Descriptive Pronouns Revisited

The semantics and pragmatics of identification-based descriptive interpretations

Thesis submitted to University College London in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Thiago Neuenschwander Galery

Research Department of Linguistics

Division of Psychology and Language Sciences

UCL

April, 2012

1

I, Thiago Neuenschwander Galery, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

_____

Signature

Aos amigos.

Abstract

This thesis confronts the semantic/pragmatic issues raised by identification - based descriptive uses of pronouns. The phenomenon, also known as deferred uses (Nunberg, 1993), arises when the correct understanding of a pronoun is dependent on the identification of a specific individual in the context that provides it with a descriptive (as opposed to a singular) interpretation. Moreover, the identification of the salient individual makes the interpretation available in a rather indirect way. For example, by pointing at a huge footprint in the sand and uttering 'He must be a giant', the speaker can convey the proposition that *the footprint maker must be a giant*, where the mental representation *footprint* (necessary for identification) and the representation *the footprint-maker* (the pronoun's interpretation) are not identical. These uses also display interesting properties when it comes to their ability to provide antecedents for other pronouns. As such, they are at the cross-road of many topics in philosophy of language and linguistics, including indexicality, anaphora, and figurative uses of language (metonymy). In this thesis, I propose that the data is best accounted for by a combination of relevance-theoretic pragmatics (Sperber and Wilson 1995, Carston 2002), certain motivated assumptions about visual information processing, and the grammar formalism of Dynamic Syntax (Kempson et al 2001; Cann et al 2005). DS models pronouns as encoding procedures that introduce a variable-like entity (e.g. a metavariable), which needs to be replaced by a semantic value (of the appropriate type), allowing for descriptive constituents, which emerge as a result of relevance-driven processes of identification and inference, to provide the pronoun with the relevant descriptive interpretation. Alternatively, the pronoun can be replaced by a singular value that communicates a descriptive proposition as an implicature. The context and the pronominal form used determine which of these approaches is the best suited.

Contents

List of Figures

Chapter 1

Identification-based descriptive uses of pronouns: the data

1.1 Introduction

This thesis deals with a phenomenon described in the literature as 'deferred ostention' (Quine, 1968, 194), 'deferred uses', 'deferrals' (Nunberg, 1993), or simply 'descriptive indexicals' (Nunberg, 2004a,b). It can be illustrated by the examples below (a use of a sentence in a particular occasion is stated in (a), whose intuitive truth-conditions are in (b); in this thesis, italics mark emphasis or the interpretation of linguistic stimuli, small caps represent concepts, quotes refer to the linguistic form (spoken or written) of a word or sentence use, and the asterisk for ungrammaticality/infelicity):

(1)   a.   *Someone pointing at a huge footprint in the sand says:* 'He must be a giant!' (Schiffer, 1981, 49)
       b.   *The person whose foot made the print must be a giant*.

According to this example, a natural language expression, in this case, a pronoun, seems to depend on some aspect of the immediate context, in this case, a footprint, in order to convey truth-conditions that are *descriptive* or *general*, as opposed to *singular*. That is, the proposition in (1)b does not say anything about a specific person, but rather about *whomever happens to have made that footprint*. I will clarify these notions shortly.

Since the two main features of cases exemplified by (1) are (i) the expression's dependence on some identifiable entity in the environment (e.g. the footprint in (1)a) and (ii) the expression's descriptive truth-conditional contribution, I shall de-

scribe the phenomenon by 'identification-based descriptive uses' of pronouns, or just 'descriptive pronouns', for short. This terminology serves a double purpose. First, it allows a neutral characterisation of the data to be reviewed below, which names such as 'descriptive *indexicals*' do not achieve. Secondly, it is a promissory note for an account that unifies this phenomenon with other types of descriptive interpretations, such as bound pronouns or cross-sentential dependencies; to be reviewed in sections below. Also, I take it that the relevant descriptive interpretations correspond to information (propositions) communicated by these pronominal uses without being committed to any particular level of content at which this takes places. Grice (1969) notoriously proposed a distinction between *what is said* by an utterance, which roughly amounts to its assertoric content, and what the speaker merely suggests or *implicates* by it. The descriptive content in question could - in principle - belong to either one of these (or other) levels of information. This distinction will be discussed in detail in chapter 2.

At this point, one may argue that such interpretations are not intrinsically tied to natural language pronouns. For example, had the speaker of (1) used 'Joe' instead of 'he' in the example above, an interpretation similar to (1)b would be communicated. I agree with the observation, but I will confine myself to a treatment of pronominal expressions only. Names and other expressions would involve complications that go beyond the aims of this thesis.

The main goal of this chapter is to introduce the interpretative properties of the pronominal uses just mentioned. In section 1.2, I review the features of the phenomenon that are more or less common ground in the literature. This constitutes the *core data* that theories of descriptive pronouns must account for. Afterwards, I will compare these data to other types of pronominal interpretation. Section 1.3 deals with deictic or indexical uses of pronouns. Section 1.4 deals with pronominal binding. Section 1.5 deals with cross-sentential anaphora or *donkey* uses of pronouns. Section 1.6 discusses generics. Section 1.7 makes some introductory relations between identification-based descriptive pronouns and metonymy. Section 1.8 discusses data which builds on the core data, namely, *the extended properties* of descriptive pronouns. Section 1.9 summarises what this chapter achieves and lays out the plan for the remaining chapters in the thesis.

Although this amounts to a long introduction to the various ways in which pronominal expressions can be used, it is far from exhaustive. For example, resumptive pronouns [1] are not mentioned. I have opted to concentrate on phenomena

---

[1]Resumptive pronouns are expressions that appear twice, often distributed over main and relative clauses, as in 'Voici l'homme que Marie lui a parlé' (French) [gloss: here is the man that Marie to

that seem to be usually confounded with identification dependent descriptive uses in order to cast light on what this type of pronominal use really is. Omissions were necessary.

1.2 The data

1.2.1 Overview

In descriptive pronominal uses, a general or descriptive (as opposed to singular) proposition is communicated via the saliency of a particular individual or object at a given occasion. I am using the words 'individual' and 'object' with enough flexibility to include inanimate objects, people, places, geographical formations, mathematical entities, and so on. In addition to (1), a variety of examples has been regarded as instances of descriptive uses in the literature. Consider:

(2)   a.   *US Supreme Court O'Connor saying:* 'We might have been liberals.' (Nunberg, 1993, 14-15)

        b.   *The US Supreme Court Justices might have been liberals.*

(3)   a.   *Bill Clinton saying:* 'The founders invested me with the sole responsibility for appointing Supreme Court Justices.' (Nunberg, 1993, 20)

        b.   *The founders invested the US President with the sole responsibility for appointing Supreme Court Justices.*

(4)   a.   *A professor pointing to a PhD thesis in his shelf and saying:* 'AHRC gave a post-doc grant to her.'

        b.   *AHRC gave a post doc grant to the author of that thesis.*

(5)   a.   *John, expecting a call from his mother, answers a phone call from his colleague and says:* 'Oh, I thought you were my mother.' (Nunberg, 1990)

        b.   *John thought that the person calling was his mother.*

(6)   a.   *Bill, to a friend at an art gallery on a Sunday:* 'John dined yesterday... with her [*pointing to a painting*].'

        b.   *John had dinner with the artist who painted that picture on the 18th of June 2011.*

---

him has talked] (Haegeman, 1994, 409), where 'que' and 'lui' refer to the same individual.

The phenomenon can be illustrated by use of the pronoun in (a), which communicates the information in (b). The pronouns and their respective interpretations can be summarised as the following pairs: 'he' - *the person whose foot made the print* in (1), 'we' - *the American Supreme Court Justices* in (2), 'me' - *the American President* in (3), 'her' - *the author of that thesis* in (4), 'you' - *the person calling* in (5), and 'her' as 'the artist who painted that picture' in (6). In some cases, such as (2) and (3) a referential interpretation about Clinton and O'Connor might also be available. This, however, does not exclude the descriptive feel of the utterances (more on that in chapter 2).

