UvA-DARE (Digital Academic Repository)

Anaphora resolved

Roelofsen, F.

Publication date
2008

Citation for published version (APA):
Roelofsen, F. (2008). *Anaphora resolved*. [Thesis, fully internal, Universiteit van Amsterdam].

Chapter 5

Resolution

This chapter presents a unified account of pronominal anaphora and VP ellipsis. The central assumption is that the meaning of anaphora is always retrieved from the context of use. In a slogan: anaphora are resolved. Moreover, it will be argued that constraints on the interpretation of anaphora follow from plausible assumptions concerning the resolution process and concerning the way people generally behave in communication.

5.1 Unification

The syntactic framework assumed in Part I did not exhibit any parallel between pronominal anaphora and VP ellipsis. In particular, VP ellipsis was taken to involve deletion, whereas pronominal anaphora was not. In the recent literature, several syntactic analyses have been proposed that do treat pronominal anaphora and VP ellipsis analogously (cf. Lyons, 1999; Elbourne, 2005b, and the references given there). I will abstract away from the details of and the differences between these individual proposals, and assume an analysis which, I think, captures the essential analogy in a most perspicuous way. I will assume that a pronoun is a determiner with an empty NP complement and that VP ellipsis involves a tense auxiliary with an empty VP complement.

(5.1)   The syntax of pronouns: (5.2)   The syntax of VP ellipsis:
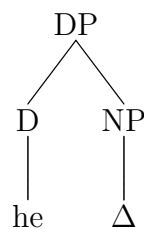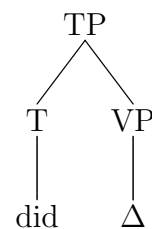
Henceforth, I will often refer to VP ellipsis as VP anaphora and to pronominal anaphora as NP anaphora. In fact, NP anaphora does not necessarily involve a pronominal determiner; many other determiners also permit NP anaphora, as illustrated by the following examples from Elbourne (2005b, p.45):

(5.3)   a.   Sue only bought two books. Mary bought at least three.
         b.   Most movies bore Mary, but she does like some.
         c.   Most MIT students build robots, and all watch Star Trek.
         d.   Many Athenians went to Sicily, but few returned.

Below, I will focus on pronominal NP anaphora, but the proposal should in principle be applicable to NP anaphora more generally.

Semantically, I will assume that pronouns are just like definite articles, apart from the fact that they encode ϕ-features (number, person, and gender). This assumption is rather common in the literature and has received strong support from cross-linguistic studies (cf. Lyons, 1999; Elbourne, 2005b).

5.2   Resolution

The crucial idea that I want to defend here is that the meaning of the empty constituents in (5.1) and (5.2) is contextually determined. Typically, the meaning of an empty NP/VP constituent is retrieved from an antecedent in the linguistic context. But it may also be retrieved from the non-linguistic context, or from an inferred antecedent. Furthermore, it is not necessarily retrieved from a single antecedent; semantic material from several sources may be combined, as long as the result is of the right semantic type. Let me illustrate this with a few examples. First, consider a case in which a pronoun is resolved to a definite description.[1]

(5.4)   The clown came in. He sat down.

The logical form of the second sentence contains an empty NP constituent Δ:

(5.5)   [the clown came in] [he Δ sat down]

The meaning of Δ may be retrieved from the noun phrase [clown]. If it is, we will say that Δ is resolved to [clown] and write:

(5.6)   [the clown came in] [he Δ sat down] $\Delta \rightarrow$ clown

Often, I will simply (and sloppily) say instead that [he] is resolved to [the clown] and write:

---

[1]Strictly speaking, I should say that this is a case in which the *complement* of a pronoun is resolved to the NP component of a definite description. This remark also applies to many cases discussed below.

(5.7)   [the clown came in] [he sat down]   he → the clown

Next consider a pronoun with an indefinite antecedent:

(5.8)   A clown came in. He sat down.

Again, the logical form of the second sentence contains an empty NP constituent, which may be resolved to the noun phrase [clown]:

(5.9)   [a clown came in] [he Δ sat down]   Δ → clown

Next consider a pronoun whose antecedent is a proper name:

(5.10)   a.   John is in good shape. I think he will win.
          b.   [John is in good shape] [I think he Δ will win]

The meaning of Δ may be retrieved from [John], but we need something extra here, because [John] denotes an individual (at least so we assumed), and Δ must denote a property. There are at least two possible ways to settle this mismatch: one is to assume that proper names do in fact not denote individuals but properties (cf. Burge, 1973; Larson and Segal, 1995; Elbourne, 2005b). The other is to maintain the assumption that proper names denote individuals, and to add the assumption that the use of a proper name N does not only make salient the individual x denoted by N, but also the *property* of being identical to x. This, then, is a property to which an empty NP constituent may be resolved.

In (5.10b) for example, Δ may be resolved to the property of being John. If it is, I will simply say that [he] is resolved to [John] and write:

(5.11)   [John is in good shape] [I think he will win]   he → John

I will assume this latter option here (a name denotes an individual, but also makes salient the property of being identical to that individual), but I am not strongly committed to it. It may turn out that names really should be analyzed as denoting properties. This would not have far-reaching consequences for the theory proposed here. It would merely simplify it.

Besides names and descriptions, I will assume that pronouns can also have traces as their antecedents. Consider:

(5.12)   a.   Every man thinks he will win.
          b.   [every man]^1 [t$_1$ thinks he Δ will win]

The meaning of Δ may be retrieved from that of [t$_1$]. Only, [t$_1$] denotes an individual, whereas Δ must denote a property. To settle the mismatch, I will assume that Δ may be resolved to the property of being identical to the individual denoted by [t$_1$], just as in the case of proper names. If Δ is resolved to the property of being identical to the individual denoted by [t$_1$], I will simply say that [he] is

resolved to [t_1] and write:

(5.13)   [every man]^1 [t_1 thinks he will win]   he $\rightarrow$ t_1

The meaning of a pronoun may also be retrieved from the non-linguistic context. For example, if I point at a certain athlete and say:

(5.14)   a.   He will win.
            b.   [he Δ will win]

Δ will generally be resolved to the property of being that athlete. There are situations in which Δ may be resolved to another property. For example, if we are watching a soccer game and I am explaining the rules of the game to you, I might point at one of the goalkeepers, say, Edwin van der Sar, and tell you:

(5.15)   a.   He is allowed to use his hands.
            b.   [he Δ is allowed to use his hands]

Then Δ is intended to be resolved to the property of being a goalkeeper and not to the property of being Edwin van der Sar.

    We have also seen cases in which the meaning of a pronoun is retrieved neither from an explicit linguistic antecedent nor from the non-linguistic context, but rather from an inferred antecedent. For example, in (5.16), Δ may be resolved to the property of being a baby.

(5.16)   [if I get pregnant, I will definitely keep it Δ]

Finally, let us consider donkey pronouns and paycheck pronouns. First, consider one of Geach's examples:

(5.17)   a.   Every farmer who owns a donkey beats it.
            b.   [every farmer who owns a donkey]^1 [t_1 beats it Δ]

For such cases I will assume, following Cooper (1979) and Heim and Kratzer (1998), that the meaning of Δ may be retrieved from several sources. For example, Δ may plausibly be resolved to the property of "being a donkey owned by t_1". Resolution to such a complex property may be forced here because resolution to a simpler property, such as that of being a donkey, would trigger the presupposition that there is a unique most salient donkey in the domain of discourse, which is not the case. Cooper (1979) observed that this strategy can also be applied to Karttunen's paycheck example:

(5.18)   a.   The man who gave his paycheck to his wife is smarter than
                  the man who gave it to his mistress.
            b.   The man who gave his paycheck to his wife is smarter than
                  the man [who]^1 [t_1 gave it Δ to his mistress]

Δ may be resolved to the property of being a paycheck of t_1. This yields exactly the intended truth-conditions.

I must remark here that a simpler treatment of paycheck pronouns may be given under the assumption that possessive DPs like [his paycheck] can have the following structure at LF:

(5.19)   [$_{DP}$ [$_D$ the][$_{NP}$ paycheck of him]]

This assumption is adopted by Heim and Kratzer (1998) and Elbourne (2005b), among others. Elbourne (2005b, p.82) points out that evidence for it can be found in work of Larson and Cho (1999), who analyze the ambiguity of DPs like [John's former house]. This phrase may denote either the object John owns that was formerly a house, or the object that was formerly a house owned by John. Larson and Cho take this ambiguity to be structural, depending on the order in which the elements of the possessive DP are semantically composed. If we first compose the meaning of [house] with that of [former], and then compose the result of this with the meaning of [John's], we get the object John owns that was formerly a house. If we first compose the meaning of [John's] with that of [house], and then compose the result of this with the meaning of [former], we get the object that was formerly a house owned by John. What is relevant for the analysis of paycheck sentences is that for [John's] to compose with [house] before [former] does, it must be in a low position at LF. Thus, there is some evidence that possessive DPs like [his paycheck] may indeed have a structure like that in (5.19).

