anaphoric; conversely, a (non–deictic) pronoun will always be strictly anaphoric, even in case of inter–sentential «accidental coreference».[28]

If both levels are relevant to anaphora, we have to assume that at both levels some mechanism keeps track of noun phrase referents. For concreteness, I adopt an informal version of the mechanism proposed by Schlenker (2005): in the course of the incremental interpretation of a sentence, the referents of noun phrases are dynamically stored in a memory register. This can be conceived as an ordered list, or a sequence, of (discourse) referents, which can be retrieved as the value of anaphoric pronouns. The precise details of the storage and retrieval operations are orthogonal to the main proposal; in the Appendix, I adopt for concreteness the sequence of evaluation and the negative indices proposed in Schlenker (2005).

In addition, I assume that, when the interpretation of the sentence is completed, the referents are passed up to a discourse–level memory register, where they become available for subsequent inter–sentential anaphora. I will dub these the *local* and the *global* register, respectively.[29]

The notion of strict anaphoricity defines the way in which the interpretation of an NP affects the memory registers:

(20)   a. A non–strictly anaphoric NP (= R–expression or deictic pronoun) adds its referent to the memory register(s).

b. A strictly anaphoric (i.e. non–deictic) pronoun must retrieve a referent contained in some accessible register.

(21) a. For any noun phrase contained in X_i, the local register of X_i is accessible.
b. For any noun phrase contained in X_i, the global register is accessible.

Note that by (21), a pronoun can retrieve a referent either from the local register or from the global one. I return to the nature of the local registers (21a) in § 5.4. In the following two subsections, I first discuss the role of the global register.

5.2. The role of focus

Let us go back to the notion of discourse anaphoricity, and to the crucial contrast between new information focus and destressing in (6)–(7). I assumed above that in (6), focus on the rightward R–expression is incompatible with a coreferential antecedent in the previous context. Why is this the case? Recall that on Williams's view, the presence of an antecedent necessarily implies anaphoric destressing, and hence, the lack of anaphoric destressing in (6) in turn implies the absence of such an antecedent. Like Williams, I assume that the presence of the "backward" pronoun implies a previous antecedent, but I suggest that the the focussed R–expression in (6) re–introduces the same discourse referent in the rheme, as new information, and this is an illegal move.[30]

I will implement this idea by assuming that focus vs. destressing constrains the way in which a non–strictly anaphoric NP affects the global register. As a qualification to (20a), I propose the following condition:

(22) a. A non–strictly anaphoric NP introduces an element in the global memory register iff it bears new information focus.
b. A non–focussed (destressed) NP does not affect the global register.

By (22), the global register only inherits from the local semantic computation of a sentence what is contained in the rhematic part of it.

Finally, I adopt the principle of Non–Redundancy proposed by Schlenker (2005):

(23) Non Redundancy: No element may be introduced twice in one and the same memory register.

Note that (22) and (23) jointly require destressing of discourse–anaphoric R–expressions: after a given referent has been introduced in the global memory register, all subsequent co–referential R–expressions are barred from bearing new

information focus; if they did, by (22) they would re–introduce their referent in the global register, giving rise to a violation of Non Redundancy.

5.3. Global violations of Non–Redundancy

At this point, we have the basic ingredients to reconsider the backward anaphora configurations in (6)–(9). For the time being, I exclusively concentrate on the management of the discourse–level register: in a nutshell, I propose that the impossibility of coreference in (6) results from a violation of Non–Redundancy at this global level. To see how this result is derived, consider the left–to–right interpretation of example (6):

(6) (Context: As for *John*, who does *his* wife really love?)
    ?* *His* wife loves *JOHN*

The backward pronoun *his* has no preceding antecedent in the noun phrase containing it, nor in the matrix clause containing that noun phrase; therefore, its referent must be retrieved from the global memory register. When the coreferential object R–expression is processed, it introduces the same referent in the local register of the matrix clause; but since the R–expression is focussed, by (22a) it also reintroduces its referent in the global memory register, violating Non–Redundancy.
    Example (7) is minimally different:

(7) (Context: –As for *John*, I believe *his* wife hates him. –You're wrong:)
    His wife LOVES *John*.

Here too, the pronoun *his* must retrieve its referent from the global register; however, the following R–expression is destressed and, by (22b), it only introduces its referent in the local register of the matrix clause, but not in the global one. Thus, Non–Redundancy is fulfilled.
    Recall now that (8c), involving two coreferential pronouns, differs from (6) in that focus does not block coreference:

(8) (Context: Who does his wife really love?)
    c. *His* wife loves *HIM*.

The reason for this contrast is that pronouns are strictly anaphoric (cf. (20b)), and never introduce a referent in the memory registers. The two pronouns in (8c) retrieve the same referent from the global register, and the focus condition (22) has no effect.

The examples (8a–b) are entirely parallel, except for the fact that they involve first and second person pronouns:

(8)   (Context: Who does your/my wife really love?)
       a. *My* wife loves ME.
       b. *Your* wife loves YOU.

I adopt Schlenker's (2005) proposal that the referents of first and second person pronouns, i.e. the speaker and the hearer/addressee(s), are always automatically available since they are given by the mere existence of the speech act. I differ from Schlenker in distinguishing global and local memory registers; specifically, I assume that these referents are contained in the global discourse register.[31] The two coreferential pronouns in (8a) and in (8b) retrieve the same referent from the global register (the speaker and the hearer, respectively).
    Let us now reconsider (9):

(9)   (Context: who does John's wife really love?)
       John's wife loves JOHN.

The minimal difference w.r.t. (6) is the fact that the first coreferential term is a destressed R–expression, rather than a strictly anaphoric pronoun. I propose that a destressed R–expression does not necessarily imply the previous availability of the relevant referent in the global register: by condition (22b), anaphoric destressing simply signals the invisibility of an R–expression for the global register. From the present perspective, in order for coreference to be allowed in (9), it must be the case that the discourse referent denoted by *John* is not contained in the global register at the point where the computation of (9) begins, although this name is indeed mentioned in the context question. (As a matter of fact, the previous occurrence of the name in the context question is destressed, and it cannot have introduced its referent in the global register.) Thus, the focussed direct object in (9) introduces this referent for the first time, in compliance with Non–Redundancy.[32]
    In sum, the set of assumptions in (20)–(23) provides an implementation of Williams's insight within the framework of an incremental left–to–right interpretation system.[33] The advantage of this implementation is that it allows us to see an affinity between the global discourse level of computation of anaphora and the "local" sentence level. To show this, I now return to the local memory registers. I will adapt Schlenker's analysis of binding, trying to make explicit the relevant notion of locality. (A more detailed presentation of my proposal based on Schlenker's formalization is deferred to the Appendix.)

5.4. Local violations of Non–Redundancy

In (21a), I have left open the definition of the local memory registers: what is the nature of X_i? Recall that by assumption, the local registers are used for the computation of reference and of truth conditions (e.g., the referent of *John* is used to compute the reference of the complex expression *John's wife* and the truth conditions of the sentence *John's wife is pregnant*).

The null assumption would seem to be that a local register is associated with each sentence. However, we can immediately see that this assumption wrongly predicts a violation of Non–Redundancy in the computation of (9), repeated here:

(9)   *John*'s wife loves *JOHN*.