These examples suggest that descriptive interpretations are available regardless of the expression's number (singular or plural), case marking (nominative or accusative or dative, data about other cases is needed), person (first, second or third), and gender (masculine vs. feminine). Issues may arise for the genderless 'it' because a clash between the features of the word and the most salient object in the context may trigger descriptive interpretations (see Nunberg 1993). The examples also show that the interpretations stated in (b) are not blocked by the type of argument taken by the verb or its adjuncts. For instance, in (1) and (2) the pronoun is the external argument of the verb 'to be' plus its modal auxiliary, whereas in (3) the pronoun is the internal argument of 'invest', in (4) 'her' is the goal argument of the ditransitive 'give', and in (6) 'her' occurs in an adjunct to the verb 'to dine'.

Moreover, descriptive interpretations seem to be available throughout a variety of speech acts and syntactic environments. Consider the following utterances.

(7)   *A politician at a press conference, worried about a certain controversial news team, utters to one of his secretaries*:

   a.   Who is he [*pointing at a copy of The Financial Times*]?
   b.   Make sure he [*pointing at a copy of The Financial Times*] shuts up!

Above, the use of 'he' contributes *the editor/journalist of the FT* to the communicated content, regardless of the speech act (assertion, question or order) used. Now, consider different ways in which this dialogue can continue, below.

(8)   *The secretary (Ann) protests to the politician's request by uttering*:

   a.   It is them [*pointing to the copy of the FT*] who make your policies known to the public!
   b.   Your policies have been made known to the public by them [*pointing to the copy of the FT*]!
   c.   They [*pointing to the copy of the FT*] seem to make your policies

known to the public.

Above, the use of 'them' conveys the interpretation *the FT people*, which is preserved under whatever mechanisms govern the interpretations of: (i) clefts, as in (8)a (note the exhaustivity effect: no one, but the FT people, makes the politician's policies known to the public, see Ward 2008), (ii) passives, as in (8)b, and constructions where the subject of the embedded clause has been 'raised' to the subject position of the main clause (raising constructions), as in (8)c.

It seems that the descriptive interpretations of pronouns are not constrained to a particular speech act or syntactic environment. This claim will be revisited when we consider the extended properties of such uses.

At this point, a few remarks on the *descriptive* status of the communicated propositions must be made. Imagine a context where the salient entity in it, say, the *Thriller* album, is related to another specific entity, say, Michael Jackson. In this scenario, illustrated below, the speaker may communicate a singular (as opposed to descriptive) proposition. Consider:

(9)   a.   *Pointing to a Michael Jackson album at a music shop:* He surely died young.
       b.   *Michael Jackson died young*.

If (9)a communicates the singular proposition in (9)b and the case above is structurally identical to the ones previously reviewed, then communicating descriptive information is not an essential property of the phenomenon. The approximation between (9), on the one hand, and the other interpretations reviewed so far, misses an important point: (9) can be reduced to a simple case of pronominal deixis (to be discussed in section 1.3). If we assume that pronouns are place holders for certain representations and that pointing to *Thriller* is an efficient way of making Michael Jackson the salient value of the expression, then there is no difference between pointing to the album or the individual himself. In other words, (9) is more similar to a situation where someone points to a picture of Michael Jackson intending to refer to him. In this case, we resist saying that there is any form of indirect mediation between individuals (i.e. the copy of *Thriller* and Michael Jackson himself). Thus, the fact that I confine my interest to cases that receive descriptive interpretations is partially justified on methodological grounds. If the interpretations were not descriptive, the phenomenon illustrated above would be an instance of indexicality. But this discussion raises an important question. If communicating descriptive propositions is one of the hallmarks of the pronominal uses discussed in this thesis,

how should this notion be understood?

Answers partially depend on one's positions regarding the nature of information that is expressed linguistically. Putting rather simplistically, there are two main takes on this issue. On the one hand, some claim that every piece of communicated information must result from the audience's attempt to retrieve what the speaker intended by her utterance. This is the *Intentionalist* position (Grice, 1957; Sperber and Wilson, 1995; Carston, 2002). On the other hand, some claim that much of what is communicated results from a similar process, but some information is determined solely in virtue of the words used and *facts* about the context (Montague 1970a; Kaplan 1977 are good examples, but many others fall here). This is the *Non-* (or semi-) *Intentionalist* position (these positions will be better discussed in chapter 2).

In the intentionalist picture, the fact that a given piece of linguistically expressed information is descriptive depends on (rationally constrained) intentions the speaker wishes to convey and the audience's ability to rationally reconstruct the intended information. If the speaker has a particular person in mind (mutually known by the audience) when using the pronoun, then the hearer faces the task to retrieve a proposition about that particular individual. If the speaker did not have a particular individual in mind, then the communicated information is about whomever fits some descriptive content (see the discussion on attributive descriptions in Donnellan 1966). In some cases, the audience may know of a particular someone who does fit the intended description, but this may fall outside the intentions of the speaker and might not be easily classifiable as *communicated*: the audience matches the descriptive content with its unique satisfier at its own expense. For example, in cases where the audience - but not the speaker - knows who made the footprint in example (1).

In the non-intentionalist picture, things are more complicated. In the literature, it has been argued that there are facts that determine whether a content expressed by an utterance is singular or not. If one or more of these facts do not hold, the expressed content is descriptive. Following Neale (1990), three types of facts - (E)pistemic, (M)etaphysical, and (L)inguistic - have been assumed, although not uncontroversially, as *definiens* of singular propositions (the definitions below are based in Galery 2008, 160).

(E) One who entertains the singular proposition must know which object is referred to by the constituent expression *e* of the utterance used to convey it (see Russell, 1910).

(M) The truth-conditional contribution of the constituent expression *e* is exhausted by the object it denotes (i.e. the existence and individuation of a singular proposition depends on the existence and individuation of the object it is about) (see Russell, 1904).

(L) If the utterance describes situations different from the actual, the constituent expression *e* still denotes the same individual in such situations (see Kripke, 1972).

For the moment, I take it that these definitions are neutral with regards to whether *e* is an expression type or token. Thus, there are many sources for the descriptive status of the relevant pronominal interpretations, according to the *non-Intentionalist* picture. Some of the previous examples, such as, (1), (4), and (6) can be described as resulting from violations of (E): the audience simply does not know who made the footprint, wrote the thesis, or painted the picture, respectively. However, many have proposed that (E) should *not* be regarded as a constraint on the expression of singular propositions (Kripke 1972; for a recent example, see Borg 2004). Still, other examples, such as (2) and (3) seem to violate (L) as opposed to (E). For example, 'we' in (2) is interpreted as *the American Supreme Court Justices*, which selects different individuals in different non-actual situations. Thus, on the assumption that our semantic/pragmatic theories must explain what is *intuitively* communicated both intentionalists and non intentionalists agree that some of the pronominal uses reviewed above are indeed descriptive (although for different reasons or via different processes). Let us now systematise the basic properties of the examples examined so far.

1.2.2 Core properties

If the paraphrases stated in (b) adequately capture linguistic intuitions concerning what is communicated from (1) to (8), it seems that descriptive uses of pronouns display the following property:

Descriptiveness: The information conveyed by the relevant pronominal uses is descriptive or general.

For the moment, I will also remain neutral on the kind of determiner that constitutes the descriptive truth-conditions. So far I have been using the definite 'the' in expressing the relevant truth-conditions in (b), but that corresponds only to an

approximation to the communicated content. Now, the descriptive uses in question seem to display other interpretative properties. Consider an example based on Quine (1968):

(10)   a.   *Someone pointing at a car whose windshield is full of tickets:* 'He is going to be sorry.'

        b.   *The owner of the car is going to be sorry.*

Let us now ask the following question. Under what circumstances is the content in (b) unavailable to the audience? Two situations spring to mind. First, consider a situation in which the audience is unable to single out the salient object as belonging to any particular kind. We could imagine the hearer seeing some amorphous blur in his visual field. Under these circumstances, what would (a) convey? Maybe the hearer would assume that the blur corresponds to a person, given the lexical material of 'he', but, in this case, the interpretation in (b) does not seem to be easily accessible. That is, descriptive pronouns seem to require the contextually salient individual to be classified by the audience as belonging to a certain *kind*. In (10) above, the relevant kind is *car* or *vehicle*. Without its identification, the interpretation in (b) would not be retrieved.