If this is the case, then the resolution of paycheck pronouns is as straightforward as can be:

(5.20)   The man [who]^1 [t_1 gave [the [paycheck of him]] to his wife] is smarter than the man [who]^1 [t_1 gave [it Δ] to his mistress]

In the first clause, [him] and [his] are resolved to [t_1], and the empty NP in the second clause is then resolved to [paycheck of t_1].

So much for NP anaphora; let us now turn to VP anaphora. First, consider a simple strict/sloppy ambiguity:

(5.21)   Max called his mother, and Bob did too.

The logical form of the source clause is:

(5.22)   [Max]^1 [t_1 called his mother]

where [his] can be resolved either to [Max] or to [t_1]. The logical form of the target clause is:

(5.23)   [Bob]^1 [t_1 did Δ too]

Δ is resolved to the VP in the source clause. If [his] was resolved to [Max] we get a strict reading; if [his] was resolved to [t_1] we get a sloppy reading.

Just like empty NP elements, empty VP elements may also be resolved deictically (see examples (4.8)–(4.11)) or to inferred antecedents (see examples (4.12) and (4.32)–(4.39)). As noted in section 4.1, deictic resolution of VP anaphora is not as common as deictic resolution of pronominal NP anaphora. I mentioned two reasons for why this may be so. First, pronouns usually carry some (gender/number/person) information, which greatly facilitates deictic resolution. VP anaphora do not convey such information. Second, there are many more situations in which a particular object is the single most salient (female/singular/third person) object than there are situations in which a certain property or activity is the single most salient property or activity. Thus, VP anaphora are generally much harder to resolve deictically than pronominal NP anaphora. Incidentally, *non-pronominal* NP anaphora seem to pattern with VP anaphora in this respect: it is quite uncommon for non-pronominal NP anaphora to be resolved deictically, and this should be expected given that non-pronominal determiners do not encode ϕ-features.

Now let us consider the resolution of anaphora to inferred antecedents in somewhat more detail. I will assume that this process imposes a higher processing load on the hearer than resolution to non-inferred antecedents. Therefore, a hearer will only consider inferred antecedents if really necessary. This could be because the context does not provide any suitable explicit antecedents, or because resolving the anaphora to any of the given explicit antecedents yields an interpretation that is incoherent or inconsistent with world knowledge and/or contextual information. As observed by Hardt (2005), such a restriction on the use of inference in resolution is necessary to explain contrasts like that between Webber's original example, repeated here as (5.24), and the variant in (5.25).

(5.24)   Irv and Mary wanted to dance together, but Mary couldn't, because her husband was there.

(5.25)   Irv and Mary wanted to dance together, but Tom and Sue didn't.

The elided VP in (5.24) may be resolved to the inferred antecedent [dance with Irv], because there is no suitable explicit antecedent. In (5.25), on the other hand, there is a suitable explicit antecedent, namely the verb phrase [want to dance together]. Therefore, the inferred antecedent [want to dance with Irv] does not come into play.

This line of reasoning also yields a natural treatment of cascaded VP anaphora. Recall that typical cases of cascaded VP anaphora, such as (5.26), do not allow mixed readings, whereas some special cases, such as (5.27) and (5.28), do.

(5.26)   Bob called his mother, and Max did too. But Tom didn't.

(5.27)   Smithers thinks his job sucks, and Homer does, too. But Marge doesn't.

(5.28)   John didn't wash his car, but Bill did, even though Harry already had.

In cases like (5.26), the elided VP in the second clause is resolved to the explicit VP in the first clause, and the elided VP in the third clause is resolved either to the explicit VP in the first clause or to the copy of that VP in the second clause. In any case, we either get a strict reading for both elided VPs or we get a sloppy reading for both elided VPs. The examples in (5.27) and (5.28) are special, because they may trigger inference. Suppose, for example, that the pronoun in the first clause of (5.27) is resolved to the trace of [Smithers], and that the elided VP in the second clause is resolved, as usual, to the VP in the first clause. This means that the second clause attributes to Homer the property $[\lambda x.\ x$ thinks that x's job sucks]. Now consider the elided VP in the third clause. Normally, this VP would be resolved to the VP in the first clause or to the copy of that VP in the second clause. This would mean that the third clause assigns the above property to Marge. But this is inconsistent with world knowledge: Marge doesn't have a job. This triggers inference. From the information that Homer has the property $[\lambda x.\ x$ thinks that x's job sucks] it can be inferred that Homer has the property $[\lambda x.\ x$ thinks that Homer's job sucks]. This inference provides a suitable antecedent for the elided VP in the third clause, and yields exactly the attested mixed reading.

A similar story applies to (5.28). Here, a sloppy interpretation of the elided VPs in the second and third clause would give rise to an incoherent discourse. In particular, the contrast indicated by *even though* and *already* would not be established. This triggers an inference parallel to the one in (5.27), which in turn yields the attested mixed reading.

To the best of my knowledge, this is the first successful analysis of mixed readings in cascaded ellipsis. Theories based on VP Identity are all too rigid (they don't allow for mixed readings at all). Theories based on NP Parallelism (Fox, 1999a; Büring, 2005b) or unification (Dalrymple *et al.*, 1991) are too flexible (they allow mixed readings even in cases like (5.26)). The same goes for theories which assume that sloppy readings arise because pronouns in the antecedent VP may be reinterpreted at the ellipsis site (Hardt, 1999; Schlenker, 2005). Of course, this flexibility may be restricted by independently motivated constraints related, for instance, to information structure (Focus Match) or discourse structure (cf. Prüst *et al.*, 1994; Hardt and Romero, 2004). But such constraints won't be able to account for the contrast between (5.26) and (5.27), because there is no pertinent difference between these two cases as far as information structure and discourse structure are concerned. The idea that inference in anaphora resolution must be triggered really seems to be the only viable explanation.

Finally, let us turn to the unexpected sloppy readings observed by Wescoat and Hardt. Consider Wescoat's example:

(5.29)   The police officer who arrested John insulted him, and
            the police officer who arrested Bill did, too.

As noted above, the source clause of (5.29) is structurally analogous to one of
Geach's donkey examples:

(5.30)   a.   Every man who owns a donkey beats it.
            b.   [every man who owns a donkey]^1 [t$_1$ beats it Δ]

I assumed above, following Cooper (1979), that Δ may be resolved to the property
of "being a donkey owned by t$_1$" in this case. As observed by Tomioka (1999),
a similar strategy can be applied to (5.29). First consider the logical form of the
source clause:

(5.31)   [the police officer who arrested John]^1 [t$_1$ insulted him Δ]

Here, Δ may be resolved to the property of "being arrested by t$_1$". Next, consider
the logical form of the target clause:

(5.32)   [the police officer who arrested Bill]^1 [t$_1$ did Δ]

Now Δ can just be resolved to the relevant VP in the source clause, which can
be glossed as [insulted the person arrested by t$_1$]. This yields the attested sloppy
reading.

      Thus, all the cases of anaphora that were problematic for the theory proposed
in Part I are now dealt with in a rather straightforward way. Moreover, there
is no longer any need to stipulate a Semantic Identity constraint on VP ellipsis.
The fact that the meaning of an elided VP must be identical to the meaning of its
antecedent (in case it is not resolved deictically and there is no inference involved
in its resolution) simply follows from the way resolution works: the meaning of
the elided VP is *retrieved* from the meaning of the antecedent VP. As a result, the
meaning of the two VPs must—in non-deictic, non-inferential cases—be identical.

5.3   Anaphoric Relations

The distinction between inherently bound pronouns and inherently referential
pronouns has been dropped. All pronouns are assumed to be definite articles
with empty NP complements, the meaning of which is contextually retrieved.
However, depending on how a pronoun is resolved, we may still think of it as
being bound, cobound, covalued, or coreferential with another DP, in a sense
that is very much in line with the way in which these terms were used in Part I.
For example, covaluation can be defined as follows:[2]

_____

[2]Context sets, $\mathcal{F}$, $\mathcal{I}$, and equivalence given $\mathcal{F}$ and $\mathcal{I}$ are defined on pages 18, 19, and 31.