By (20a), the two R–expressions would introduce twice the same referent in the sentence–level local register, violating Non–Redundancy; however, in (9) coreference is perfectly acceptable.

It is possible to solve this problem by assuming that the local memory registers are not sentence–level, but rather, they are local to smaller pieces of structure.[34] Let us adopt the following informal definition of semantic phases:

(24)   Every minimal noun phrase or clause is a semantic phase, i.e., a subpiece of a *top–down, left-to-right* semantic computation.

Note that phases, as subunits of a computation, are obviously contained in one another: in (9), the subject NP phase is contained the matrix clause phase. I will dub a nonmatrix phase a *subordinate* phase, and the immediately superordinate phase, the *containing* phase.

Let us assume, then, that the local memory registers are not associated with sentences, but rather, with semantic phases. Following Chesi (2004a, 2007), I assume that the local phase registers are linked to one another by the following inheritance relation:

(25)   a. A subordinate phase p_S inherits the content of the local register of the containing phase in the state it has reached when the computation of p_S begins.
         b. A containing phase, instead, does *not* inherit the content of the local register of a subordinate phase, but only the result of its overall semantic computation.

Let us now reconsider example (9) in the light of (25) (the square brackets indicate the phase containment relations):

(9)   [[*John*'s wife] loves *JOHN*].

The subordinate subject NP phase will have a separate memory register from that of the matrix clause, and by (25b), the latter will only receive the referent of the whole subject NP phase, *John's wife*, but not the referent of the subexpression *John*; hence, the subsequent processing of the object R–expression *John* will not cause a violation of Non–Redundancy in the matrix clause register. On the other hand, by (25a) the subordinate subject NP phase inherits the matrix clause register in the state it has reached when the processing of the subject begins: at this point, the matrix clause register does not yet contain the referent of the object R–expression. Thus, Non–Redundancy is satisfied by both the local registers.[35]

(24)–(25) can also account for the Principle C effects in (10), repeated here:

(10)   a. * [*John* claims [that Mary hates *JOHN*]]
       b. * [*John* claims [that Mary HATES *John*]]

Here the first R–expression is contained in the matrix clause phase, and the second R–expression is contained in the embedded clause phase. By (25a), the embedded clause register inherits from the register of the containing matrix clause the referent of the matrix subject R–expression: then, Non–Redundancy is violated at the local level of the embedded clause register, when the direct object *John* re–introduces the same referent which was inherited from the matrix clause phase. Focus vs. destressing is irrelevant here, because the focus condition (22) only concerns the introduction of a referent in the global register; however, a strictly non–anaphoric R–expression always introduces its referent in the local register, even when it is destressed. In other terms, a violation of Non–Redundancy at the local level is not sensitive to the information structure status of the R–expressions.

The inheritance relation defined in (25) also allows for an implementation of binding in terms of retrieval of a local antecedent:

(26)   [*John* loves [*his* wife]].

Here the subject R–expression introduces its referent in the local register of the matrix clause. The content of this register is inherited by the local register of the direct object noun phrase, *his wife*: the pronoun *his* can thus retrieve its referent directly from the local register, rather than from the global one.

One final clarification is necessary with regard to (10). As is well known, a Principle C effect also arises when the first coreferring term is a pronoun:

(10')   a. * [*He* claims [that Mary hates JOHN]].
       b. * [*He* claims [that Mary HATES *John*]].

This suggests that a pronoun retrieving its referent from the global register introduces that referent in the local register of the phase containing it:[36] the referent of the matrix subject pronoun is inherited by the embedded clause local register and yields a violation of Non–Redundancy in the embedded clause phase, when the object R–expression is processed. As the reader can easily verify, this assumption does not overgenerate in the case of (6)–(8), because the pronoun is embedded in the subject NP phase, whose local register is distinct from that of the matrix clause phase, which contains the coreferential object.

In sum, within a left–to–right oriented interpretation process the inheritance relation between phases essentially captures the effects of c–command (see Chesi 2004 for general discussion). This, however, only holds for local anaphora; as (21b) makes clear, the global register is always accessible from any phase.

5.5. Interim conclusions: local and global anaphora

 The discussion of (9) and (10) has emphasized the role of the computation at the local level, where some locality notion is relevant (here semantic phases, linked by an inheritance relation which replaces c–command). At this point, we can rethink Reinhart's divide between binding and «accidental coreference» in terms of local vs. global anaphora: a bound pronoun is a pronoun that retrieves its referent from the phase–local register (where it may have been inherited from a previous phase); a coreferential pronoun, instead, retrieves its referent from the global discourse–level register. However, the mechanism of storage and retrieval are the same in the two cases.

The local and the global registers stand in a quasi–hierarchical relation, since the global register is superordinate to all the phase–local registers.[37] On the other hand, despite the similarity of the storage and retrieval mechanism, these objects serve quite different purposes, and pertain to different levels of organization of the grammar. The local registers are essentially semantic, in that they serve the purpose of reference and truth assessment. The global register is, instead, a repository for potential discourse topics, and as such it is affected by discourse constraints that are instead irrelevant to the «local» semantic computation. One such constraint, I suggest, is the focus condition (22). This general idea is compatible with technically different views of a (structured) global repository of referents; from the present perpective, the only necessary constraint is that this register is intrinsically dynamic and compatible with a left–to–right intraclausal interpretation. Let me stress once

again that from this perspective, left-to-right orientation must be a property of the grammar itself, and not merely of processing, since grammatical constraints such as (23) and (25) are sensitive to it.

6. Further consequences and open problems

6.1. Deictic pronouns

In (20a) above, repeated here, I have assumed that deictic pronouns (those accompanied by some sort of ostension, which depend on the physical context of the speech act for their interpretation) differ from anaphoric pronouns in that they introduce a referent rather than retrieving it from some register.

(20)   a. A non–strictly anaphoric NP (= R–expression or deictic pronoun) adds its referent to the memory register(s).

In this respect, deictic pronouns are assimilated to R–expressions; hence, (22)–(23) predict that they should give rise to a focus effect in a backward anaphora configuration. This prediction seems to be correct: an example parallel to (8c), but with a deictic use of the pronoun, disallows backward anaphora, much as the focussed R–expression in (6).

(27) (Context: Among these men here, could you point out one that is such that his wife loves him?)
    * *His* wife loves ☛*HIM*!

6.2. Backward anaphora with a novel antecedent

As is widely assumed, indefinite noun phrase introduce novel referents, which are not available in the previous linguistic context. The present approach predicts the impossibility of backward anaphora with a novel antecedent. My introspective judgement supports this claim:

(28)   a. * *His* wife loves *a judge.*
        b. * *His* wife said that *a judge* was corrupt.

In (28), the non-deictic pronoun *his* must recover its referent from the global register, and the following indefinite expression reintroduces the same referent in it, violating Non–Redundancy.

Carden (1982: 367 ff.)[38] pointed out that in real text examples, backward anaphora often has an indefinite rightward antecedent:

(29)   When *she* was five years old, *a child of my acquaintance* announced a theory that she was inhabited by rabbits (Carden's (12)).

Here, the referent of the indefinite subject is (by hypothesis) not in the global register at the point when the pronoun is computed, and it cannot be recovered from the local register of the adjunct clause either. As stressed by Carden, this kind of example cannot be accounted for by the «forward–pronominalization only» approach, since there cannot be any antecedent preceding the pronoun.