This intuition concerning the availability of descriptive interpretations can help to determine another necessary requirement that an account of descriptive uses must meet. More specifically, it motivates:

Identification Dependency: Descriptive uses of pronouns depend on the identification of a particular entity as belonging to a certain kind.

The property above illuminates the *identification-based* nature of descriptive pronouns. Now, the notion of *identification* has raised a great amount of controversy in the philosophical literature. At this point, I assume, following Strawson (1959) and Evans (1982, ch. 5), that an agent can employ various cognitive mechanisms, sensory modalities, memory, lexical information and full blown communicative practices (testimony), to identify an individual across a variety of situations[2]. Still, this thesis focuses on cases exploiting visual information. Empirically grounded notions of visual individuation and identification will be proposed in chapter 3.

Note also that, under a certain understanding, deictic uses of pronouns can also

---

[2]I personally think, however, that this unification requires empirical support and does not sustain merely on conceptual grounds.

be described as identification dependent uses. For example, when Bill points to Susan and says 'She is brilliant', the proposition expressed depends on the identification of a particular person. Whether this is done by context and word alone, on the one hand, or by the audience, on the other, is another matter. Despite such similarities, there are reasons to claim that the forms of identification present in descriptive and deictic cases are slightly different from one another. Kahneman et al. (1992) consider the example below:

(11)   *Someone seeing an object moving very fast in the sky utters*: 'It's a bird... It's....a plane.... Ohh no...It's superman' (Kahneman et al., 1992).

In this case, 'it' seems to refer to the same thing regardless of the kind (*bird, aeroplane, superhero*) which the speaker classifies the object as belonging to. That is, in some cases of deixis, a bare individual seems to be all that is needed for the expression of the proposition (this notion will be developed in chapter 3).

It may be also noted that the entities mentioned so far have clear boundaries. But are descriptive interpretations only available when prototypical instances of solid, bound material objects are demonstrated? The answer seems to be negative. Consider the following example:

(12)   a.   *A student in a demonstration pointing at a cloud of tear gas thrown by the police warns two citizens who walk into the violent demonstration:* 'They are bringing reinforcements'.
         b.   *The policemen are bringing reinforcements.*

In this scenario, a huge amount of tear gas has been used, so that both students and the surprised citizens were engulfed in it. There is no single *entity* that is singled out in the context, just some amorphous stuff. This situation, however, poses no problem for the definition of *Identification Dependency*. The audience seems to identify something in the environment - some stuff - as an instance of the kind *gas*. This classification makes some knowledge accessible, namely that the police uses tear gas to suppress violent demonstrators, which then paves the way to the right descriptive interpretation of the pronoun. That being said, it seems that most examples of identification-based descriptive uses of pronouns seem to depend on the saliency of prototypical material objects.

Now, let us take a look at a second reason for the unavailability of descriptive interpretations, namely, the audience's lack of knowledge associated with the concept used in the identification of the salient individual. Imagine a scenario where

19

the audience does classify the salient object in (10) as an instance of the kind *car*, but lacks the knowledge that cars have owners, who are typically responsible for paying fines (providing reasons for one's being sorry), as in a case where the hearer does not have a proper sense of ownership, nor a good idea about the financial burden that comes with illegal parking. That is, the hearer is able to discriminate cars in the environment, but lacks that right kind of information the speaker seems to hint at. In this case, the proposition that *the owner of the car is going to be sorry* is not so easily accessible. This scenario shows that classifying an individual under a concept serves as a gateway to information (world knowledge) that is relevant for establishing the the right descriptive interpretation. If the concept is not associated with the right body of information, the hearer fails to interpret the pronoun correctly. These observations motivate the following interpretative property:

Connection: The mental representations used to identify or classify the salient object in the context must make information available so it can provide the right descriptive content for the pronominal use.

Connection seems to be a property relating the other two, more basic, interpretative properties. That is, an adequate account of descriptive uses must explain the relation between *Descriptiveness* and *Identification Dependency*, namely, how the concepts used to represent objects as falling under a certain kind make the descriptive interpretation available in a principled way. These three features of descriptive pronominal interpretations constitute their *core interpretative properties*.

Now that we have covered a variety of examples of descriptive cases and highlighted their key features, I will compare the examples stated here with other pronominal uses.

1.3   Deictic or indexical uses

Similar to identification-dependent uses, deictic or indexical uses of pronouns (the names will be used interchangeably) depend on certain features of the occasion of speaking (or writing). For example, the truth-conditional contribution of words like 'I' depends on who utters it and similar rules can be written for 'you', 'she', 'now', 'tomorrow', 'today', and analogous expressions. Given the properties discussed in the previous chapter, indexical uses seem to share with descriptive pronouns an interpretative property similar to *Identification Dependency*. That is, the understanding of the utterance below depends on the identification of the referent of the

pronoun.

(13)   *Thiago Galery introducing himself to an undergrad student:* 'I'm your backup tutor'.

The understanding of the utterance above attributes to a certain identified person, who happens to be speaking, the property of *being a backup tutor*. However, in the definition of *Identification Dependency*, the salient individual must be classified as belonging to a certain kind. As mentioned in the previous section, one could imagine a scenario where someone sees an object moving very fast in the sky and utters 'It's a bird... It's....a plane.... Ohh no...It's Superman', where 'it' refers to the same entity throughout, even though the concepts used to classify the object change. Examples like this motivate the idea that at least some cases of deictic pronominal interpretation do *not* require the notion of *identifying as C*, where *C* is a place holder for some concept, used in the definition of *Identification Dependency*. Or else, these cases do involve such a notion, but the concept that performs the necessary identification is highly general, like MATERIAL OBJECT or PERSON. I will refrain from commenting on these issues until the notions of individuation and identification are properly discussed in chapter 3. For now, let us assume that the first interpretative property of indexical uses of pronouns is a more basic, possibly disjunctive, form of identification dependency, namely:

Identification Dependency (indexical): Deictic or Indexical uses of pronouns depend on the identification of an object as a bare particular or as an entity that belongs to a very general kind (e.g. MATERIAL OBJECT or PERSON).

There are two ways to understand what the necessary *identification* in the definition above does. On the one hand, one could claim that the words themselves (not the hearer) identify the relevant object in the context, as in the non-intentionalist picture mentioned in the previous section. Consequently, the audience does not need to know *who* the pronoun picks out in order for information to be expressed (Montague 1970a; Kaplan 1977, that is, condition (E) on singular thinking, in the previous section, does not govern the level of information determined by indexical uses). On the other hand, one could say that the notion of identification is communicator-based, that is, in order for information to be expressed, the hearer must know *who* the pronoun picks out (indexically induced thoughts would obey condition (E) in the previous section). Besides its intuitive appeal, this position might be motivated by certain differences between utterances that are true in the same circumstances.

As Frege (1967, 1948, 1956) observes, the thought constituents determined by indexicals display a peculiar form of *cognitive significance*. For example, Anna, an undergrad student who is addressed by my utterance of 'I'm your backup tutor' (i.e.(13)) may learn something new, whereas an utterance of 'The UCL pragmatics backup tutor is the UCL pragmatics backup tutor' does not inform anyone of anything, even though this utterance and (13) have the same truth-conditions (in the actual world / current time, say).

There are many proposals in the literature in support of each of the two ways of understanding the mechanisms responsible for identification (on the words as individuators side, see Kripke (1972); Kaplan (1977); Salmon (1986), on the communicator or Fregean *senses* as individuators side, see Perry (1977); Evans (1982); Wettstein (1986), among many others). I will not review this controversial topic here, although the discussion in chapters 2 and 3 argues for the latter position. On more neutral grounds, one could argue that whatever the mode of identification in indexical uses is, it must be *non-descriptive*. This can be attested by the following contrast, pointed out by Kripke (1972).

(14)   a.   *A person pointing at Pelé, the great Brazilian football player in the 50s and 60s:* He could have been Argentinian.

          b.   *Someone utters out of the blue:* The best football player in the 50s and 60s could have been Argentinian.

As Kripke (1972) observed, (14)a asserts that the history of the world could have been such that Pelé - the best football player in the 50s and 60s - was born somewhere in Argentina. *But in addition to this reading*, (14)b can convey that the history of the world could have been such that an Argentinian, different from Pelé, is the best football player in the 50s and 60s. In order to capture the presence of this ambiguity[3] in (14)b and the lack of it in (14)a, a distinction between expressions: referring or singular vs. descriptive or general has been proposed. An argument for this distinction can be schematically represented as follows (see Kripke 1972, 48-49, 71-77):

(15)   *Modal argument*

          a.   Premise 1: If an expression e differs from an expression $e*$ in terms of modal behaviour, they belong to different classes of expression (e.g.