5.1. DEFINITION. [Covaluation]
Let C be a context and let s_c be the context set of C. Then, two expressions A and B are covalued in C iff for every $w \in s_c$, $[\![A]\!]^C(w)$ is equivalent to $[\![B]\!]^C(w)$ given $\mathcal{F}$ and $\mathcal{I}$.

Coreference can be seen as a special case of covaluation, namely the one involving only referential expressions (expressions of type *se* whose translation does not contain any free variables).

5.2. DEFINITION. [Coreference]
Two expressions corefer in a context C iff they are referential and covalued in C.

Binding may be defined as follows:

5.3. DEFINITION. [Binding]
A moved DP always binds its own trace. Moreover, if X is a logical form constituent, A a moved DP in X, B a pronoun in X, and C a context, then A binds B in LF/C iff:

 i B is covalued with the trace of A in C;

 ii A c-commands B in X;

 iii A does not c-command any other NP in X which satisfies i and ii.

This notion of binding is very similar to the one defined in Part I, and therefore also very similar to what Heim and Kratzer (1998) and Büring (2005a) call *semantic binding* and what Reinhart (2006) calls *A-binding*. The crucial difference is that the present notion is not defined in terms of indices, but rather in terms of covaluation between a pronoun and a trace.
Finally, cobinding can be defined in terms of binding:

5.4. DEFINITION. [Cobinding]
If X is a logical form constituent, A and B two nodes in X, and C a context, then A and B are cobound in LF/C iff there is a third node which binds both A and B in LF/C.

Let me illustrate these notions by means of a simple example:

(5.33)   [John]^1 [t_1 thinks he will win]

If [he] is resolved to [John] then [he] and [John] are covalued and even coreferential (because [John] is a referential expression). On the other hand, if [he] is resolved to [t_1] then (i) [he] and [t_1] are covalued (though not coreferential), (ii) [he] is bound by [John], and therefore (iii) [he] and [t_1] are cobound.

Thus, the familiar notions of anaphoric relatedness can be maintained, even though pronouns are no longer assumed to be either inherently bound or inherently referential. One consequence of this is that the ambiguities in (1.3), (1.4), and (1.5) can still be explained just as they were in chapter 1. A second consequence is that the formulation of Movement Economy in section 3.4 is still valid. Thus, we don't need a new account of crossover effects. Dahl's puzzle and Condition B effects, however, do require a new analysis.

5.4   Dahl's Puzzle

In section 4.6 it was pointed out that Dahl's puzzle requires two kinds of explanations. First, it should be explained why, in neutral contexts, across-the-board readings are preferred over mixed readings. Second, it should be explained why the sloppy-strict reading is easier to accommodate than the strict-sloppy reading.

I propose that the first issue, the preference for across-the-board readings, is due to a general preference for *local* resolution. The idea that such a preference exists is plausible given the incremental nature of the interpretation process, and the limited capacity of short-term memory. To see how this explains the preference for across-the-board readings, consider the source clause of Dahl's example:

(5.34) $[\text{Max}]^1$ $[t_1$ said that $[\text{he}]^2$ $[t_2$ called his mother$]]$

Assuming a preference for local resolution, the pronoun [his] will preferably be resolved to $[t_2]$ or to [he] rather than to $[t_1]$ or [Max].[3]   Thus, the preferred resolutions are:

(5.35)   a.   he $\rightarrow$ Max   c.   he $\rightarrow$ t_1
                his $\rightarrow$ he   his $\rightarrow$ he

           b.   he $\rightarrow$ Max   d.   he $\rightarrow$ t_1
                his $\rightarrow$ t_2 his $\rightarrow$ t_2

These resolutions all give rise to across-the-board readings of the elided VP in the target clause: (5.35a) and (5.35b) yield the strict-strict reading, while (5.35c) and (5.35d) yield the sloppy-sloppy reading. Hence, the preference for local resolution explains the preference for across-the-board readings in Dahl's puzzle.

The idea that resolution is preferably local is of course reminiscent of Fox's Locality constraint. But the two are really quite different. Locality is a grammatical principle, which classifies certain logical forms as ungrammatical. The local

---

[3]There may also be a slight preference for $[t_2]$ over [he] and for $[t_1]$ over [Max], but I will assume that this preference is negligible. In general, if two possible antecedents are directly adjacent, I will assume that the difference between them is too small to induce a noticeable preference.

resolution preference is an interpretive preference, which explains why certain interpretations of a given sentence are more accessible than others. Dahl's example is one case in which the two yield different predictions. Another example which is worth highlighting is the one discussed in section 3.2:

(5.36)   Every man said that he called his mother and that Bill did too.

This sentence has the following two readings:

(5.37)   a.   Every man x said that x called x's mother
             and that Bill called Bill's mother too.   [sloppy]
         b.   Every man x said that x called x's mother
             and that Bill called x's mother too.   [strict]

It was observed in section 3.2 that this example is problematic for Locality and many other accounts of Dahl's puzzle, because they all predict the strict reading in (5.37b) to be unavailable. This prediction does not follow from the local resolution preference. To see this, consider the following logical form of (5.36):

(5.38)   [Every man]^1 [t$_1$ said [[that [he]^2 [t$_2$ called his mother]] and

         [that [Bill]^2 [t$_2$ did Δ]] too]]

We are only interested of course in readings in which [he] and [his] are anaphorically related to [every man]. Thus, [he] must be resolved to [t$_1$], and [his] must be resolved to [t$_1$], [he], or [t$_2$]. Now, if resolution is preferably local, [his] will preferably be resolved to [he] or to [t$_2$] (and not to [t$_1$]). These two possibilities lead exactly to the strict and the sloppy reading in (5.37a) and (5.37b). This is another case, then, in which the local resolution preference makes different, and more desirable predictions than Locality.

Now let us turn back to Dahl's puzzle. It must still be explained why one of the mixed readings is easier to accommodate than the other. Both these readings are harder to get than across-the-board readings, but many people find the sloppy-strict reading significantly more accessible than the strict-sloppy reading. I propose the following account of this contrast. First, it should be noted that people generally need some time to decide whether the mixed readings are acceptable or not (the across-the-board readings are generally judged ok without much reflection). It seems that people use this time to try and figure out a specific context in which the reading they are considering is likely to be the intended reading. We could say that people try to find a context which *supports* the reading under consideration, where a context C is defined to support a reading R of a sentence S iff R is a likely reading of S in C. Now, in the case of Dahl's puzzle, it is relatively straightforward to find a context which supports the sloppy-strict reading. For example, if the question under discussion is:

(5.39)   Who called Max's mother?

then the sloppy-strict reading (Max said that Max called Max's mother and Bill said that Bill called Max's mother) is likely to be intended.

There are also contexts which support the strict-sloppy reading. One such a context was given in (2.33) in Part I, repeated here as (5.40):

(5.40)   a.   Did Max call everyone's mother?
          b.   Well, I don't know. . .
          c.   Max said he called his mother, and Bob did too.
          d.   But I haven't heard from Sue and Mary yet.

Other contexts supporting the strict-sloppy reading were given in (4.50) (Hardt's lawsuit case) and (4.51) (Reuland's gambling case). However, there are good reasons to believe that these contexts are much harder to evoke than contexts that support the sloppy-strict reading. The question in (5.39) is relatively simple: its logical structure can be represented as $?x.R(x,m)$ (read: *which x stand in relation R to m?*), where R is a simple relation, namely that of calling, and m is a simple individual, namely Max's mother. The question in (5.40a) is rather more complex. First of all, it is ambiguous between the two readings given in (5.41) and (5.42):

(5.41)   a. $?\forall x.R(\mathrm{max}, \mathrm{mother}(x))$
          b.   Is it true that for every x, Max called x's mother?
          c.   Possible answers: *yes*, *no*.
(5.42)   a. $\forall x.?R(\mathrm{max}, \mathrm{mother}(x))$
          b.   For every x, is it true that Max called x's mother?
          c.   Possible answer: well, Max called Max's mother, and he called
                Bob's mother, but I don't know whether he called Sue's mother
                and Mary's mother.