However, note that there is something special about these examples: as stressed by Calabrese (1986: (43)), the indefinite antecedent must be in a highly prominent subject position (Calabrese's Thema position),[39] which has scope over a clause–initial adjunct phrase, though linearly following it:[40] as a matter of fact, the same configuration also allows for quantificational binding of a pronoun within the adjunct (Carden 1982: 368 ff.). A natural solution is to assume that the preposed adjunct clause can be reconstructed into a position that follows the subject, i.e., its interpretation can be «delayed» in the left-to-right interpretation. An account along these lines remains to be worked out. [41] [42]

We have seen above that (22) implies the necessary destressing of discourse–anaphoric definite noun phrases. Conversely, indefinite noun phrases are constrained in the opposite way: in order to introduce their novel referent in the global register, they have to bear new information focus. This is consistent with the well known observation that novel indefinites occur in the rhematic part of the clause.

6.4. Reinhart's contrast

Another piece of data to be accounted for is Reinhart's contrast in (17), repeated here:

(17)   a. When *he* entered the room, $^{\sqrt{}}$*Max* greeted $^{\sqrt{}}$*Bill.*
        b. When *he* entered the room, $^{\sqrt{}}$*Max* was greeted by **Bill.*
        c. # When *she* entered the room, Max was greeted by *Kora*. [43]

A relevant observation is that the *by*–phrase in (17c) belongs to the rhematic part of the clause, and hence, it introduces its referent in the global memory register, giving rise to a global violation of Non–Redundancy.

6.5. The strict/sloppy ambiguity

As is well known, Reinhart's distinction between binding and accidental coreference is usually invoked to account for the ambiguity between the strict and sloppy reading of pronouns contained in elliptical VPs, as in (13) above and in (31) below:

(31)   John loves his mother, and Bill does too.

Following Schlenker (2005: 31 ff.), it is possible to derive the sloppy reading of the elliptical VP if the pronoun contained in the latter retrieves its referent from within the (local) memory register of the second conjunct: this register contains the referent of the subject noun phrase *Bill*.[44] On the other hand, in order to derive the strict reading, Schlenker stipulates that «an elided pronoun may optionally introduce in the sequence the value that its antecedent did»: namely, the elided pronoun in the second conjunct may retrieve the same referent that the pronoun in the antecedent VP had retrieved from the local memory register of the first conjunct, i.e. John.[45]

This solution remains somewhat stipulative in Schlenker's approach, but it is interesting to note that it is equivalent to the analysis of strict readings recently proposed by Safir (2007b). According to Safir, the strict reading arises when the pronoun in the elided VP is identity–dependent on the parallel pronoun in the antecedent VP (32b), rather than on the local subject (32a). (Dependent identity is represented by means of arrows.)

(32) a. *John* loves *his* mother and ***Bill*** does [love ***his*** mother] too   (sloppy)
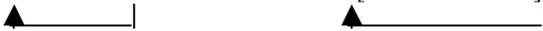
  b. John loves his mother and Bill does [love his mother] too   (strict)

Safir (2007b) argues that the strict reading cannot be analysed in terms of coreference: this is shown by the possibility of a strict reading of (33) (his (6)), where the strict interpretation of the elided pronoun involves a bound variable (cf. Heim 1998).

(33) Everyone thinks that he loves his pet and that St.Francis does too
      a. Every x, x a person (x thinks that (St.Francis y (y loves y's pet))) (sloppy)
      b. Every x, x a person (x thinks that (St.Francis y (y loves x's pet))) (strict)

This example is not problematic in Schlenker's account, as he assumes that a quantifier introduces a variable in the (local) memory register: this variable can be exceptionally retrieved in reading (b) by the same mechanism that allows the strict reading of (31).

It is fair to say that the exceptional retrieval of a non–local antecedent in the strict reading is at odds with the locality of memory registers which I have assumed (and which Schlenker also implicitly assumes). The two conjoined clauses of (31) have two separate local registers; thus, Safir's insight cannot be expressed by the mechanism of local anaphora. In the approach I have sketched here, the strict reading would rather have to be mediated by the global register.[46] The problem remains open (see the Appendix to Safir 2007b for extensive discussion of the strict reading).

6.6. Exceptional violations of Principle C

Since Non–Redundancy is as categorical as Principle C is, it inherits all the well–known problems for Principle C which were originally pointed out by Evans (1980). One case in point is the exceptional coreference in (33):

(33)   (We believed that everybody would vote for Joel, but in the end it turned out that…) only *JOEL* voted for *Joel*.

In this particular example, we may avoid a local violation of Non–Redundancy by stipulating that the complex phrase [*only Joel*] is a semantic phase distinct from the matrix clause, and hence the two occurrences of the R–expression introduce their referent in two distinct local registers. But the assumption is dubious, as the focussed R-expression in (33) need not be accompanied by a focussing adverb.

In this respect, the present proposal shares the rigidity of Reinhart's original principle, as well as that of Schlenker's (2005) proposal. Reinhart herself explained away Evans-type counterexamples by observing that in contexts like (33), the coreference option is not equivalent to the binding option (i.e., the two predicates $\lambda x.\mathbf{vote}(x,x)$ and $\lambda x.\mathbf{vote}(x, \mathbf{joel})$ give rise to distinct truth conditions), and argued that this is why Rule I does not apply. Along similar lines, Schlenker (2006) too has replaced Non-Redundancy with a non–categorical principle, Minimize Restriction, which is somewhat closer in spirit to the «competition» approaches to anaphora. (Safir (2004, 2007a) proposes another non-categorical principle, Expected Non–Covaluation.) Despite being highly marked, Evans-type counterexamples are generally accepted as grammatical in the literature, and remain a counterexample for the present Redundancy-based approach.

7. Concluding remarks

Although the plot of this paper may have seemed quite complex, it is actually very simple. By bringing together four empirical observations about backward anaphora, I have tried to convince the reader that a proper understanding of the phenomenon must take into account the interplay of two factors: the information structure status of the coreferring noun phrases, and their intrinsic «referential content». By reconsidering Reinhart's Topic Antecedent hypothesis, I have come to the conclusion that anaphora must be thought of in terms of a left–to–right interpretation process. I have then adopted Williams's insight about the focus effect, according to which there is no real backward anaphora; I have suggested that a focussed rightward 'antecedent' illegally reintroduces a referent which was already available at the point where the backward pronoun was processed. I have implemented this idea within a modified and informal version of the interpretation system proposed by Schlenker (2005).

Hardly anything in this note is entirely original; however, I hope that by combining insights from various sources, I have succeeded in making two points. First, it seems to me that the distinction between the semanticist's strict anaphoricity and the discourse theorist's discourse anaphoricity is not always clearly stated, which may lead to some misunderstanding of one another's viewpoint. If my line of reasoning is correct, in the domain of anaphora we have to recognize the coexistence of two levels of organization: the local level of the semantic computation, which I have argued is divided in phases, and a global discourse level, which is at the interface between semantics and pragmatics. Both levels are dynamic and incremental; the standard generative divide between syntactic binding and pragmatic coreference has been restated in terms of local versus global anaphora. Secondly, I hope I have made a convincing case for an incremental left–to–right semantic interpretation at the intraclausal level, as well as at the discourse level. This view of semantic interpretation supports an architecture of the grammar in which the syntactic and semantic computation proceed in parallel and, crucially, in a top–down, left–to–right fashion. From the syntactic viewpoint, this implies an incremental structure building along the lines of Phillips (2003) or Chesi (2004a,b, 2007). A left–to–right architecture incorporates the central insight that the temporal dimension is not just a fact about processing/performance, but it is intrinsic to the constitution of natural language grammars.