---

[3]Note that this ambiguity can also be observed if we replace the modal operator by other scope taking elements, such as negation. However, the notion of rigidity does not coincide with that of scope, because scope is understood as relative to other scope taking elements in the sentence and rigidity is not a relative notion.

referring vs. descriptive) (assumption) .

    b.   Premise 2: Pronouns and definite descriptions differ in terms of modal behaviour (motivated by the examples in (14)a and (14)b).

    c.   Conclusion: Pronouns and definite descriptions belong to different classes (types) of expression.

In the first premise, two expressions e and $e*$ differ in modal behaviour if and only if one of them designates the same individual in every possible world (i.e. ways the world could have been) and the other does not. Kripke (1972) calls the former kind of expression a *rigid designator*. His point is the following. The content expressed by an utterance corresponds to the utterance's *possible world truth-conditions*. When we embed pronouns and definite descriptions under modal operators, the resulting utterances express different modal behaviour: the first arguably select the same individual across worlds, while the later selects different individuals according to different worlds. Therefore, these expressions must be distinguished[4]. The following property, which Kripke took to hold of lexical types, seems to be motivated.

Rigidity: Deictic pronouns refer to the same entity in every possible world/situation.

    Finally, let us consider issues of existential order. Suppose I hallucinate a person in front of me and say 'He has a funny hat', whilst pointing to the illusory man. Have I conveyed information or have I conveyed nothing at all? To many, empty cases of deixis carry no information and thus the following property of deictic interpretations has been proposed:

---

[4]Kaplan (1977) laid out a machinery that is able to capture the contributions of indexicals that roughly correspond to the intuitions invoked by Kripke's argument. The gist of the proposal follows a strategy originally pursed by Kamp (1971), namely, *double-indexing*. It involves the assumption that a given parameter (possible world) figures twice in the formal description of the utterance content (sentences relativised to contexts). Kaplan distinguishes two relevant aspects of semantic evaluation: the context of utterance, modelled as a set of parameters which includes the actual world, and the circumstance of evaluation, modelled as a set of parameters which includes a possible world slot (not necessarily restricted to the actual). The linguistic meaning of indexicals (characters) is then modelled as functions from contexts to contents, the latter being functions from circumstances of evaluation to truth-values. Since the context only includes the actual world, pronouns only refer to entities (who are speaking, being addressed, or demonstrated) in the actual world. This content then is evaluated for truth or falsity against a circumstance of evaluation. This guarantees that pronouns refer to the same thing in different worlds, i.e. they are rigid designators. Descriptions on the other hand are not restricted to select their referents from the context set and can pick out different entities according to parameters in the circumstances of evaluation.

Non-emptiness: The content of utterances containing indexicals depends on the existence and individuation of the objects they refer to.

The definition above roughly corresponds to the constraint (M) on singular thinking stated in the previous section. *Dependence on existence* means that if the referent of an indexical does not exist, the utterance where it occurs does not have any content. This rules out the empty case just considered. Now, *dependence on individuation* means that the content of utterance containing an indexical is individuated according to the individuation of its referent, and here we could either say that words themselves are individuators or that communicators/agents/senses are. It is worth noting that a circularity emerges at this point. Non-emptiness is supposed to be a metaphysical aspect of indexical content, but it draws upon an epistemic notion, namely, *individuation* (i.e. objects are individuated by cognitive agents), in its definition. Rather than solving such circularity, I will assume a meta-property, namely, *Content singularity*: which states that utterances containing indexicals express singular (i.e. not descriptive) content about a contextually salient entity (i.e. the sum of *Identification Dependency* (indexical), *Rigidity* and *Non-emptiness*, whatever the best way for these to be understood is).

On the basis of this property, one can conclude that indexical uses of pronouns differ from descriptive uses. One way to see this is the following: if the referents of deixis do not exist or cannot be individuated, it could be argued that nothing is expressed. In this picture, if both speaker and hearer hallucinate the referents of deixis, no information is expressed, because *Content Singularity* is violated. Now, like pronominal deixis, identification-based descriptive interpretations require an individual to be individuated (under a kind), but they do not communicate content about any specific individual. Thus, if speaker and hearer hallucinate an individual, which is classified under the same concept in the context, then it could be argued that some form of descriptive content is communicated (i.e. descriptive cases are not constrained by *Content Singularity*[5]). In conclusion, deixis and identification-based descriptive uses seem to be constrained by different contextual requirements [6].

---

[5]In other words, like violations of epistemic constraints on singular content, the violation of metaphysical constraints may characterise the content imparted by the utterance as *descriptive*.

[6]As will emerge in chapter 2, I assume an intentionalist picture of reference. Thus, I would be prepared to bite the bullet and say that in cases where speaker and hearer jointly hallucinate an object or experience virtual reality, say, singular or descriptive content can be communicated (as long as the hearer figures what is intended by the speaker).

This asymmetry poses a challenge for accounts of pronominal interpretations. How can pronouns express descriptive readings in certain contexts but singular ones in others? From a meta-theoretical point of view, a unified theory is preferable. Now, if we take a look at Kripke's argument above, its conclusion states that pronouns *must* be different from descriptions because of their modal behaviour. Important consequences follow from this. If indexical uses convey singular interpretations in virtue of the linguistic rules they encode, the descriptive interpretations reviewed in the previous section (and others to come) cannot be explained solely in virtue of their linguistic meaning.

Two ways of accommodating this consequence are worth mentioning. First, one could say that the pronouns are ambiguous between lexical types that govern singular and descriptive readings. The desired explanatory unification would be lost. Second, one could claim that the indexical uses are governed by the linguistic meaning of pronouns, while identification-based descriptive pronouns would be accounted for by rules governing cooperative communicative behaviour: the level of Gricean conversational implicatures (see Stokke 2008 for a recent defence). I reserve a full assessment of this possibility for chapter 2, but for now I would like to say that there are other descriptive uses of pronouns that seem to be governed by linguistic rules, such as bound uses and cross-sentential dependencies. On the assumption that such cases exist, why cannot such rules be extended to cover identification-based descriptive uses? This strategy would be supported by the fact that the descriptive interpretations reviewed in the previous section constitute evidence against Kripke's claim that indexical pronouns behave like rigid designators in virtue of their linguistic type (see Nunberg 1993 for the expression of such view). In other words, the modal argument may be invalid, because premises 1 and 2 in (15) are not true. Although descriptive uses share the identification-based aspect of interpretation with indexical uses, the two seem to be quite different phenomena. Let us now turn to the bound uses of pronouns.

1.4   Pronominal binding

In many cases, pronouns display a form of dependency that is different from that of deictic interpretations. Instead of picking out a salient individual in the context, the interpretation of the pronoun depends on another expression in the neighbouring linguistic environment. Consider:

(16)   John thinks he is smart.

a. John thinks John is smart.

  b. John thinks Bill is smart.

The utterance in (16) has two possible interpretations: one where the pronoun refers to a salient individual in the context, (16)b, another where the pronoun is *bound* by an antecedent expression, such as 'John' in (16)a. The bound reading is usually represented by sameness of index (e.g. John$_1$ thinks he$_1$ is smart) and the non-dependent reading by different indices. Dependency here indicates an asymmetric relationship: the bound pronoun depends on an antecedent (the binder) in order to establish its interpretation, but not vice versa. Natural numbers usually indicate this relationship, but letters or any other symbol to achieve this goal could be used[7]. Here I take these symbols only to indicate the *prima facie* intuitive relation between expressions.

For my purposes, it is interesting to see whether bound pronouns convey descriptive or singular interpretations, hence siding with either the descriptive or indexical interpretations reviewed previously. The example above, however, cannot illuminate this issue, for it does not rule out the following possibility. The proper name 'John' is used to refer to a particular man in the context and then the pronoun is just used to refer to the same entity. In short, there is nothing that blocks (16)a, which we would describe as an instance of binding, from being assimilated to cases of indexicality (as the reading in (16)b may be). In what follows, I will look at environments where the claim that bound pronouns result from coincidental reference with their antecedents seems to give the wrong predictions, namely, dependencies in complex antecedents and dependencies in ellipsis sites. Both raise some representational properties of binding that underlie the complementary distribution of pronouns and reflexives; to be examined immediately afterwards.