Only if the quantifier takes wide scope, as in (5.42), does the question license Dahl's sentence as a possible (partial) response. Notice, however, that there is a rather strong preference for the narrow scope reading in (5.41) over the wide scope reading in (5.42).[4] Moreover, even if the quantifier *is* given wide scope, it is unlikely that someone would use Dahl's sentence as a *complete* response to (5.40a) (the discourse in (5.40) becomes very odd if (5.40b) and (5.40d) are left out). Presumably, this is because *everyone* is unlikely to quantify just over Max and Bob. In any case, the relevant observation is that the question in (5.40a), and the way in which it may support the strict-sloppy reading of Dahl's sentence, is not nearly as straightforward as the question in (5.39), and the way in which

---

[4]The work of Groenendijk (2007) provides an interesting explanation for this preference. Roughly speaking, less inquisitive questions are generally preferred over more inquisitive questions (just as more informative assertions are generally preferred over less informative assertions) and the question in (5.41) is indeed less inquisitive than the one in (5.42) (a complete answer to the second question always entails an answer to the first, but not the other way around).

it supports the sloppy-strict reading of Dahl's sentence. Clearly, this observation also applies to Hardt's lawsuit case and Reuland's gambling case, where the strict-sloppy reading of Dahl's sentence is supported not just by a simple question, but rather by a whole plot.

I suggest, then, that in general, the level of accessibility of a particular reading for a given sentence will correlate with the complexity of the contexts that support this reading, and that, in particular, this is what explains the contrast in accessibility between the two mixed readings in Dahl's example.

5.5   Condition B Effects

Throughout the first part of this dissertation it was assumed, following Reinhart and many others, that Condition B effects are to be accounted for by two distinct mechanisms. The first accounts for Condition B effects on binding; the second accounts for Condition B effects on other kinds of codetermination. The primary piece of evidence in favor of such a two-level approach is Reinhart's observation that coreference is sometimes exceptionally allowed in Condition B environments. The relevant cases are repeated below:

(5.43)   Only Max himself voted for him.

(5.44)   I know what John and Mary have in common.
            John hates Mary and Mary hates her too.

(5.45)   If everyone voted for Oscar, then certainly Oscar voted for him.

However, as discussed in section 4.7, this judgment is disconfirmed by many informants. (5.43), (5.44) and (5.45) are generally felt to be ungrammatical on a coreferential reading, even though it is typically considered likely that such a reading is in fact intended.

It might be possible to formulate a two-level theory which accommodates this assessment of (5.43)–(5.45). But it wouldn't make much sense to do so. Reinhart's assessment of (5.43)–(5.45) was adduced as primary evidence for a two-level approach. If this assessment turns out to be inaccurate, the motivation for the whole approach goes up in smoke. This really concerns the approach in general, not just Reinhart's or anyone else's theory in particular. If there is no significant motivation for a two-level approach[5], we may as well pursue a

---

[5]It must be noted here that, apart from the alleged acceptability of coreference in constructions like (5.43)–(5.45), the two-level approach has also been supported by certain findings in the acquisition literature (cf. Chien and Wexler, 1990; Grodzinsky and Reinhart, 1993). More recently, however, the validity of these findings has been disproved quite convincingly (Elbourne, 2005a; Conroy *et al.*, 2007). Thus, I take the constructions in (5.43)–(5.45) to constitute the only alleged piece of evidence for a two-level approach to Condition B effects in English (see Heim, 2007; Grodzinsky, 2007; Conroy *et al.*, 2007, for concurrent views).

simpler, "one-level" explanation of Condition B effects in English. And such an explanation can indeed be given.

I think that the essential source of Condition B effects in English is the fact that speakers of English have come to use the marker *-self* to indicate that a pronoun should be resolved to one of its coarguments, and as a consequence, hearers have come to assume that, whenever a speaker does not use such a marker, interpretations that would result from coargument resolution are not intended.

I will refer to marked pronouns such as *himself* and *herself* as *self-pronouns*. Many authors refer to *himself* and *herself* as *anaphors*, following Chomsky (1981), or *self-anaphors*, following Reinhart and Reuland (1993). I have chosen not to adopt these terms for two reasons: (i) to avoid confusion with the term *anaphora*, and (ii) to remain neutral with respect to the claim that words like *himself* in English share some essential characteristics with words like *zich* in Dutch and *sig* in Icelandic, which are also called anaphors. Other authors refer to *himself* and *herself* as *reflexive pronouns* or simply as *reflexives*. I avoid these terms because *himself* and *herself* are not only used to indicate that a reflexive interpretation is intended. They are also used as intensifiers, marking prominence and contrast, possibly among other things (cf. Baker, 1995). In fact, in the history of English, the use of self-pronouns as intensifiers preceded their use as reflexivity markers (cf. König and Siemund, 2000).

I will assume that if a self-pronoun is used to mark reflexivity, it may be resolved either to one of its coarguments, or to the trace of one of its coarguments. Thus, self-pronouns may be interpreted as bound variables, but also referentially. In the literature, it is often assumed that self-pronouns can only be interpreted as bound variables. However, this assumption is problematic: it wrongly predicts that the question-answer pair in (5.46) below is incongruent, and that the sentence in (5.47) (adapted from Dalrymple, 1991) does not have a strict reading (saying that Bill's lawyer couldn't defend Bill against the accusations). These examples clearly show that self-pronouns cannot only be interpreted as bound variables, but also referentially.[6]

(5.46)   a.   Who evaluated John?
          b.   He evaluated himself.

(5.47)   Bill defended himself against the accusations because his lawyer couldn't.

---

[6]It should be remarked here that there are also cases of VP ellipsis involving self-pronouns which do not admit strict readings. For example:

(i)   John defended himself, and Bob did too.

However, the contrast between (i) and (5.47) can be explained on independent grounds (see Kehler (2002) and the discussion on page 127 below).

I will refer to any interpretation that results from resolving a (self-)pronoun to (the trace of) one of its coarguments as a *reflexive interpretation*. Finally, I will refer to the convention that speakers always use a self-pronoun if they intend a reflexive interpretation as the Reflexivity Convention.

5.5. DEFINITION. [Reflexivity Convention]
If a reflexive interpretation is intended, this is indicated by using a self-pronoun.

The Reflexivity Convention does not only account for Condition B effects on binding, but also for Condition B effects on other kinds of codetermination. An unmarked pronoun will never be interpreted as codetermined with one of its coarguments, because this would yield a reflexive interpretation, and such an interpretation could only be intended if the speaker had used a self-pronoun.

The idea that the Reflexivity Convention is the source of Condition B effects in English is strongly supported by the following two facts. First, in a broad range of languages, the existence of Condition B effects correlates with the existence of reflexivity markers (cf. Levinson, 2000; Huang, 2000). In particular, languages without reflexivity markers do not exhibit Condition B effects. Second, languages like English have gradually developed from an earlier stage, without reflexivity markers and without Condition B effects, to the current stage, with reflexivity markers and with Condition B effects (cf. Levinson, 2000; König and Siemund, 2000; van Gelderen, 2001; Keenan, 2002). The same development has been observed in several Creole languages (cf. Carden and Stewart, 1988; Levinson, 2000). Levinson (2000) provides a particularly attractive explanation of the crucial steps in this evolutionary process.

At an early stage, a language may not have any reflexivity markers and unmarked pronouns may freely be resolved to coarguments. This was the case, for instance, in Old English. However, there is a general tendency, even at such a stage, *not* to resolve pronouns to coarguments, for the simple reason that the agent and the patient of most actions are stereotypically distinct. Then, reflexivity markers gradually come into existence as "markers of the unusual": if a speaker intends a reflexive interpretation, he uses a marked construction (e.g., a self-pronoun) to signal to the hearer that something unusual is intended. This is an instance of what Horn (1984) called *the division of pragmatic labor*: unmarked forms are associated with stereotypical interpretations, while marked forms are associated with non-stereotypical interpretations. It should be noted that some verbs describe actions whose agent and patient are stereotypically identical (e.g., grooming verbs like *shaving* and *washing*). It should be expected, then, that a reflexive interpretation of such verbs does not necessarily involve special marking at this stage. This has indeed been observed, for example in Middle English (Faltz, 1985, p.242) and in Frisian (Reuland, 2001, p.478). Over time, though, the association between reflexive interpretations and reflexive marking becomes stronger and stronger and eventually leads to the Reflexivity Convention.

Levinson (2000) provides a wide range of cross-linguistic and diachronic data to support this hypothesis. Thus, the idea that Condition B effects in English stem from the Reflexivity Convention is well-motivated and well-supported.

Let us now return to the disputed Condition B effects in (5.43), (5.44) and (5.45). Two observations should be explained. First, these examples are generally felt to be ungrammatical on a reflexive interpretation. Second, however, informants often feel that a reflexive interpretation may nevertheless be intended.

The first observation is explained by the Reflexivity Convention. If a reflexive interpretation is intended, this should be indicated by a reflexive marker, and such a marker is not present in (5.43), (5.44) and (5.45).