Appendix: an implementation based on Schlenker (2005)

A1. A summary of Schlenker (2005)

The sequence of evaluation represents the linguistic context w.r.t. which a constituent is evaluated. It includes
– those objects that are given by the mere existence of the speech act, i.e. the speaker and an addressee,
– the objects that have been linguistically introduced, i.e., which are the denotation of the terms that have been processed.

A sequence of evaluation represents *a state of a memory register*, which is constructed as a sentence is processed, top down, in accordance with the following rules:

i) R–expressions (= proper names, demonstrative pronouns and definite descriptions): when an R–expression is processed, its denotation is added at the end of the register. (Demonstrative pronouns bear a *positive* index).
ii) Non–demonstrative (= indexical and bound) pronouns: when a non–demonstrative pronoun is processed, some element of the register is recovered and moved to the end of the register, leaving an empty cell in the original position. Non–demonstrative pronouns bear *negative* indices, which indicate "how far back" in the sequence the denotation is to be found.
iii) The speaker/author and the addressee/hearer are always the first elements of any sequence; they are singled out by the superscripts ^{A} and ^{H} (e.g. $j^{A \wedge} m^H$)
iv) An n–place predicate is evaluated w.r.t. the last n elements of the sequence. (i.e., it is interpreted after all its arguments, in the lowest position of VP shells.)

(34) Ann runs.
$⟦\text{Ann run}⟧w\, j^{A \wedge} m^H = 1$ iff $⟦\text{run}⟧w\, j^{A \wedge} m^{H \wedge} a = 1$, iff $a \in I_w(\text{run})$

(35) *# Bill* likes *him$_{-1}$* : $j^A \wedge m^{H \wedge} b \rightarrow j^A \wedge m^H \#^\wedge b$
$⟦\text{like}⟧\, w, j^{A \wedge} m^{H \wedge} \#^\wedge b$: the interpretation of the 2–place predicate w.r.t. the last two cells is undefined (Principle B effect)

(36) Non–Redundancy: No object may occur twice in the same sequence of evaluation.[47]

Non–Redundancy is justified by the observation that the speaker and addressee – which are always given at the beginning of any sequence – cannot be referred to by an R–expression, which would reintroduce them in the sequence, by (i).

(37) a. # John is happy (said by John to Mary): $j^{A^\wedge}m^{H^\wedge}j$ (* Non–Redundancy)
      b. #Mary is happy (said by John to Mary): $j^{A^\wedge}m^{H^\wedge}m$ (* Non–Redundancy)

(38) a. # *Bill* likes *Bill* : $j^{A^\wedge}m^{H^\wedge}b^\wedge b$: (* Non–Redundancy: Principle C effect)
      b. # *He₁* likes *Bill*: $j^{A \wedge}m^{H^\wedge}b^\wedge b$: (*Non–Redundancy: Principle C effect)

Schlenker's analysis of backward anaphora configurations is as follows:

(39) [*Bill*'s teacher] likes *Bill.*

There is no violation of Non–Redundancy: «the key is that the VP *hates Bill* is evaluated under a sequence that contains Bill's teacher but not Bill himself, with the result that Non–redundancy is satisfied (p. 16)». This implies that the subject NP must be evaluated w.r.t. a different sequence from that of the matrix predicate.

(40) [*His* teacher] likes *Bill.*

Schlenker (note 15) argues that here *his* is a demonstrative pronoun, not an anaphoric one (from what sequence could it possibly recover its antecedent?)[48]

A2. A revision of Schlenker's system

I propose the following two modifications:
i) The semantic computation must be made sensitive to *phases*: each phase has a *local* sequence of evaluation.
ii) In addition, there is a *global* sequence of evaluation which keeps track of potential antecedents (actual referents or discourse referents); a pronoun can retrieve an element either from a local sequence (Reinhart's A–binding) or from the global one (Reinhart's accidental coreference).
    Let's see in which way exactly Schlenker's sequences must be local. Suppose we assume a single sentence–level sequence of evaluation:

(41) Bill likes [Ed's teacher]
      $j^{A^\wedge}m^{H^\wedge}b$ →
      $j^{A^\wedge}m^{H^\wedge}b^\wedge e$ →
      $j^{A^\wedge}m^{H^\wedge}b^\wedge e^\wedge d$ (d=e's teacher)
      $[\![like]\!]^w\ j^{A^\wedge}m^{H^\wedge}b^\wedge e^\wedge d$

By (iv) we obtain the wrong truth conditions (equivalent to: Ed likes Ed's teacher). Rather, the object DP must be evaluated w.r.t. a separate sequence *s*, and returns to the matrix sequence only the result of the evaluation (=d).

(41) $[\![\text{like}]\!]^w \; j^A {}^{\wedge} m^H {}^{\wedge} b {}^{\wedge} d$, where d= $[\![\text{Ed's teacher}]\!]^w \; s$

In (42), in order for the pronoun *his* to recover its antecedent, the sequence of the object DP must *inherit* what is contained in the matrix clause sequence:

(42)   Bill likes [his$_{-1}$ teacher].
       $[\![\text{like}]\!]^w \; j^A {}^{\wedge} m^H {}^{\wedge} b {}^{\wedge} d$, where d= $[\![\text{his}_{-1} \text{ teacher}]\!]^w \; j^A {}^{\wedge} m^H {}^{\wedge} b$

In (43), in order for a violation of Non Redundancy to occur, the object DP must inherit the content of the matrix clause sequence, which already includes b:

(43)   * *Bill* likes [*Bill*'s teacher]
       $[\![\text{like}]\!]^w \; j^A {}^{\wedge} m^H {}^{\wedge} b {}^{\wedge} d$, with d= $[\![\text{the } Bill \text{ teacher}]\!]^w \; j^A {}^{\wedge} m^H {}^{\wedge} b$ →
       $[\![\text{teacher}]\!] \; j^A {}^{\wedge} m^H {}^{\wedge} b \; {}^{\wedge} b$ (*Non-Redundancy)

On the contrary, in (44) the matrix clause sequence does *not* inherit what is contained in the subject DP sequence, so that Non–Redundancy is not violated:

(44)   [*Bill*'s teacher] likes *Bill*
       $[\![(17)]\!]^w \; j^A {}^{\wedge} m^H = [\![\text{like Bill}]\!]^w \; j^A {}^{\wedge} m^H {}^{\wedge} t =$
       $[\![\text{like}]\!]^w \; j^A {}^{\wedge} m^H {}^{\wedge} t \; {}^{\wedge} b$
       with t= $[\![\text{the Bill teacher}]\!]^w \; j^A {}^{\wedge} m^H$ (Schlenker 2005, 16)

In sum:
a) Sequences of evaluation must be local for any argumental constituent – minimally, CP and NP.
b) The sequences are linked by inheritance relations which capture the effects of c–command: the sequences of the embedded constituents inherit the content of the matrix sequence, but only transmit back to it the result of their semantic computation.
   Let us adopt Chesi's (2004) definition of top–down phases:

(45)   A phase is the minimal part of a top–down computation in which all the functional and selectional specifications associated to a given lexical phase head (minimally, N or V) are satisfied (cf. Grimshaw 1991).[49]

(46) i. Each CP or NP phase has a local sequence of evaluation.
ii. A subordinate phase n_a inherits the local sequence of the containing phase n, in the state it has reached when the computation of n_a begins; n_a does *not* transmit back its local sequence (with any elements added) to the superordinate phase n, but only the result of the overall computation.