The first environment in which the *descriptive* nature of bound pronouns can be attested emerges when we consider utterances with complex antecedents (as pointed out by Geach 1962). Consider the examples below:

(17) Binding by quantificational antecedents.

  a. Every lawyer believes he is smart.

  b. No mother thinks she is wrong.

---

[7]Note that in the Chomskyan tradition, (Chomsky, 1981; Heim and Kratzer, 1998; Büring, 2005; Elbourne, 2005) the use of numerals as indices distinct from variables paves the way for a distinction between syntactic and semantic binding, *modulo* a specific understanding of these terms.

The examples above have a reading where the pronoun does not depend on any antecedent ('he' refers to John and 'she' to Mary, say) as well as the bound reading. The former will not be considered. Now, if pronominal binding reduced to referring to whatever the antecedent refers to (see Evans, 1977), the cases above would be difficult to explain. In the case of (17)a, the best candidate for this alleged reference relation would be all the individuals belonging to the set of lawyers (in the model). However, this misses an important point. (17)a does not express that all lawyers think of themselves (as a collection) to be smart. It rather expresses that for each lawyer, that lawyer thinks he is smart (and may be sceptical about the smartness of other lawyers). Things are even more complicated for (17)b. The antecedent 'No mother' does not refer to anything, so how can the pronoun be co-referential to it [8]? Still, what is expressed by this second utterance is straightforward: there is no individual such that this individual is a mother and thinks herself to be wrong.

These examples highlight an important interpretative property of bound uses; one which is one of the hallmarks of *Descriptiveness*, namely:

Co-variation: the interpretation of the pronoun co-varies with another expression in discourse.

If bound pronouns co-vary with other expressions in the utterance, no proposition about a specific entity (i.e. a singular proposition) is expressed. In short, the presence of *Co-variation* entails *Descriptiveness*.

In addition to environments containing complex quantificational antecedents, the idea that binding reduces to coincidental co-reference can be rejected, when we consider evidence from VP ellipsis (Δ indicates unpronounced but interpreted material, * ungrammatical/infelicitous utterances, and strikethrough reconstructed material)[9]. Consider:

(18) John saw his sister and Bill did Δ, too.

    a. John$_1$ saw his$_{1,i.e.John's}$ sister and Bill$_3$ did ~~see his$_{1,i.e.John's}$ sister~~, too. (strict)

    b. John$_1$ saw his$_{1,i.e.John's}$ sister and Bill$_3$ did ~~see his$_{3,i.e.Bill's}$ sister~~, too. (sloppy)

    c. John$_1$ saw his$_{4,i.e.Andrew's}$ sister and Bill$_3$ did ~~see his$_{4,i.e.Andrew's}$ sister~~, too. (identical deictic)

---

[8]Evans (1977) replies by assuming that the pronoun refers to the empty set.
[9]Focus can also be used to illustrate the same point, see Büring (2005).

d. *John$_1$ saw his$_{4, i.e. Andrew's}$ sister and Bill$_3$ did ~~see his$_{5, i.e. Marc's}$ sister~~, too. (different deictic)

According to the view that pronominal binding amounts to cases where the pronoun refers to whatever the antecedents refer to, sameness of indices does not entail any structural constraint: it just indicates that two expressions co-refer in virtue of the context, say. This assumes that linguistic information is completely neutral with regards to the pronoun's referential possibilities: the pronouns would induce a variable like entity that could refer to anything in principle.

On the assumption that ellipsis is an operation where some linguistic material is unpronounced but still interpreted, it is natural to take the reconstruction of the elided fragment as having the same interpretative possibilities as the material it depends on (against this claim, Fiengo and May 1994 argue for the idea that ellipsis targets co-arguments). Now, if the copied structures were free to refer to anything, the interpretation in (18)d would be possible. However, an utterance of (18) cannot give rise to it. Thus, many have assumed that binding cannot merely be a by product of 'accidental' co-reference. Rather, binding seems to exploit some structural configuration between antecedent and pronoun which VP ellipsis is sensitive to. On the basis of this claim, one could ask: under which conditions can a bound pronoun co-vary with its antecedent? Consider the following binding possibilities:

(19) a. John$_2$ loves himself$_2$.
      b. *John$_2$ loves him$_2$.
      c. *John$_2$ thinks Peter$_3$ despises himself$_2$.
      d. John$_2$ thinks Peter$_3$ despises him$_2$.

The pattern above is usually referred to as 'the complementary distribution' of pronouns and reflexives. It indicates that some antecedents are accessible for binding purposes depending on the type of pronominal expression (pronoun or reflexive pronoun). This motivates the following interpretative property:

Accessibility (binding): Certain expressions/representations can provide antecedents for (bound) pronouns, while others cannot.

As for the binding possibilities above, it seems that whether an expression is *accessible* as a binder of a pronominal depends on the position it occupies in some form of structure. In other words, the complementary distribution of pronouns and reflexives is one of the early evidences for the commonly held assumption that lin-

guistic representations are structured. Tree-like representations are one way to capture such structure, allowing the definition of constraints on the binding between pronouns and their antecedents, as illustrated below (I am using a toy phrase structure loosely based on an old labelling system of transformational grammar, where sentences, verb phrases and noun phrases are represented by their corresponding letters. For the sake of simplicity, many important nodes, i.e. tense, complementiser, have been omitted):

(20)   Tree structure I

```
                    S
              ┌─────┴─────┐
          NP[John₁]      VP
                     ┌────┴────┐
                 V[loves]  NP [himself₁/*him₁]
```

(21)   Tree structure II

```
                    S
              ┌─────┴─────┐
          NP[John₁]      VP
                     ┌────┴────┐
                 V[thinks]     S
                          ┌────┴────┐
                      NP[Peter₃]    VP
                                ┌────┴────┐
                          V[despises]  NP[himself₃/*₁/him*₃/₁]
```
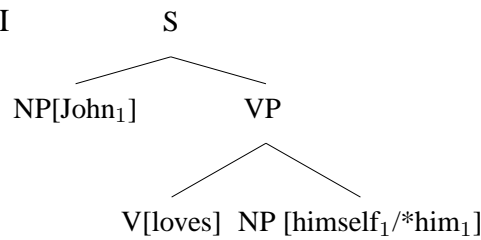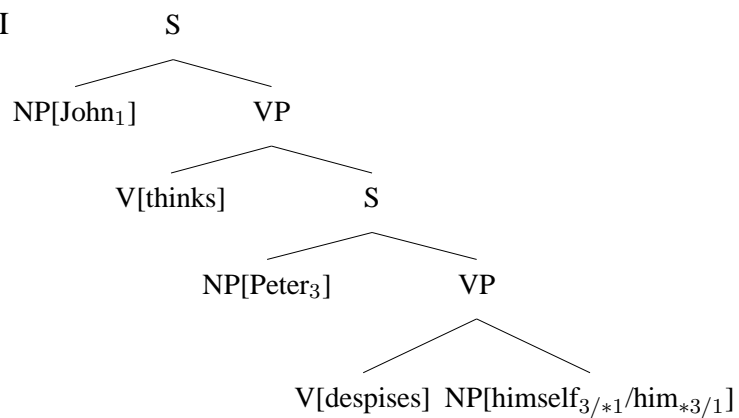
The point is quite simple. Based on these structures, it is quite easy to see that, in English, reflexive ('self') pronouns require a closer antecedent, whereas the accessible referents of non reflexive pronouns cannot be that close.

In the tradition of generative grammar (Reinhart, 1976; Chomsky, 1981), the distribution between pronouns and reflexives motivates the notion of *c-command*, which was then used as one of the ingredients for capturing a plethora of phenomena related to linguistic structure. Loosely following Reinhart (1976), a tree node **A** *c-commands* a tree node **B** if and only if: (i) **A** is a different tree node from **B**, (ii) **A** does not dominate **B** and **B** does not dominate **A**, and (iii) the first branching node dominating **A** also dominates **B**. *Domination* occurs under node motherhood: if tree node **A** has **B** and **C** as daughters, **A** dominates **B** and **C**. It is also transitive: if **B** has **E** as daughter and **A** dominates **B**, then **A** dominates **E**.