There are several reasons why a reflexive interpretation may nevertheless seem to be intended in these examples. The case of (5.44) is relatively straightforward: the second sentence in (5.44) is supposed to convey what John and Mary have in common. If [her] is resolved to [Mary], there is indeed a property that is attributed to both John and Mary, namely that of hating Mary. If [her] is resolved in some other way, the sentence does not tell us which property John and Mary have in common. Therefore, it seems likely that a reflexive interpretation is intended, even though it is not properly expressed. An additional indication that this is the case is the use of the particle [too]. If [her] is resolved to [Mary], the use of this particle is justified, but if [her] is resolved in some other way, it is hard to see why [too] should have been used here.

The case of (5.45) is different but equally straightforward: only if [him] is resolved to [Oscar] does the sentence present a valid argument. If [him] is resolved in some other way, the sentence presents a nonsensical argument. Thus, a reflexive interpretation is probably intended, even though it is not properly expressed.

Example (5.43) is more subtle. I think that the crucial element here is not so much the focus-sensitive particle *only*, but rather the intensifier *himself*.[7] When confronted with examples like (5.43), informants quite often report that a reflexive interpretation is probably intended. But when confronted with examples like (5.48) (without intensifier), they don't.[8]

(5.48)   Only Max voted for him.

---

[7]Self-pronouns can be used as intensifiers in several ways. For instance, in (5.43) and in (i) below the self-pronoun is used as an *adnominal* intensifier, while in (ii) below it is used as an *adverbial* intensifier:

(i)   The President himself opened the exhibition.

(ii)   The President opened the exhibition himself.

I am only concerned here with adnominal intensifiers, to which I will simply refer as intensifiers.

[8]This contrast has, to the best of my knowledge, not been noted previously, perhaps because it cannot be accounted for by any two-level theory of Condition B effects.

Thus, there must be something about intensifiers that makes the reflexive interpretation in (5.43) particularly salient. Let me try to pin down what this is.

The standard analysis of adnominal intensifiers, due to Eckardt (2001) and Hole (1999) (see also Gast, 2006; Eckardt, 2006; König and Gast, 2006), is that they adjoin to DPs and denote the identity function on the domain of individuals. Thus, the denotation of *Max himself* is obtained by applying the identity function to the denotation of *Max*. Intensifiers, then, do not make any contribution to the ordinary semantic value of a sentence. However, they do make a significant contribution to the focus semantic value: intensifiers are always in focus (accented), and therefore, just like other focused elements, evoke a set of alternatives. These alternatives are contextually determined functions, other than the identity function. For example, in (5.49) the contextually triggered alternative function is the one mapping people to their family members and in (5.50) it is the one mapping kings to the members of their court.

(5.49)   John and his family are deciding where to spend their holidays. John himself wants to go to Greece.

(5.50)   The king himself opened the door.

Intensifiers interact with focus-sensitive particles like *only* just like other focused elements do. For example, (5.51) entails that John's family members do not want to go to Greece (see Eckardt (2001) for more illustrations of the fact that intensifiers behave just like other focused elements).

(5.51)   John and his family are deciding where to spend their holidays. Only John himself wants to go to Greece.

The crucial difference between intensified nominals and simply focused nominals is that the referent of an intensified nominal must be particularly *prominent* (Baker, 1995). This prominence may come from several sources. One possible source is world knowledge. For example, nominals like *the king* and *the President* refer to individuals who are prominent because of the role they play in society. Another possible source is the discourse. In particular, the prominence of a referent may be due to its being the discourse topic or the so-called subject of consciousness (the person whose perspective is taken in the discourse). In Baker's (1995) terms, the prominence of a referent may be justified either *externally* (i.e., by world knowledge) or *internally* (i.e., by the discourse).[9] The importance of this prominence-factor is illustrated by the following contrast:

---

[9]As observed by Baker, most examples of intensification in the linguistic literature involve nominals like *the king* and *the President*, such that prominence is justified externally. The prominence of intensified nominals "in the wild", however, is mostly justified internally, i.e., by the surrounding discourse.

(5.52)   a.   Eric Clapton is working on a new album with his band.
                 The members of the band are showing up at the studio
                 every morning around 9am.
                 Clapton himself usually joins them in the afternoon.
           b.   Eric Clapton is working on a new album with his band.
                 Most members of the band are showing up at the studio
                 every morning around 9am.
                 # The drummer himself usually joins them in the afternoon.

What is especially relevant for examples like (5.43) is that, if a sentence is considered in isolation, and if the prominence of an intensified nominal in that sentence is not justified externally (i.e., by world knowledge), then it is supposed to be justified internally (i.e., by the (missing) surrounding discourse). Let me illustrate. Suppose the second sentence in (5.49) is considered in isolation:

(5.53)   John himself would like to go to Greece.

The prominence of John is not justified externally, so it must be justified internally: the (missing) preceding discourse must be one in which John is particularly prominent, for example, one in which John figures as the discourse topic.

   Now let us turn back to example (5.43), repeated here:

(5.43)   Only Max himself voted for him.

This sentence tells us two things about the kind of discourse context in which it may occur. First, the use of *only* and the focus on *himself* indicate that the preceding discourse must be one in which, for some person p, the issue:

(5.54)   Who voted for p?

has been raised. This is the issue, then, that (5.43) addresses.

   Second, the use of the intensifier in (5.43) indicates that the preceding discourse must be one in which Max is particularly prominent. Given these two indications, the simplest assumption to make is that the discourse preceding (5.43) is one in which the following issue has been raised:

(5.55)   Who voted for Max?

But if this is the issue that (5.43) addresses, then the pronoun must be resolved to Max, and this yields a reflexive interpretation. This is, I think, the reason why informants sometimes feel that a reflexive interpretation might be intended in (5.43), even though it is not properly expressed.

   Thus, we have an explanation for why (5.43), (5.44) and (5.45) are felt to be ungrammatical on a reflexive interpretation (in terms of the Reflexivity Convention), but also for the fact that these sentences evoke the impression that a reflexive interpretation may nevertheless be intended.

It should be emphasized that *reflexive interpretations* are interpretations which result from resolving a pronoun to one of its coarguments *or* from resolving a pronoun to the trace of one of its coarguments. We could call these two kinds of reflexive interpretations *coreferential* and *bound*, respectively, hopefully without causing confusion. Now, when confronted with examples like (5.43) and (5.44), informants often feel that a *coreferential* reflexive interpretation may be intended, but they never feel that a *bound* reflexive interpretation may be intended. For example, (5.43) could possibly be supposed to mean that Max was the only one who voted for Max, but it can certainly not be supposed to mean that Max was the only "self-voter" (the only one with the property $[\lambda x.x$ voted for $x]$). This observation is accounted for by the explanations given above. In particular, the statement that Max was the only self-voter does not address the issue in (5.55), and being a self-hater cannot be the property that John and Mary have in common according to (5.44).

Finally, let me remark that the analysis of Condition B effects proposed here differs from that of Levinson (2000), even though Levinson's work provides much support for it. The crucial difference is this: the Reflexivity Convention says that whenever a reflexive interpretation is intended, this must be indicated by means of a self-pronoun. Levinson assumes that (i) a self-pronoun in argument position must be resolved to one of its coarguments, (ii) a self-pronoun is *more informative* than an unmarked pronoun, and therefore (iii) the use of an unmarked pronoun *implicates* that a reflexive interpretation is not intended (just as *some students passed the test* implicates that *not all students passed the test*).[10] The main problem with this proposal, in my view, is that Condition B effects are not *cancelable* in the way implicatures generally are. To see this, consider the contrast between (5.56) and (5.57):

(5.56) a. Some students passed the test.
        b. In fact, it is possible that all of them passed.

(5.57) a. John thinks that Bill voted for him.
        b. ??In fact, it is possible that John thinks that Bill voted for himself.

(5.56a) implicates that not all students passed the test. This implicature is canceled in (5.56b). It is a characteristic feature of implicatures that they are cancelable in this way. Thus, if Condition B effects are implicatures, as Levinson suggests, we should expect that they are cancelable too. Example (5.57) shows that this is not the case. Thus, although from a historical perspective it is plausible that the pragmatic inference patterns described by Levinson have played an important role in the realisation of the Reflexivity Convention, I don't think that they provide a suitable account of Condition B effects in present-day English.

---

[10]A summary of Levinson's proposal can be found on pp.347–348 of his (2000) book.

The issues raised in chapter 4 have now been resolved. A unified treatment of pronouns and VP ellipsis has been established. The stipulative Identity condition on VP ellipsis has been eliminated. Pronouns which could not be classified as either bound or coreferential, and instances of VP ellipsis which could not be dealt with in terms of VP Identity are no longer problematic. Dahl's puzzle has received a refined treatment. And finally, Condition B effects have been dealt with in a satisfactory way.