A *global* register, accessible from any phase, stores the contextually available elements and keeps track of potential antecedents for inter–sentential anaphora:

(47) At any given point, the global register contains
i. the discourse referents for the speaker and addressee(s), singled out by the superscripts [A] and [H] (replaces Schlenker's (iii)),
ii. Salient familiar discourse referents, in the sense of Roberts (2003, 320–337), i.e., discourse referents pertaining to the immediate question under discussion in the hierarchical discourse structure of Roberts (1996).

Non–salient familiar discourse referents, as well as novel referents, get to be stored in the global register only if they satisfy the following necessary condition:

(48) Only an NP bearing new information focus can introduce an element in the global register.

Note that an NP will introduce a referent in the global register only if it bears a *positive* index. Anaphoric pronouns do not introduce anything, even when they bear focus.
Given (48), Schlenker's (i)–(ii) are modified as follows:

(i') An anaphoric pronoun copies an element from the global register into the last position of the local sequence of evaluation of the current phase, or retrieves an element from within the local sequence and moves it to the last position.
(ii') A focussed R–expression introduces an element both in the last position of the local sequence and in the global register.
(iii') A non–focussed (destressed) R–expression introduces an element in the last position of the local sequence, but not in the global register.

A3. Analysis of (6)–(10)

In (6), the possessive pronoun copies an element *g* from the global register into the local sequence of the NP phase. As this is subordinate in the matrix CP, no sequence inheritance obtains. However, when the focussed R–expression is

processed, by (ii') it introduces g in the local sequence of CP, and also in the global register: the latter violates Non–Redundancy.

(6)   * [$_{CP}$ [$_{NP}$ *His$_{-1}$* wife] loves *JOHN*]
       GLOBAL: $s^{A}{}^{\wedge}a^{H}{}^{\wedge}(m)^{\wedge}g$ (= g and possibly his wife m already introduced in previous discourse)
       NP: g (i')
       CP: m(iii')^g (ii')
       GLOBAL, at the end of computation: $s^{A}{}^{\wedge}a^{H}{}^{\wedge}m^{\wedge}g^{\wedge}g$(ii') (*Non–Redundancy)

The computation of (7) differs minimally: by (iii'), the destressed R–expression introduces g in the local sequence of the matrix CP, but not in the global register; hence, no violation of Non Redundancy arises.

(7)   [$_{CP}$ [$_{NP}$ *His$_{-1}$* wife] LOVES *John*]
       GLOBAL: $s^{A}{}^{\wedge}a^{H}{}^{\wedge}m^{\wedge}g$ (as above)
       NP: g (i')
       CP: m(iii')^g(iii') (global register unaffected)

In the examples (8), both the coreferring DPs are pronominal, and irrespective of focus, they copy an element from the global register into the local sequences of the DP and NP phases: these remain distinct, by (46ii), and Non–Redundancy is fulfilled.

(8) a. [$_{CP}$ [$_{NP}$ *My* wife] loves *ME*].
        GLOBAL: $s^{A}{}^{\wedge}a^{H}(^{\wedge}m)$
        NP: s^{A} (i')
        CP: m(iii')^s^{A} (i')

(8) c. [$_{CP}$ [$_{NP}$ *His$_{-1}$* wife] loves *HIM$_{-1}$*].
        GLOBAL: $s^{A}{}^{\wedge}a^{H}{}^{\wedge}g$
        NP: g (i')
        CP: m(iii')^g(i')

In (9), by (iii') the first, destressed R–expression merely introduces g in the local sequence of the subordinate DP phase, but it does not require that g be contained in the global sequence. The second, focussed R–expression introduces g both in the local sequence of the matrix CP and in the global sequence. At the local level, no sequence inheritance obtains between DP and CP, by computational nesting; at the global level, g is introduced only once. Hence, Non–Redundancy is fulfilled.

(9) $[_{CP} [_{NP} John$'s wife] loves JOHN]
        GLOBAL: $s^{A}{\wedge}a^{H}{\wedge}m$
        NP: g (iii')
        CP: m(ii'')^g (ii')
        GLOBAL, at the end of computation: $s^{A}{\wedge}a^{H}{\wedge}m{\wedge}g(ii')$

In (10a), the matrix subject introduces the element *g* in the local sequence of the CP1 phase. This sequence is inherited by the sequential complement CP2 phase, s–selected by the matrix V *claim*. When the lower R–expression is processed, it introduces in the local sequence of CP2 the very same element *g* that the matrix subject had already introduced: a violation of Non–Redundancy arises at the local level.[50]

(10)   a.* $[_{CP1} John$ claims $[_{CP2}$ that Mary HATES *John*]]
        GLOBAL: $s^{A}{\wedge}a^{H}{\wedge}(g)$
        CP1: g (iii')
        CP2: g(inherited from CP1)^m(ii'i)^g(iii')   (* Non–Redundancy)

References

Barker, Chris & Shan, Chung–chieh. 2007. Donkey anaphora is simply binding. Ms., NYU & Rutgers. semanticsarchive.net/Archive/ 2Y2ODI4Z/barker-shan-donkeys.pdf (8 November 2007)

Belletti, Adriana, Bennati, Elisa & Antonella Sorace. 2007. Theoretical and developmental issues in the syntax of subjects: Evidence from near–native Italian. *Natural Language and Linguistic Theory* 25, 657-689.

Bianchi, Valentina. 2001. Antisymmetry and the Leftness Condition: leftness as anti-c-command. *Studia Linguistica* 55, 1–38.

Bianchi, Valentina & Chesi, Cristiano. 2006. Phases, left branch islands, and computational nesting. *U. Penn Working Papers in Linguistics* 12.1, 15–28.

Büring, Daniel. 2004. Crossover situations. *Natural Language Semantics* 12, 23–62.

Büring, Daniel. 2005. *Binding Theory*. Cambridge University Press.

Büring, Daniel. 2005b. Bound to bind. *Linguistic Inquiry* 36, 259–274.

Burzio, L. 1991. The morphological basis of anaphora. *Journal of Linguistics* 27, 81–105.

Burzio, Luigi. 1998. Anaphora and soft constraints. In Barbosa, Pilar et al. (eds.), *Is the best good enough?*, Cambridge, MA, MIT Press, 93–113.

Calabrese, Andrea. 1986. Pronomina. *MIT Working Papers in Linguistics* 8, 1–46.

Carden, Guy. 1982. Backwards anaphora and discourse context. *Journal of Linguistics* 18, 361–387.

Carminati, Maria Nella. 2002. *The processing of Italian subject pronouns.* Doctoral dissertation, University of Massachusetts at Amherst.