In our tree structures above, all potential antecedents - the NPs 'John' and 'Peter' c-command the pronouns. We go up one node from the NP and down in the other direction in the tree as much as we like and we find the node that could be occupied by 'him' or 'himself'. Thus, although the notion of c-command captures some hierarchical relation between antecedent and pronoun (remember, binding is asymmetric) it is not *sufficient* for explaining the complementary distribution alluded to. In the generative tradition, this theoretical demand was met by principles that invoked the notion of c-command in their formulation. Following Büring (2005, ch. 1), the general format of such principles is:

General Form of Binding Theory: An **expression class** must (or must not) be bound by a **c-commanding** or **expression class** within its **domain**.

Expression class includes pronouns (e.g. 'him'), reflexives ('himself'), and R-expressions (i.e. full determiner phrases, such as 'some man', 'that girl' or ,'John'). **Domain** would be a part of the whole structure which binding principles seem to be sensitive to: it aims to capture a notion of *locality* with regards to possible antecedents. In tree structure II, our toy tree labelling suggests that *S* (sentential clause) might be a binding domain, but a more minimal constituent containing tense information may be better suited. This allows the definitions of principles such as:

(22)   Binding Principles (Chomsky, 1981)

   (A) An anaphor (e.g. 'himself') must be bound in its binding domain.
   (B) A pronoun (e.g. 'he') must not be bound in its binding domain.
   (C) An R-expression (e.g 'John') must not be bound.

Rather simplistically, the binding principles correspond to the generative grammarian's way of capturing the *Accessibility* property of bound pronouns[10] . It is by no means the *only* way (see Pollard and Sag 1992 and also Kempson et al. 2001; Cann et al. 2005 for alternatives).

The take home lesson that I would like to draw attention to is this. The complementary distribution of pronouns and reflexives, which motivated *Accessibility*, necessitates some form of structured representation that underlies not only binding

---

[10]I do not intend to debate on the most appropriate notion of **expression** or **domain** to be used in the formulation of the principles nor on the status of these principles *vis-a-vis* the grammar architecture. For example, Chomsky (1995) proposes that the principles should be explicable in terms of the lexical properties of the expressions (pronouns, reflexives, r-expressions). Some have argued that the binding patterns follow from more general *principles of economy* (Reinhart and Reuland, 1993; Fox, 2000; Levinson, 2000).

possibilities above but also the interpretation of VP ellipsis, and a variety of other phenomena. In order to capture this structure, the Chomskyan used tree structures to define the notion of *c-command* and formulated the binding principles using this notion. These offer some characterisation of the structure underlying linguistic representations[11]. Note, though, that trees themselves are not properly explained (and are rather presupposed) by these theoretical elements. A proper explanation of tree-structured representations will emerge in chapter 5. With all the properties of bound pronouns in mind, let us now compare such uses to the pronominal uses reviewed in previous sections.

From the discussion in this section, it is clear that bound uses are more similar to identification-based descriptive uses than indexical uses in virtue of *Descriptiveness*. This was supported by the fact that bound uses co-vary with other expressions in the utterance (*Co-variation*). However, this is what also makes bound uses different from descriptive uses. The latter do not seem to co-vary with anything in previous discourse, nor seem to be subject to the *Accessibility* constraint that governs bound pronouns, which exploits linguistic structure unavailable in the case of identification-based descriptive pronouns. On the other hand, bound pronouns do not seem to display *Identification Dependency* as descriptive pronouns do. Some of these claims will be revisited in the last sections of this chapter. Still, the conclusion is that all the pronominal uses reviewed so far seem to be distinct phenomena. Let us move now to the interpretative properties underlying unbound dependencies (i.e. donkey pronouns).

1.5 Donkey pronouns

Donkey pronouns constitute a linguistic phenomenon whereby a pronoun depends on an antecedent, but the nature of this relation cannot be easily explained in terms of pronominal binding (as described in the previous section). The term 'donkey pronouns', which I will use throughout and interchangeably with 'unbound/cross-sentential dependency/anaphora', is an alternative terminology and was introduced in the literature [12] due to the enormous quantity of examples involving donkeys, as illustrated below (see Geach 1962; relevant dependencies in boldface):

(23)   a.   If a farmer owns **a donkey**, he feeds **it**.

           b.   Every farmer who owns **a donkey** feeds **it**.

---

[11]For the moment, I am neutral with regards to whether structure here is syntactic or semantic in nature.

[12]Additional terminology includes 'E-type pronouns' (Evans, 1977).

c.   A farmer bought **a donkey**. He feeds **it**.

Above, the interpretation of the pronoun 'it' depends on the antecedent 'A don-key'. Intuitively, the pronoun seems to co-vary with its antecedent. For example, (23)b conveys that the donkeys which are fed covary with the farmers who own them. Thus, like pronominal binding, donkey pronouns display *Co-variation* and therefore also display *Descriptiveness*.

Despite the similarities between donkey and bound pronouns, the former can-not be so easily reduced to instances of the latter. As reviewed in the previous section, an explanation of the complementary distribution of (bound) pronouns and reflexives (i.e. *Accessibility*) depends on some structural configuration between an-tecedent and pronoun within the same *structured representation*). In the framework of generative grammar (Chomsky, 1981), accessible binders would have to occupy a node that c-commands the node where the pronoun sits. In (23)b, 'a donkey' occurs in a relative clause that modifies the subject. This means that the first branching node that dominates 'a donkey' *does not* dominate 'it': there is no c-command re-lation. The example in (23)c illustrates this more dramatically, for the dependency occurs *across* sentences. According to Chomsky (1986), grammar can be roughly described as a set of rules (comprising an I-language) induced from the stimuli pro-vided by a specific linguistic environment *modulo* a universal set of constraints (a Universal Grammar). For him, the rules of this I-language are limited to the gener-ation of well formed *strings*. This assumption [13] together with the assumption that pronominal binding involves structural configurations described by the grammar (see previous section) precludes the cross-sentential dependency in (23)c from be-ing governed by grammatical principles and must be different from that of binding *proper*.

The last point can also be made in frameworks different from generative gram-mar. For example, if we attempt to treat the formal properties of language by trans-lating a fragment of it into first order predicate logic (see Geach 1962 for an at-tempt and an appraisal of the limitations involved), the same problem would arise. The first string of (23)c would translate as $[\exists x \ \exists y \ (farmer(x) \land donkey(y) \land owns(x,y))]$ and the second as $(feeds(x,y))$. In the second formula the variables occur free and in principle could not be bound by the quantifiers. Of course, one could stipulate that quantifier scope has been extended, but this cannot be a matter

---

[13]Note that it is perfectly possible to reject a string based view of grammar even within the gen-erative tradition, see Büring (2001) for the idea of discourse trees. The grammar model eventually adopted to describe the data raised by descriptive pronouns, namely, *Dynamic Syntax* (Kempson et al. 2001; Cann et al. 2005, described in my chapter 5) departs from this string-based view.

of stipulation; we want a principled mechanism that can account for the conveyed readings (see Kamp 1981, for an explanation based on the idea that meanings are updates of representational states).

These remarks suggest that donkey pronouns do not relate to their antecedents in the same way as bound pronouns do. Three environments reinforce this claim (i.e. that donkey pronouns do not display *Accessibility* (binding)). Consider:

(24)   *Negative quantifiers*

    a.   No lawyer thinks he earns enough.

    b.   *No lawyer earns enough. He does work, though.

(25)   *Incremental and (occasionally) maximal interpretations*

    a.   A donkey thinks it is happy.

    b.   If a farmer feeds a donkey, it gets happy.

In (24), the quantifier 'no lawyer' can be an antecedent for a bound pronoun but not for a donkey pronoun. Moreover, in (25)a 'it' seems to have existential force, that is, conveys that *there is an individual that is a donkey and thinks to be happy*. The utterance could describe a situation where not all donkeys think they are happy. In (25)b, on the other hand, 'it' seems to have universal force, that is, it is interpreted as *all the donkeys a certain farmer feeds get happy*. This asymmetry was prominently noted by Evans (1977) and Cooper (1979) and it suggests two interrelated points.