I now turn to a brief discussion of how the central ideas proposed here are related to previous and ongoing work of others.

5.6   Related Work

Early Ancestors.   In the early days of generative grammar, Wasow (1972) proposed a theory of anaphora that has remarkably much in common with the theory defended here. In particular, Wasow argued that pronominal anaphora and VP ellipsis should be treated in a unified manner, and that anaphora involves resolution rather than deletion. Another early proponent of resolution, especially for the case of VP ellipsis, was Williams (1977).

Diversification.   These ideas should have been standard ever since. But instead, much energy has been devoted to exploring several alternatives. As mentioned in section 4.1, one important reason for exploring such alternatives was the work of Hankamer and Sag (1976), who argued for a fundamental distinction between *deep anaphora* and *surface anaphora*. Pronouns were classified as deep anaphora and analyzed in terms of resolution, while VP ellipsis was classified as surface anaphora and analyzed in terms of deletion. I already pointed out that Hankamer and Sag's main argument has been refuted and that even Sag himself recently proposed that VP ellipsis should be dealt with in terms of resolution rather than deletion. Other authors who have argued for a resolution approach to VP ellipsis include Hardt (1993, 1999) and Kehler (2002). But the deletion approach is still quite widely adopted (cf. Heim and Kratzer, 1998; Merchant, 2001).

The other reason why many authors have departed from a unified theory of anaphora is the fact that Reinhart (1983) and many others have argued for a distinction between bound and referential pronouns, as discussed at length in Part I. Such a distinction does of course not permit a unified analysis of pronouns, let alone of pronouns and VP ellipsis.

Re-unification.   The move I made in this chapter was to replace the idea that pronouns are *inherently* bound or referential by the alternative conception that pronouns may *end up* either as bound variables or as referential expressions, depending on how they are resolved (e.g., to a trace or to a referential antecedent).

Heim and Kratzer (1998) made a similar move. That is, they also suggested that pronouns should not be treated as inherently bound or referential, but rather end up as bound variables in some contexts and as referential expressions in others.

Heim and Kratzer's implementation of this idea, however, is different from mine. I treat all pronouns as expressions whose meaning is to be determined contextually. In particular, a pronoun is interpreted as a variable iff it is resolved to a trace. Heim and Kratzer propose that all pronouns are treated as variables. These variables, then, may end up bound, or remain free, in which case they are interpreted as referring to some contextually salient entity.

My proposal has at least three advantages over Heim and Kratzer's. First, as Heim and Kratzer (1998, chapter 11) show in detail, certain pronouns cannot be treated as plain variables (examples of such pronouns are donkey pronouns and paycheck pronouns, see (4.29)–(4.31) above). Thus, Heim and Kratzer do not establish a completely unified analysis of pronouns. Second, pronominal anaphora are treated very differently in Heim and Kratzer's system from non-pronominal NP anaphora and VP ellipsis. The theory I have proposed treats all these kinds of anaphora in a unified manner. Finally, certain cases of VP ellipsis force Heim and Kratzer (1998, p.254) to stipulate an additional constraint on logical forms: "no LF representation must contain both bound occurrences and free occurrences of the same index". This constraint does not have any independent motivation. Indeed, it only arises because referential pronouns are embodied as free variables in Heim and Kratzer's system. On my proposal, it does not have to be stipulated.

Elbourne (2005b) elaborates on Heim and Kratzer's work in order to overcome the first two problems. He analyzes pronouns as definite articles whose NP complement is either an index or a full NP which is deleted at PF under identity with some other NP in the discourse. The indexed pronouns are translated as variables, which may end up either bound or free (referential), just as in Heim and Kratzer's system. Pronouns that cannot be analyzed as bound or referential, such as donkey and paycheck pronouns, are captured by NP-deletion. Thus, Elbourne establishes a uniform account of pronouns which is very much reminiscent of—and can indeed be unified with—a deletion approach to VP ellipsis.

The crucial difference with my proposal is that on Elbourne's view, certain pronouns have an indexical complement and others have a full NP complement which is deleted at PF, while on my view, all pronouns have an empty NP complement whose meaning is contextually retrieved.

One advantage of my proposal, then, is that it does not need to postulate the existence of indices as "lexical items". Another advantage has to do with the fact that resolution provides more freedom in the interpretation process than NP-deletion does. Elbourne argued that this freedom is problematic, but I will counter Elbourne's arguments below and show that the freedom provided by resolution is indeed needed.

NP-deletion versus Resolution.   Elbourne presents two arguments against resolution. The first is based on pairs of sentences like those in (5.58) and (5.59) (Elbourne, 2005b, p.64, originally from Heim, 1982, 1990).

(5.58)   a.   Every man who has a wife is sitting next to her.
           b. #Every married man is sitting next to her.

(5.59)   a.   Someone who has a guitar should bring it.
           b. #Some guitarist should bring it.

Elbourne's deletion theory predicts that the NP complement of a pronoun can only be deleted if there is an identical NP elsewhere in the discourse. Thus NP deletion is licensed in (5.58a) and (5.59a) but not in (5.58b) and (5.59b). Resolution is more liberal: empty NP complements can in principle be resolved to any salient property. Elbourne argues that (5.58b) and (5.59b) show that this is too unconstrained.

But I am not convinced. There are many examples which do require the freedom provided by resolution. Some such examples were discussed in section 4.3. One of these is repeated below in (5.60), and three additional examples are given in (5.61)–(5.63). (5.61) resembles (5.59b), and has already been noted in the literature at several occasions (according to Geurts, 1999, p.74, it dates back to Lakoff and Ross (1972)). (5.62) is designed to resemble (5.58b), and (5.63) is a similar 'real-life' example, taken from a website called *The Real Keys to a Happy Marriage* which, crucially, does not contain any occurrence of the word *husband*.[11]

(5.60)   If I get pregnant, I'll definitely keep it. (overheard in conversation)

(5.61)   John became a guitarist because he thought that it was a beautiful instrument.

(5.62)   Some men have been married for more than twenty years and still don't know what her favorite breakfast is.

(5.63)   If you don't know what his favorite movie is, you should plan to find out and watch it with him at the earliest convenience.

I do not have a very precise account for why resolution works much better in examples like (5.60)–(5.63) than it does in (5.58b) and (5.59b). It may be relevant that the conversational purpose of (5.58b) and (5.59b), considered in isolation, is far from clear. Thus, to make sense of (5.58b) and (5.59b) it is most naturally assumed that these sentences are part of larger discourse segments, and that the pronouns they contain refer to entities that are discussed not only in (5.58b) and (5.59b), but rather in those containing discourse segments. The conversational purpose of (5.60)–(5.63), considered as stand-alone utterances, is much clearer.

---

[11]`http://balancepdx.com/article_detail.php?article_id=32`

There is much room here for an improved account, but I don't think that such an account should make all too black-and-white predictions. Resolution may be much easier in some cases than in others, but there is a large gray area, with many gradations. A deletion theory such as Elbourne's appears to be much too strict in this respect.[12]

Elbourne's second argument is based on the following example:

(5.64)   In this town, every farmer who owns a donkey beats it, and the priest does too.

According to Elbourne (2005b, p.69), this sentence does not have a sloppy reading (which would say that the priest also beats the donkey he owns). The analysis of donkey pronouns proposed in section 5.2, which is essentially that of Cooper (1979), predicts that the donkey pronoun in the source clause can be resolved to the property of being a donkey owned by t_1. In the source clause, t_1 is bound by [every farmer who owns a donkey]. In the target clause, it may in principle be bound by [the priest], and this would give rise to the sloppy reading that Elbourne claims not to exist.

However, I think that sloppy readings should not be ruled out by the grammar in examples like (5.64). A sloppy reading is just somewhat implausible in this particular example. In many examples that are structurally analogous to (5.64), sloppy readings are readily available:

(5.65)   In this town, every farmer who has a spare room rents it out to tourists, and the priest does too.

(5.66)   Most men who own a car like to show off with it. But Peter doesn't.

This view has also been voiced by Maier (2006), who uses the following example to disprove Elbourne's claim:

(5.67)   Every male farmer who owns a donkey beats it, but farmer Mary doesn't.

Again, a sloppy reading is readily available here, contrary to the predictions of Elbourne's NP-deletion theory.

Very much related to these examples are Wescoat's and Hardt's examples discussed in section 4.4, repeated here:

---

[12]Jeroen Groenendijk pointed out to me that the pairs in (5.58) and (5.59) may in fact not be the right pairs to consider. The problem is that these pairs are not really *minimal*. Take the pair in (5.59). A much more minimal pair is the one in (i) below:

(i)   a. #Some guitar player should bring it.
       b. #Some guitarist should bring it.