Chesi, Cristiano. 2004a. *Phases and Cartography in Linguistic Computation.* Doct. diss., University of Siena*. (*http://www.ciscl.unisi.it/pubblicazioni.htm)

Chesi, Cristiano. 2004b. Phases and Complexity in Phrase Structure Building. *Proceedings of the 15th Meeting of Computational Linguistics in the Netherlands.*

Chesi, Cristiano. 2007. An introduction to Phase–based Minimalist Grammars: why *move* is Top–Down from Left–to–Right. In Moscati, E. (ed.), *Studies in Linguistics (CISCL Working Papers in Linguistics)* 1. CISCL, University of Siena.

Chomsky, Noam. 1976. Conditions on rules of grammar. *Linguistic Analysis* 2.4, 303–351.

Di Eugenio, Barbara. 1990. Centering theory and the Italian pronominal system. *COLING* 1990, vol 2, 270–275.

Elbourne, Paul. 2001. E–Type Anaphora as NP–Deletion. *Natural Language Semantics* 9, 241–288.

Erteschik–Shir, Nomi. 1997. *The Dynamics of Focus Structure.* Cambridge University Press.

Gordon, Peter C. & Hendrick, Randall. 1997. Intuitive knowledge of linguistic co–reference. *Cognition* 62, 325–370.

Grimshaw, Jane. 1991. Extended projection. Ms., Brandeis University.

Grosz, Barbara J., Joshi, Aravind K. & Weinstein, Scott. 1983. Providing a unified account of definite noun phrases in discourse. *Proceedings of the 21st annual meeting of the ACL,* 44–50.

Grosz, Barbara J., Joshi, Aravind K. & Weinstein, Scott. 1995. Centering: A framework for modelling the local coherence of discourse. *Computational Linguistics* 21, 203–225.

Hardt, Daniel. 2003. Sloppy Identity, Binding, and Centering. *Proceedings of SALT 13.* http://www.id.cbs.dk/~dh/publications.html (8 November 2007)

Hardt, Daniel. 2004. Dynamic Centering. *Reference Resolution and its Applications*, ACL–Barcelona. http://www.id.cbs.dk/~dh/publications.html

Heim, Irene. 1990. E–type pronouns and donkey anaphora. *Linguistics and Philosophy* 13, 137–178.

Heim, Irene. 1998. Anaphora and semantic interpretation: a reinterpretation of Reinhart's approach. *MIT Working Papers in Linguistics* 25, 205–246.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in Generative Grammar.* Oxford, Blackwell Publishers.

Von Heusinger, Klaus. 2004. Choice functions and the anaphoric semantics of definite NPs. *Research on Language and Computation* 2, 309–329.

Huang, Yan. 2000. *Anaphora: A Cross–linguistic Study*. Oxford: Oxford University Press.

Huang, Yan. 2007. *Pragmatics*. Oxford: Oxford University Press.

Kazanina, Nina. 2005. *The acquisition and processing of backwards anaphora*. Ph.D. dissertation, University of Maryland, College Park.

Kiss, Katalin E. 1998. Identificational Focus versus Information Focus. *Language* 74, 245–273.

Levinson, Stephen. 1991. Pragmatic reduction of the binding conditions revisited. *Journal of Linguistics* 27, 107–161.

Reinhart, Tanya. 1981. Pragmatics and linguistics: an analysis of sentence topics. *Philosophica* 27.1.(Also distributed by Indiana University Linguistic Club)

Reinhart, Tanya. 1983. *Anaphora and Semantic Interpretation*. Chicago: The University of Chicago Press.

Reinhart, Tanya. 1986. Center and periphery in the grammar of anaphora. In: Lust, Barbara (ed.), *Studies in the Acquisition of Anaphora*, vol I, 123–150.

Reinhart, Tanya. 2000. Strategies of anaphora resolution. In: Bennis, Hans, Everaert, Martin & Reuland, Eric (eds.), *Interface Strategies*, Amsterdam, North Holland.

Reinhart, T. 2004. Topics and the conceptual interface. In: Kamp, Hans & Barbara Partee (eds.), *Context Dependence in the Analysis of Linguistic Meaning*, Elsevier, 275–305

Reuland, Eric & Sergej Avrutin. 2005. Binding and Beyond: Issues in Backward Anaphora. In A. Branco, T. McEnery and R. Mitkov (eds.), *Anaphora Processing: Linguistic, cognitive and computational modelling.*, John Benjamins.

Roberts, Craige. 2003. Uniqueness in definite noun phrases. *Linguistics and Philosophy* 26, 287–350.

Safir, Ken. 2004. *The Syntax of (In)Dependence*. Cambridge, MA, MIT Press.

Safir, Ken. 2005. Abandoning coreference. In J.L. Bermudez (ed.), *Thought, Reference and Experience: Themes from the Philosophy of Gareth Evans*, Oxford, Oxford University Press.

Safir, Ken. 2007b. Strict readings. Ms., Rutgers University. Available at: http://ruccs.rutgers.edu/~safir/strict.pdf

Schlenker, Philippe. 2003. A plea for monsters. *Linguistics and Philosophy* 26, 29–120.

Schlenker, Philippe. 2005. Non–redundancy: towards a semantic reinterpretation of binding theory. *Natural Language Semantics* 13, 1–92.

Schlenker, Philippe. 2006. Minimize Restrictors! (Notes on definite descriptions, Condition C and epithets). In Emar Maier, Corien Bary & Janneke Huitink (eds.), *Proceedings of Sinn und Bedeutung 9*, 385–416. Nijmegen, NCS.

Schwarzschild, Roger. 1999. GIVENness, AvoidF and other constraints on the placement of accent. *Natural Language Semantics* 7, 141–177.

Shan, Chung–chieh & Barker, Chris. 2006. Explaining crossover and superiority as left–to–right evaluation. *Linguistics and Philosophy* 29, 91–134.

Sorace, Antonella & Filiaci, Francesca. 2006. Anaphora resolution in near – native speakers of Italian. *Second Language Research* 22, 339 – 368.

Williams, Edwin. 1997. Blocking and anaphora. *Linguistic Inquiry* 28, 577–628.

Zubizarreta, Maria Luisa. 1998. *Prosody, Focus and Word Order*. Cambridge, Mass., MIT Press.

Author's address: Dipartimento di Scienze della Comunicazione, Università degli Studi di Siena, via Roma 56, I-53100 Siena, Italy . E-mail: bianchi10@unisi.it

[1] The material in this paper was presented at the main session of the 30th GLOW Colloquium (Tromsoe, April 12-14th 2007), at the "Workshop on Romance-Japanese: Comparative syntax and language acquisition" (Siena, May 4-5th 2007), and in a seminar held at the Scuola Normale Superiore, Pisa (19 October 2007). I thank the audiences at those event for various important comments and suggestions. Thanks also to Mara Frascarelli, Ken Safir and Philippe Schlenker for making available to me some yet unpublished papers of theirs. This work was inspired by discussions on Weak Crossover with Michelangelo Falco, and especially by joint work with Cristiano Chesi on his top-down minimalist grammar (Chesi 2004, 2007); my warmest thanks to both of them. I am entirely responsible for any remaining errors.

[2] See Heim (1998), Büring (2005) for general discussion. See Büring (2007) for a recent overview of truth–conditional approaches to pronoun interpretation. An accurate and comprehensive presentation of the empirical domain can be found in Huang (2000).