The first concerns the cardinality of the interpretation induced by the dependent pronoun. In bound cases, the pronoun seems to be entirely governed by the quantificational force of the antecedent ('it' in (25)a is interpreted as some donkey via quantifier binding). In cases of cross-sentential anaphora, the pronoun, although dependent on antecedents with certain quantificational force, *may*[14] express readings about the maximum number of entities in the relevant set (i.e. maximal or universal readings: 'it' in (25)b is interpreted as all the donkeys fed by a certain farmer). This is yet another reason for rejecting the naive translation in first order predicate logic suggested a few paragraphs above. If donkey pronouns were bound by existential quantifiers, this interpretative effect could not be accounted for. The second point is related to the incremental nature of donkey pronouns (described as 'e-type' effects, after Evans 1977). While the interpretation of bound uses seem to rely solely on that of their antecedents, donkey pronouns seem to pack a lot of information from

---

[14]They must not express such readings. Consider: 'If a farmer owns a donkey and beats it, he will be prosecuted' (Breheny, 1999). Here, 'it' depends on 'a donkey', but it does not seem that all the donkeys must be beaten in order for their owner to be prosecuted.

previous discourse, that is, 'it' in (25)b is not only interpreted as *the donkey*, but actually *the donkey which is fed by a certain farmer*. Evans (1977) and Cooper (1979) account for these properties by postulating that donkey pronouns are disguised definite descriptions (i.e. *e- or d- type* theories of unbound dependencies), which under certain theories (e.g. Russell 1905) are maximal quantifiers[15]. Many have built on this proposal (e.g. Neale 1990; Heim 1990; Elbourne 2005).

Thus, it seems that donkey pronouns cannot be reduced to bound cases because they differ with respect to how their antecedents are accessible. One could go further and argue that donkey pronouns do not related to their antecedents in any formal way, but this would be a hasty conclusion. As early as Postal (1969), the following type of example was used to illustrate that the relations between donkey pronoun and antecedent are rather intricate:

(26)   a.   Every man who has a wife should bring her to the party.
        b. #?Every married man should bring her to the party (intended reading: as (a) above).

(27)   a.   Followers of McCarthy are now puzzled by his intentions (from Postal 1969, 213).
        b. #?McCarthyites are now puzzled by his intentions (intended reading: as (a) above).

Such cases motivate the following argument. If there were no formal relationship between antecedents and donkey pronouns, the utterances in (b) could be used to express what the utterances in (a) do. Since the consequent of this conditional statement is not true, there must be a formal relationship between pronoun and antecedent (known as *the formal link condition*, after Kadmon 1987). The first attempt in the literature to capture what this formal relationship is comes from Postal (1969), below:

Anaphoric Island Constraint: An anaphorically dependent pronoun must have: (i) an overt NP antecedent (ii) that is not part of a word.

---

[15]For Russell, the meaning of *The F is G* corresponds to $\exists x \ (F(x) \ \wedge \ \forall y \ (F(y) \ \rightarrow \ x = y) \ \wedge \ G(x))$. This formula contains three different assertions: (i) that there is an F (Existence), (ii) that at most one thing is F (Uniqueness), and (iii) that everything that is F is G (Maximality). Strawson (1950) followed up by arguing that some of these components (Uniqueness) should hold at the level of what the utterance presupposes rather than asserts. See Elbourne (2005) for an updated version of the Strawsonian position.

Condition (i) of the constraint is meant to rule out (26)b and condition (ii) (27)b. Note that Postal's constraint also applies to bound cases. However, bound uses additionally require a *c-commanding* antecedent, whereas donkey cases do not. This goes to show that the *Accessibility* feature seems to be *relative* to the kind of dependency. Bound pronouns require accessible antecedents that occur in certain positions within the structure of the string itself. Donkey pronouns require antecedents that are less stringently described. In conclusion, both bound and donkey uses seem to display *Accessibility*, although the ways in which antecedents are available are specific to each case. Whether Postal's anaphoric constraint figures in the right characterisation of the *Accessibility* properties of donkey pronouns will be discussed in the section on the extended properties of identification-based descriptive uses and in chapter 3.

At this point, the properties of donkey uses should be compared to the other pronominal data reviewed previously. Like bound pronouns, donkey cases are similar to identification-based descriptive cases in virtue of *Descriptiveness*. However, the similarities seem to stop here. On the one hand, identification-based descriptive uses do not depend on any antecedent in previous discourse, simply because there is not any. On the other hand, donkey cases do not seem to depend on the identification of an individual in the environment or context in a broad sense of the term (a property that indexical and identification-based descriptive uses might share). As the examples in (26) and (27) indicate, donkey pronouns might depend solely on a linguistically given antecedent. Therefore, let us assume for now that donkey pronouns and descriptive pronouns that rely on singling out a contextually salient object are different phenomena. This tentative conclusion will be revisited in the section on the extended properties of the data. We now turn to other cases that might be similar to descriptive uses, namely, generic pronouns.

1.6 Pronouns and generics

There is another phenomenon that might be similar to descriptive uses of pronouns: genericity. In some generic uses, the pronoun is interpreted as a bare plural noun (Carlson, 1977). Nunberg (1993, 12) compares some descriptive interpretations of 'we' to uses which Carlson (1977) analyses as bare plurals. For example, in the same way that in (28), below, 'we' can be interpreted as *women*, O'Connor's use of 'we' in (2), reprinted below as (29)a, can be interpreted as *American Supreme Court Justices*. That is, the same bare plural construction that seems to be available for the generic use, below, also could be used for the descriptive case. Consider:

(28)   *Generic Use*

    a.   *Said by a woman:* We contract less diseases than men (based on Carlson 1977).

    b.   *Women contract less diseases than men.*

(29)   *Descriptive uses as generics*

    a.   *US Supreme Court O'Connor saying:* We might have been liberals (Nunberg, 1993, 14-15).

    b.   *American Supreme Court Justices might have been liberals.*

The idea is this. Since pronouns can be used in generic statements conveying bare plural interpretations, as in (28), the descriptive feel of descriptive uses might result from a similar interpretative process, as suggested in (29). Although to some speakers (Robyn Carston, personal communication) the information in (29)b does not quite capture what is expressed by (29)a, we must evaluate the approximation between generics and descriptive uses from a theoretical standpoint. In order to do so, we must first understand what generics are, besides being uses that can convey bare plural readings.

Essentially, generic statements express certain generalisations. The speaker of (28) seems to express that, *in general*, women contract less diseases than men. Similar expressive power can be attained by uses of adverbs such as 'usually', 'generally', or, better yet, 'roughly'. Adding any of these to the generic statement above would not change its truth-conditions. Such generalisations, however, are not so easily captured in terms of quantification. The statement in (28) is true even though not all women satisfy the predicate: some of them might contract diseases quite easily[16]. In short, the relevant generalisations are immune to counter examples and seem to be connected to certain stereotypical properties of the relevant individuals or kinds, as, for example, women's tendency to take better care of their health, diet and well being. This is supported by the observation that genericity is connected to 'well-established' or 'essential' properties. Consider the following example (attributed to Barbara Partee by Carlson 1977):

(30)   a.   The coke bottle has a narrow neck.

    b.   ?The green bottle has a narrow neck.

---

[16]In many cases, the relevant generalisation does not seem to be true of most members of a class. For example, the generic statements 'Mosquitoes carry diseases' might be true even though only 25 % of mosquitoes carry diseases.

In (30)a the generic reading is quite salient, whereas in (30)b, it seems that the speaker is talking about a unique bottle. Thus, the first property that generic uses seem to display is:

Expressive Genericity: Generic uses express generalisations about the relevant individuals or kinds that seem to be immune to certain counter examples and are connected to stereotypical or well established properties.

The definition above suggests another property of generics. They can be used to refer to either individuals or the kind to which they pertain. The two forms of reference are described as individual *vs* stage level generics (Carlson, 1977) or I- vs D- generics (Krifka, 1987). Consider the pronominal/bare plural uses below:

(31)   *Reference to individuals or kinds*

   a.   They [*pointing to tigers*]/Tigers are striped.
   b.   They [*pointing to tigers*]/Tigers are extinct.