Contrary to what the NP-deletion theory predicts, there is not much of a difference between (ia) and (ib), even though the former contains the NP [guitar] and the latter does not.

(5.68)   The police officer who arrested John insulted him, and the one who
          arrested Bill did, too.

(5.69)   If Harry has trouble at school, I will help him. But if John has trouble
          at school, I won't.

Elbourne (2005b, p.89–91) notes that his NP-deletion account of donkey pro-
nouns wrongly prohibits sloppy readings for these examples. In reaction to this
problem, he observes that sloppy readings do not only arise with pronominal NP-
deletion, but also with other kinds of NP-deletion. For example, sloppy readings
are available in:

(5.70)   The police officer who arrested some murderers insulted at least three,
          and the police officer who arrested some burglars did too.

Thus, Elbourne argues, the fact that examples like (5.68) and (5.69) allow for
sloppy readings does not show that his account of pronouns is wrong, but rather
that such an account must rely on a theory of NP-deletion that is flexible enough
to license the sloppy readings in question. However, Elbourne does not provide
such a theory of NP-deletion. In fact, throughout his book he assumes a theory of
NP-deletion which is based on LF identity, and the very point he wants to make
when discussing examples like (5.58b), (5.59b), and (5.64) is that such a strict
identity constraint on NP-deletion is necessary. How, then, would it be possible
to build in the flexibility that is apparently required to account for the sloppy
readings in (5.68) and (5.69)?

   As mentioned in section 5.2, Tomioka (1999) already observed that an analysis
of donkey pronouns à la Cooper (the one adopted here) straightforwardly explains
the sloppy readings in (5.68) and (5.69). Elbourne launches an argument against
such an analysis, based on the observation that the following variant of (5.68)
does not have a sloppy reading:

(5.71)   Every police officer who arrested a murderer insulted him,
          and Officer Jones did too.

I agree that a sloppy reading is highly inaccessible in this case, but I do not
think this should be explained on grammatical grounds. In fact, there is a very
plausible pragmatic explanation. Namely, on a sloppy reading, the second clause
of (5.71) would be completely *redundant*. It would convey information that is
already conveyed by the first clause. If the example is slightly changed to avoid
this redundancy, the sloppy reading reappears:

(5.72)   Almost every police officer who arrested a murderer insulted him,
          but Officer Jones didn't.

(5.73)   Every police officer who arrested a murderer insulted him.
          Even Officer Jones did.

Thus, Elbourne's arguments in favor of NP-deletion have all been countered. The flexibility provided by resolution appears to be necessary in general, although it may be constrained in certain specific cases by pragmatic factors. Elbourne's NP-deletion alternative is too strict, and there does not seem to be a straightforward way to add the necessary flexibility to it.

Complementary Theories. Of course, anaphoric mechanisms interact with many other linguistic mechanisms. Therefore, it should be expected that certain phenomena involving anaphora cannot be explained purely in terms of a theory of anaphora. Instead, a theory of anaphora must often interact with theories of other linguistic mechanisms in order to accomplish satisfactory explanations.

One important mechanism that interacts with anaphora resolution was already discussed in section 1.8, namely the encoding (and decoding) of *information structure*. The fact that Focus Match affects the resolution of VP anaphora should be seen as one particular consequence of this interaction. It is to be expected that there are many more such consequences, but these have, as far as I know, not yet been studied in much detail.

Another mechanism that interacts with anaphora resolution is the establishment of discourse coherence. This point has forcefully been made by Hobbs (1979), Prüst *et al.* (1994), Asher *et al.* (2001), and Kehler (2002), among others. To illustrate, I will consider some ways in which this interaction affects the resolution of VP anaphora, as described by Kehler (2002).

The first observation Kehler focuses on is that VP anaphora may exhibit a so-called *voice mismatch*. Sometimes, the target clause is in the passive voice, while the source clause is in the active voice, or the other way around. Kehler considers the following examples:

(5.74)   In March, four fireworks manufacturers asked that the decision be revised, and on Monday the ICC did.
(from an official document originally cited by Dalrymple (1991))

(5.75)   This problem was to have been looked into, but obviously nobody did.
(Vincent della Pietra, in conversation)

(5.76)   Of course this theory could be expressed using DRSs, but for the sake of simplicity we have chosen not to.
(from text of Lascarides and Ahser (1993))

(5.77)   Actually I have implemented the system with a manager, but it doesn't have to be.
(Steven Ketchpel, in conversation)

(5.78)   Just to set the record straight, Steve asked me to send the set by courier through my company insured, and it was.
(posting on the internet)

Similar examples can be found in (Dalrymple *et al.*, 1991) and (Hardt, 1993). The fact that VP ellipsis allows voice mismatches has been used by Dalrymple *et al.*, Hardt, and others to argue against the idea that VP ellipsis consists in PF deletion under a syntactic identity constraint (à la Sag, 1976), or in copying syntactic material at LF (à la Williams, 1977). Rather, they argue the resolution of VP ellipsis involves the recovery of *semantic* material.

However, this argument is problematic, because there are also cases of VP ellipsis that do not allow voice mismatches. Kehler gives the following examples:

(5.79)   #This problem was looked into by John, and Bill did too.

(5.80)   #This theory was expressed using SDRSs by Smith, and Jones did too.

(5.81)   #John implemented the system with a manager, but it wasn't by Fred.

Such examples could be used to argue exactly the opposite of what Dalrymple *et al.*, Hardt, and other proponents of a semantic approach argued, namely that the syntactic structure of the source clause *is* relevant for VP ellipsis.

Kehler shows us a way out of this impasse. He observes that there is a crucial difference between examples (5.74)–(5.78) on the one hand and examples (5.79)–(5.81) on the other. Namely, the kind of discourse relation between the source and target clauses in (5.74)–(5.78) is fundamentally different from the kind of discourse relation between the source and target clauses in (5.79)–(5.81). The clauses in (5.74)–(5.78) stand in a *Cause-Effect* relation, while the clauses in (5.79)–(5.81) stand in a *Resemblance* relation. Kehler argues on independent grounds that the establishment of Cause-Effect relations does not involve the reconstruction of syntactic material, while the establishment of Resemblance relations does. Thus, it is for the purpose of establishing discourse coherence (rather than for the purpose of resolving VP anaphora) that syntactic material must be reconstructed in (5.79)–(5.81) (and not in (5.74)–(5.78)). This explains why the voice mismatches in (5.79)–(5.81) are problematic, while the ones in (5.74)–(5.78) are not.

Similar observations can be made concerning VP anaphora with nominal antecedents. We have already seen some examples of this phenomenon in section 4.4. Kehler considers the following examples:

(5.82)   This letter deserves a response, but before you do, ...
          (attributed to Gregory Ward)

(5.83)   Today there is little or no official harassment of lesbians and gays by the national government, although autonomous governments might.
          (Hardt, 1993)

However, Kehler observes that VP anaphora with nominal antecedents is not always possible:

(5.84)   #This letter provoked a response from Bush, and Clinton did too.

(5.85) #There is unofficial harassment of lesbians and gays by the American government, and the Canadian government does too.

Again, the contrast can be explained in terms of discourse coherence establishment. The clauses in (5.82)–(5.83) stand in a Cause-Effect relation, and the establishment of such a relation does not involve the reconstruction of syntactic material. Thus, the VP anaphora may be resolved to nominal antecedents. The clauses in (5.84)–(5.85), however, stand in a Resemblance relation, and the establishment of such a relation does involve the reconstruction of syntactic material. This is why the VP anaphora in these sentences cannot be resolved felicitously to the nominal antecedents.

Yet another manifestation of the interaction between VP anaphora resolution and discourse coherence establishment surfaces when the antecedent VP contains a self-pronoun. For example, a reflexive interpretation is forced in the target clause in examples (5.86)–(5.87), but not in examples (5.88)–(5.89).

(5.86)   John defended himself, and Bob did too.

(5.87)   Fred voted for himself, and Gary did too.

(5.88)   John defended himself, because his lawyer couldn't.

(5.89)   Fred voted for himself, even though no one else did.

The clauses in (5.86)–(5.87) stand in a Resemblance relation. The establishment of a Resemblance relation involves reconstruction of syntactic material. Hence, the self-pronoun is reconstructed and forces a reflexive interpretation in the target clause. The clauses in (5.88)–(5.89) stand in a Cause-Effect relation, the establishment of which does not involve reconstructing syntactic material. Therefore, a non-reflexive interpretation is allowed in the target clause.

Incidentally, Kehler (2002, p.58) notes that there are certain borderline cases. For example:

(5.90)   The alleged murderer defended himself, and his lawyer did too.