[3] Here I gloss over the problem of coindexing: see Heim (1998) and Reinhart (2000) for relevant discussion. Even if the proper name and the pronoun in (2) bore the same index, this would not be a true binding configuration for lack of c–command.

[4] The c–command requirement is widely assumed, leaving aside E–type pronouns as a separate category (Heim 1990, Elbourne 2003, Büring 2004 among others). Two important alternatives are the Independence Principle (Safir 2004) and the Leftness Condition (see Shan & Barker 2006, Barker & Shan 2007): see note 10 below.

[5] The impossibility of coreference in c–command configurations of "Principle C violation" is accounted for by Reinhart's Rule I, to which I return below.

[6] This is not true of psycholinguistic and acquisition research: see, among many others, Gordon & Hendrick (1997), Kazanina (2005), Reuland & Avrutin (2005), Sorace & Filiaci (2006), Belletti et al. (2007). Note that the backward cases are the only cases where the lack of c–command is uncontroversial; for virtually all the forward cases, it is possible to propose a phrase structure analysis that satisfies the Bound Anaphora Condition. I refer the reader to Jackendoff (1990), Pesetsky (1995) and Bianchi (2001) for discussion (this issue is no longer so hot nowadays as it used to be in the Nineties).

[7] Chomsky (1976, 344) provided the following examples, with no clear indication as to the kind of focus involved:

(i)   The woman *he* loved BETRAYED *John*.

(ii) * The woman he loved betrayed JOHN.

[8] In the following examples, italics indicate coreference, small caps indicates stress.

[9] This is parallel to Schwarzschild's (1999) (47), which involves contrastive focus:

(i)   – John's mother praised Bill. – No, John's mother praised JOHN.

[10] As discussed by Lasnik (1989), in languages like Thai coreference is possible when an R–expression is c–commanded by another R–expression. This is a well known problem for Principle C, which will also be inherited by my proposal below (see § 6.5).

[11] On Reinhart's (1983) view, a Weak Crossover configuration simply violates the Bound Anaphora Condition: Pronoun binding can only take place from a c–commanding A–position (Büring 2004: 24, after Reinhart 1983). In (11), the A– position occupied by the variable fails to c–command the pronoun. Safir (2004, 2007a, 2007b) rejects the c–command requirement and proposes instead an anti–c–command requirement, the Independence Principle, whereby a

pronoun cannot depend on an antecedent if either the pronoun itself or some constituent containing it c–commands the antecedent. Chomsky's (1976) Leftness Condition may derive from the basic asymmetry of left–to–right interpretation, as argued by Shan & Barker (2006). A full development of this point exceeds the limits of the present discussion; for our purposes, any formulation of the Weak Crossover constraint may do, and I will stick to Reinhart's proposal in my critical discussion.

[12] This corresponds to the fact that in the context question of (8), the possessive pronoun is not bound by the wh–phrase.

[13] Reinhart (2000, (26)) reformulates Rule I as follows: α and β cannot be co–valued in a derivation D if (a) α is in a configuration to A(rgument)–bind B, and (b) α cannot A(rgument)–bind β, and (c) the covaluation interpretation is indistinguishable from what we would obtain if α A(rgument)–binds β. In configurations like (11), clause (a) does not hold (cf. Reinhart's (27), p. 16). Suppose instead that WCO is a possible (A')–binding configuration, although actual binding is excluded by an independent WCO constraint: then, the (2000) formulation of Rule I would correctly exclude coreference in (11). However, this assumption cannot account for the examples in (8)–(9): here too the rightward term is focussed, yet coreference is not ruled out.

[14] The notion of sentence–topic is that of Reinhart (1981), and it is defined in terms of an aboutness relation. Very roughly, the sentence topic indicates which "file card" in the conversational common ground will be updated by the propositional content of the asserted sentence. This notion has been adopted and elaborated by Zubizarreta (1998).

[15] See Erteshik–Shir (1998: 78 ff.) for an elaboration of this account.

[16] Note that this does not correspond to Reinhart's (1983) original view of unbound pronouns. See , again, Heim (1998) for relevant discussion.

[17] As for (8a–b), involving first and second person pronouns, it is controversial whether these should be assimilated to free variables (Schlenker 2003, 2004) or not (Safir 2004). If they are variables falling in the domain of the assignment function, they will be analysed as (8c), modulo additional presuppositions concerning the speaker/hearer role borne by the individuals that are assigned as values by the assignment function.

[18] The neo–gricean competition approaches reduce disjoint reference effects to the Q(uantity)–implicatures arising from the hierarchy of referentiality between anaphoric expressions. Burzio's (1991, 1998) competition approach is based on a morphological hierarchy of anaphoric expressions; Safir's (2004) Form–to–Interpretation Principle is based on the Most Dependent Scale. All these approaches are based on the same hierarchy (syntactic anaphor – pronoun – name).

[19] Cf. the notion of backward– vs. forward–looking center in Centering Theory.

[20] And, less conclusively, by Schlenker (2005) (cp. his § 7.2).

[21] The same observation is also found in Calabrese (1986: 14).

[22] Carden (1982) explicitly argues against this view, which he dubs the "forward [pronominalization] only" hypothesis and he defines "a theoretical garden path" (see also Reinhart 2004, § 14.3.1) Some of Carden's evidence, based on the analysis of a corpus of spontaneously produced discourse, will be discussed in §6.2. One real exception to Williams's view is the mystery story case discussed by Safir (2007: (49a)):

(i) *His* back was turned to us when we came in. *He* swivelled in *his* chair to face us. The penetrating eyes of *Count Marzipan* were trained upon us.

As noted by Safir, in these cases "a pronoun is introduced that we have no referent for and wer must wait for a plausible candidate to appear that supplies a value for the pronoun". However, Carden (1982: 372) agrees that such cases of "withheld antecedent" are qualitatively different from garden–variety backward pronominal-ization. (Similarly, in Gordon & Hendrick's 1997 experimental investigation of backward anaphora, the relevant examples are introduced without any context, hence, as the authors note, "the ability to co–referentially interpret a name and a preceding pronoun depends on a suspension of full interpretation of the initial pronoun".) Safir (2007) argues that backwards dependence is possible, but it is disfavoured by the principle Preferred Covaluation, which creates a tension until the point where the referential antecedent is introduced.

[23] Though speaking of a "fixed referent" is obviously an oversimplification, it will do for present purposes.

[24] This notion of semantic anaphoricity is not to be confused with the syntactic classification of noun phrases by the features [±anaphoric] and [±pronominal] in standard versions of Binding Theory.

[25] Elbourne (2005) however conceives of all pronouns as covert definite descriptions.

[26] Schwarzschild (1999) notes that pronouns can be focussed (not GIVEN) even though they are in some sense necessarily "not new". In my terms, non–deicic pronouns are strictly anaphoric and also discourse anaphoric; hence, Schwarzschild's GIVENness cannot be equated with discourse anaphoricity.

[27] As observed by Roberts (2003), anaphoric uses of definite descriptions are most felicitous when there is a certain distance between the definite description and the linguistic antecedent. See also von Heusinger (2004) for an analysis of anaphoric definite descriptions in terms of a global choice function.

[28] On this second point, I agree with Safir (2004, 2007a,b), who allows for inter–sentential dependencies, and I differ from Schlenker (2005), who analyses all free pronouns as demonstrative pronouns and assimilates them to R–expressions.