(5.91)   Bush voted for himself, and his campaign manager did too.

Kehler reports that many of his informants find a non-reflexive interpretation of the source clause in these examples at least marginally acceptable, although a majority of them report that these sentences are not completely natural.

This seems to be the same kind of judgment that my informants reported when faced with examples like (5.92) and (5.93) (see section 5.5):

(5.92)   Only Max himself voted for him.

(5.93)   I know what John and Mary have in common.
          John hates Mary and Mary hates her too.

Kehler's assessment of (5.90) and (5.91) is indeed analogous to my assessment of (5.92) and (5.93). The fact that the elided VPs in (5.90) and (5.91) must be reconstructed in the process of discourse coherence establishment forces a reflexive interpretation of the target clause. But other factors (in this case world knowledge) strongly suggest that such a reflexive interpretation is not intended. As a result, informants generally feel that a non-reflexive interpretation is intended, even though it is not properly expressed.

It is to be expected that there is a variety of interactions between anaphora resolution and discourse coherence establishment, as well as other mechanisms, that remain to be explored. Such explorations, however, are left for future work.

Auxiliaries as proforms or pronouns as determiners?   The central ideas defended in this paper are (i) that NP anaphora and VP anaphora should be analyzed in a unified manner, and (ii) that the interpretation of anaphora primarily involves resolution. I have proposed a particular implementation of these ideas, but alternative implementations are possible of course. The most important feature of the implementation proposed here is that it assimilates the case of pronominal anaphora to the case of non-pronominal NP anaphora and VP anaphora by assuming that pronouns are determiners with an empty NP complement, and that it is really this empty NP complement whose meaning is contextually determined.

The alternative is to proceed the other way around, namely to assimilate the case of VP anaphora to the case of pronouns. This would mean that neither pronouns nor auxiliaries that figure in VP ellipsis have empty NP/VP complements. Rather, the pronouns/auxiliaries themselves are resolved. Such an alternative unified (and resolution based) analysis of pronouns and VP ellipsis has been proposed by Hardt (1999).

One reason for assimilating pronouns to non-pronominal anaphora, rather than the other way around, is that much work in syntax and typology supports the idea that pronouns are definite articles, usually with empty NP complements. One relevant observation, which dates back to Postal (1966), is that English pronouns actually have overt NP-complements in some constructions:

(5.94)   a.   we linguists
          b.   you troops
          c.   them guys (dialect)

A comprehensive argument, which involves data from many languages other than English, can be found in (Lyons, 1999).

A second reason to treat pronouns as determiners, rather than auxiliaries as proforms, is that this allows for a unified analysis not only of pronouns and VP ellipsis, but also of non-pronominal NP anaphora. To the best of my knowledge, a proform theory of non-pronominal NP anaphora has not been proposed yet, and I find it hard to imagine one.

Sloppy readings: binding or reinterpretation? Sloppy readings arise in several constructions. In the first part of this dissertation, we concentrated on focus constructions such as (5.95) and (5.96), and on elliptical constructions such as (5.97).

(5.95)   MAX called his mother.

(5.96)   Only MAX called his mother.

(5.97)   Max called his mother. Bob did too.

In the second part, we saw another construction in which sloppy readings arise, namely so-called paycheck sentences:

(5.98)   Max spent his paycheck. Bob saved it.

Notice that (5.98) is very similar to (5.97): in both cases the sloppy reading arises because an anaphoric expression is resolved to an antecedent which itself contains another anaphoric expression. In (5.97), the empty VP in the target clause is resolved to [called his mother], which contains the pronoun [his]. In (5.98), the empty NP in the target clause is resolved to [his paycheck], which contains the pronoun [his]. For ease of reference, let me call such constructions *embedded anaphora constructions.*

There are essentially two ways to account for sloppy readings in embedded anaphora constructions. The one originally suggested by Keenan (1971) and preserved in the present proposal assumes that the embedded anaphoric expression in the source clause can be interpreted as a variable. This variable, then, is bound by one element in the source clause and by another in the target clause.

The alternative is to assume that the embedded anaphoric expression in the source clause may be *reinterpreted* in the target clause. Such an approach is pursued by Hardt (1999) and by Schlenker (2005).[13]

My main reason to account for sloppy readings in terms of binding and not in terms of reinterpretation is that sloppy readings do not only arise in embedded anaphora constructions. The focus constructions in (5.95) and (5.96), for example, are not embedded anaphora constructions, but do exhibit sloppy readings. As we have seen, a theory in which anaphoric expressions may be interpreted as bound variables immediately accounts for this fact. A reinterpretation theory of sloppy readings does not.

Another problem with reinterpretation accounts of sloppy readings is that they are, in principle, too weak to rule out certain impossible readings of embedded anaphora constructions. For example, they do not rule out mixed readings in

---

[13]The theories proposed by Hardt and Schlenker are cast in a *dynamic* framework. But a reinterpretation account of sloppy readings could just as well be implemented in a static framework. And vice versa, a bound variable account of sloppy readings could just as well be implemented in a dynamic framework.

cascaded ellipsis. I argued in section 5.2 that such readings should generally be ruled out (see, in particular, example (5.26)). A related problematic prediction that is made by reinterpretation theories is that the second sentence in (5.99) has readings such as (5.100) and (5.101):

(5.99)   John didn't come to work. Bill called his boss, but Peter didn't.

(5.100)   Bill called John's boss, but Peter didn't call Peter's boss.

(5.101)   Bill called Bill's boss, but Peter didn't call John's boss.

It may be possible of course to devise additional machinery to rule out such readings. But if sloppy readings are accounted for in terms of binding rather than reinterpretation, no such additional machinery is needed. The right predictions are automatically generated.

Sloppy readings: additional support for a unified theory of anaphora. In (5.95)–(5.98), a sloppy reading arises because a pronoun is interpreted as a bound variable. If NP and VP anaphora are essentially the same, then we should also expect to find constructions in which a sloppy reading arises because an elided VP is interpreted as a bound variable. Such constructions do indeed exist. Kratzer (1991) discusses example (5.102), a VP analogue of (5.95):

(5.102)   I only went to TANGLEWOOD after you did.

Schwarz (2000) considers example (5.103), a VP analogue of (5.97):

(5.103)   When John had to cook he didn't want to.
           When he had to clean he didn't either.

Hardt (1999) considers a similar example, originally due to Carl Pollard:

(5.104)   I'll help you if you want me to.
           I'll kiss you even if you don't.

Hardt also provides a VP analogue of (5.98):

(5.105)   When Harry drinks, I always conceal my belief that he shouldn't.
           When he gambles, I can't conceal it.

Schwarz (2000) shows that these cases are satisfactorily accounted for by assuming that VPs can be raised, just like DPs, and that VP anaphors can be bound by (in my terminology: can be resolved to the trace of) a raised VP.

I take this as further support for the view that NP and VP anaphora should be analyzed in a unified way (see also Charlow, 2008).

Pointers to collaborative work. Finally, I would like to point to two projects that I have been working on in collaboration with others, and which bear some relation to the work presented here.

First, (Nesson, Roelofsen, and Grosz, 2008) investigates the mechanisms that underly the generation and interpretation of anaphoric referring expressions in discourse. In particular, we focus on certain phenomena that have previously been considered in the light of Centering Theory (Grosz *et al.*, 1995). Very much in the spirit of this dissertation, we develop a theory which is based on very general principles about human behavior in communication. The theory explains why the main generalization embodied by Centering Theory is generally very robust, but also why it systematically fails in certain cases.

A second piece of related joint work is (Roelofsen and Aloni, 2008). This paper is concerned with a completely different empirical domain, namely that of *concealed questions*. But there is an interesting connection: the mechanism that has been argued here to play a crucial role in the interpretation of NP and VP anaphora, namely the contextual resolution of a property that is not overtly expressed, turns out to play an essential role in the interpretation of concealed questions as well.

5.7   Summary

In this chapter, a unified analysis of NP and VP anaphora has been proposed. All the issues that were raised in previous chapters have been resolved. In particular, the stipulative Semantic Identity condition on VP ellipsis has been eliminated; pronouns which could not be classified as either bound or coreferential, and instances of VP ellipsis which could not be dealt with in terms of VP Identity are no longer problematic; Dahl's puzzle has received a refined treatment; and the disputed Condition B effects have been dealt with in a satisfactory way. It has been argued that the proposed account improves on existing proposals, in particular the deletion-based account of Elbourne (2005b) and the resolution-based account of Hardt (1999). Moreover, the proposed theory neatly ties in with theories of information structure and discourse coherence establishment.