[29] Schlenker (2005) does not distinguish local and global registers.

[30] This idea was inspired by what Levinson (1991: 114) reports as "Bolinger's maxim": "Don't reintroduce (i.e. lexically specify) the topic of the theme in the rheme".

[31] In Schlenker (2005), the speaker and the hearer are introduced as the first elements of any memory register.

[32] This view of the problem entails a prediction: a sentence like (9) with a focussed rightward R–expression should be impossible in a context where there is a focussed previous mention. This prediction seems to be correct:
(i) ! I've had bad news about *JOHN*: *John*'s wife wants to break away from *JOHN*.

[33] David Adger (p.c.) points out the following example, with focus on the first R-expression:
(i) JOHN's wife loves John.
This seems to be an instance of contrastive focus. As stated in (22a), I assume that only new information focus allows an NP to introduce its referent in the global register: contrastive focus, instead, is not an appropriate way to introduce a new potential topic. Note that, since the second R-expression is destressed, it would not produce a violation of Non-Redundancy

even if the first R-expression were to introduce its referent in the global register. This last remark does not apply, however, to the following example, which I owe to another member of the audience at the 30th Glow Colloquium:

(ii) I don't know about Bill's wife, but JOHN'S wife loves JOHN.

Here again, the R–expression embedded in the subject is contrastive: possibly, this is not contrastive focus but a contrastive topic (Mara Frascarelli, p.c.); as a matter of fact, it is not part of the rheme, as required by condition (22a).

[34] This is also assumed by Schlenker (2005), who, however, does not make fully explicit the relevant notion of locality (cp. note 37 below). The definition in (24) is inspired by Chesi (2004a, 2004b).

[35] Cf. Schlenker (2005, 16):

(i) [*Bill*'s teacher] likes *Bill*.

"The key is that the VP *hates Bill* is evaluated under a sequence that contains Bill's teacher but not Bill himself, with the result that Non–Redundancy is satisfied." This implies that the subject noun phrase and the predicate are evaluated with respect to different sequences: see (44) below.

[36] The assumption that the retrieved referent is copied from the global register into the local one is consistent with the idea that the local registers are used for the computation of reference and truth conditions: see the Appendix for discussion.

[37] The global register is intended to model Roberts's (2003) strong familiarity. The notion of weak familiarity corresponds to a more general repository of referents – all those whose existence is somehow entailed by the Common Ground.

[38] See also Reinhart (2004: 296) for recent discussion.

[39] "(43) In an utterance U, a pronoun in U can have either the index of a referent in the Universe of Discourse of U, or the index of the Thema in U, where the Thema of U is the Thema of the main clause"; Calabrese (1986, (81)): "The T–domain of a Thema is given by the c–domain of its predication" (including preposed adverbial clauses).

[40] This is supported by the observation that at least in Italian, the relative position of the Thema and of a clause–initial adjunct can be freely reversed.

[41] This account predicts that in case of backward anaphora, a novel antecedent must be in Thema position. Notice that this position typically hosts specific indefinites, rather than nonspecific ones. Carden (1982: 366) also observes a substantial number of "first mention" definite antecedents, corresponding to Roberts's (2003) weakly familiar referents. I have tentatively assumed above that the global register only contains strongly familiar referents; weakly familiar antecedents may perhaps be included as well, if they are somehow made salient salient (see Roberts 2003, § 3 for extensive discussion of the salience requirement for pronouns' antecedents).

[42] Sorace & Filiaci (2006) and Belletti et al. (2007), in experimental studies of backward and forward anaphora interpretation by English near–native L2 speakers of Italian, observe the very strong tendency of backward subject pronouns to take the matrix subject as antecedent, and attribute it to an interface processing strategy: since backward anaphora has an intrinsic processing computational load, a processing strategy favors coreference with the first possible referent that is encountered, i.e. the matrix subject.

[43] Note that in the Italian equivalent of (17a), a null subject in the adjunct clause could only take the matrix subject as an antecedent. On the other hand, even an overt subject pronoun would disallow coreference with a *by*–phrase, in the Italian eqivalent of (17b–c). The relevance of the opposition between null and overt subjects is confirmed by the L2 acquisition studies by Sorace & Filiaci (2006), Belletti et al. (2007); see also Carminati (2002) and Frascarelli (2007) for relevant discussion.

[44] This proposal would also subsumes Reinhart's claim that the sloppy reading is only allowed in c–command configurations: this is because the sloppy reading requires the retrieval of an antecedent which is contained in the (local) memory register of the clause containing the sloppy pronoun. See however Hardt (2003) and Safir (2007b) for criticism of the c–command condition on the sloppy reading.

[45] Schlenker (2005, 31 ff.) argues that this solution also yields an account of Dahl's puzzle, whereby in a series of pronoun c–commanding one another, a strict reading of one pronoun blocks a sloppy reading of all the other pronouns c–commanded by it.

(i) Max said he saw his mother, and Oscar did too [say that he saw his mother].

a. Ok sloppy – sloppy: Oscar said that Oscar saw Oscar's mother
b. Ok strict – strict: Oscar said that Max saw Max's mother
c. Ok sloppy – strict: Oscar said that Oscar saw Max's mother
d. *strict – sloppy: Oscar said that Max saw Oscar's mother

The interpretation in (d) cannot be derived: if the first subject pronoun *he* in the elided VP of the second conjunct has a strict interpretation, the following pronoun *his* (with negative index −1) can only retrieve this same referent, and cannot refer back to Oscar.

[46] Such a move would not be unprecedented: Hardt (2003, 2004) proposed a Centering-based account of the strict/sloppy ambiguity, whereby the elided pronoun refers to the current Center: the strict reading results from Center retention between the first and the second conjunct (i.e., John is retained as the current Center in the second conjunct), whereas in the sloppy reading there is a Center shift from the first to the second conjunct (i.e.., Bill becomes the current Center in the second conjunct in the place of John). As mentioned above, the global register discussed here is conceptually similar to the list of backward and forward–looking Centers. (See Hardt 2003, § 5 for an analysis of Dahl's puzzle, mentioned in note 45.)

[47] Apparent violations of Non-Redundancy are due to different *guises*, which must count as different elements in the register (cf. Heim 1998). Schlenker (2006) also points out that his principle Minimize Restriction (replacing Non–Redundancy) can be violated when an R–expression serves certain pragmatic purposes (e.g. adding expressive/evaluative information) which could not be achieved by using a pronoun:

(i) *John* ran over a man who was trying to give *the idiot* directions.

[48] As a matter of fact, Schlenker reproduces in different terms Reinhart's dichotomy between bound and unbound pronouns. But in his system A–binding no longer involves introducing a variable in the antecedent position and co–binding of this position with the coindexed pronoun by the same λ–operator; hence, it seems implausible that A–binding and coreference should have such a radically different status as in Reinhart's account. This point is acknowledged by Schlenker (2006), who extends his analysis to anaphora across discourse.

[49] Chesi (2004a), Bianchi & Chesi (2006) draw a distinction between sequential and nested phases which is irrelevant for the present purposes. I will collectively call them *subordinate* phases, as opposed to the *superordinate* or containing phase.

[50] The same goes for (10b), where in addition, the lowest R–expression, being focussed, also re–introduces the element *g* in the global register.