In (31)a, 'they' or 'tigers' refer to the individuals that belong to the kind *tiger* and predicates that they are striped. In (31)b, this form of reference is not possible: it is not the individuals that are extinct (they can only be dead), but the kind itself, *tigerness*, say, that cannot be found around the globe. Now the kind-denoting ability of generics has a lot of interesting features. To begin with, establishing the relevant kind does not depend only on the encoded information of the noun (e.g. 'tiger') or the concept that emerges in the perceptual experience (e.g. TIGER). For example, in the generic interpretation of an utterance of 'Tigers are extinct', the speaker might not convey that all tigers are extinct, but that a sub-class of tigers, say, Caspian tigers, are (these are taxonomic readings according to Krifka 1987, which has strong similarities to the process of enrichment or narrowing in the post Gricean pragmatics literature, see Sperber and Wilson 1996; Recanati 1993; Carston 1997, 2002). Moreover, the kind-denoting and individual-denoting generic statements have different entailment patterns. For example (see Carlson and Pelletier 1995), on individual denoting interpretations, if 'Caspian Tigers are on the front lawn' is true, so is 'Tigers are on the front lawn' (e.g. a set to superset). However, kind-denoting interpretations have the reverse entailment patterns, if 'Tigers are extinct' is true, so is 'Caspian Tigers are extinct' (e.g. set to subset). Thus, the following property is motivated.

Reference to Individuals or Kinds: Generic statements can refer to either individuals or kinds, each of these having different entailment patterns.

Finally, generic statements seem to be associated with information conveyed by tense. That is, the expressed generalisations seem to range over a significant number of instances. If the utterance used describes a specific event (i.e. episodic utterances), a generic statement cannot be made. Consider the contrast below (Carlson and Pelletier, 1995, 12):

(32)   *Episodic utterances*: The/An Italian/Italians/Luigi is (are) drinking wine with his dinner.

(33)   *Non-episodic utterances*: The/An Italian/Italians/Luigi drink(s) wine with his dinner.

Generic reading of the utterances in (32) are not possible, whereas generic readings of the present tense utterances in (33) are. The following property seems to be motivated.

Event Neutrality: Generic statements can only be made by utterances that do not describe specific eventualities or situations.

The two sets of examples above also illustrate that generic statements are relatively neutral with respect to the type of determiner used: plurals, definites, indefinites and proper names can convey similar readings. This could be formulated as a separate property known as *Determiner Neutrality*, but I would formulate this here for a simple reason. Given the properties of generics reviewed in this section, namely, *Expressive Genericity, Reference to Individuals or Kinds and Event Neutrality*, generic and descriptive interpretations seem to be quite different phenomena. The descriptive interpretations in section 1.2, do not seem to express generalisations that are immune to counter examples, denote kinds, nor are neutral with regards to the eventualities described by their utterances. We will now move to an assessment of the relationship between identification-based descriptive uses and metonymy.

1.7 Descriptive pronouns and metonymy: preliminary remarks

Given the indirect relationship between the individual salient in the context (e.g. the footprint) and the interpretation (e.g. *the footprint maker*), identification-based descriptive pronouns bear strong similarities to metonymy and hence we should consider whether the former is an instance of the latter. Roughly, metonymy is 'a figure [of speech] in which one word is substituted for another on the basis of some material, causal, or conceptual relation' (Brogan et al. 1993, quoted in Papafragou 1996, 169; square brackets mine). Typical substitutions include event - cause, object - possessor, artefact - creator, concrete entities - abstract properties they exemplify. As an illustration, consider Nunberg's famous metonymy example:

(34)　　*In a restaurant, the waiter warns the cook:* The ham sandwich is getting restless. (Nunberg, 1978, 186)

Here the use of 'the ham sandwich' contributes *the ham sandwich orderer* to what is communicated, where the culinary item and the person who ordered it would be connected via some form of (contiguity or causal) relation. It is important to observe that the same interpretation can be conveyed by a descriptive use of the pronoun in the right context. Consider:

(35)　　*In a restaurant, the waiter warns the cook:* He [*pointing to the ham sandwich*] is getting restless.

Arguably, the identification of the ham sandwich provides the audience with a mental representation about it (e.g. a HAM SANDWICH concept) that provides a gateway to the interpretation *the ham sandwich orderer*. Note that in (34) the same thing might occur, the only difference is that the relevant representation (e.g. the HAM SANDWICH concept) emerges via the lexical properties of the words 'ham sandwich'. Thus, identification-based descriptive pronouns could be regarded as a form of *identification-based* (as opposed to lexical) metonymy.

　　Given the possibility of non-literal interpretations, one could ask which constituents in an utterance are interpreted figuratively. As Recanati (2004, 34-36) points out, the utterance below allows for two non-literal interpretations, in (a) and (b), respectively.

(36)　　The city is asleep.

<ol type="a">
The inhabitants of the city are asleep.
The city is quiet.

Prima facie, *either* argument or predicate (verb) could be, in principle, non-literal interpretations. That is, 'the city' could be metonymically interpreted as referring to its inhabitants, as in (36)a, or 'is asleep' could be non-literally interpreted as denoting the property of being quiet, as in (36)b (to some, e.g. Robyn Carston, p.c., the latter seems to be more like a case of metaphor than metonymy). If the observation is correct, many of the descriptive uses mentioned so far could be described as behaving like proper indexicals combined to predicates that have been interpreted non-literally. This strategy is somewhat counter-intuitive for many, if not all, descriptive uses of third person singular pronouns lack an obvious literal interpretation (e.g. 'He [*pointing to a footprint in the ground*] must be a giant', where the context lacks a suitable referent). Descriptive uses of plural pronouns (e.g. Justice O'Connor saying 'We might have been liberals') are problematic, because even though a literal interpretation of the pronoun is possible (i.e. *the actual American Supreme Court Justices*), the metonymic interpretation of the predicate is unclear. For example, which figurative interpretation of 'might have been liberals' could have been intended by the speaker? Still, the assumption that the predicate is interpreted in a non-standard way[17] could work for some descriptive uses of the first person singular pronoun, as argued by Nunberg (1995). His crucial example is this:

(37)   *Driver holding a set of car keys*: I'm parked out back (Nunberg, 1995, 111).

The interpretative options are the following: 'I' could be interpreted descriptively as *the car* or 'to be parked out back' could be interpreted, metonymically, as denoting the property of *being a driver whose car is parked out back*. Nunberg (1995) proposes four tests aimed to show that the second option is the right one: (i) predicate coordination, (ii) apposition insertion, (iii) agreement, and (iv) substitution by descriptions or names. Let us take a look at these tests and assess whether they can be used to support the claim that the descriptive feel of certain utterances containing pronouns emerge from non-standard interpretations of the verb (based on Barrios 2011). Consider:

---

[17]I prefer to describe such cases as being non-standard rather than non-literal, because drawing the literal/non-literal is a theoretical matter. Ways in which this distinction can be drawn will be discussed properly in chapter 2. As I shall argue later, descriptive pronouns are not cases of non-literal meaning, but they might be classified as 'non-standard' in an intuitive way.

(38) *Predicate Coordination* (judgement in Nunberg 1995, 111):

    a. I am parked out back and have been waiting for 15 minutes.

    b. *I am parked out back and may not start.

Above, it seems that we can coordinate a structure which describes the referent of 'I', but if we try to co-ordinate a structure which involves the descriptive interpretation of the pronoun, the resulting utterance is infelicitous. Nunberg (1995) argues that this shows that it is the predicate that is not interpreted in its standard way. Moreover, consider the insertion of appositions.

(39) *Apposition Insertion* (Barrios, 2011, sect. 6) *A guest who has not been exposed to the weather to a parking valet*: I'm coated with ice.

    a. I, the blue Chevrolet, am coated with ice (based on Mount 2008).

    b. *I, the hostess' husband, am coated with ice (where the intended reading describes the guest's car).

In this case, we cannot insert appositions that modifies the standard interpretation of the pronoun. It is interesting to note that the apposition test appears to present the opposite pattern to the co-ordination test above (although, see Nunberg 1993, 39-40, ex. 87-88). However, Nunberg could claim that the apposition is somehow connected to the predicate. That is, (39)a would be interpreted as *John is the blue Chevrolet coated with ice*, where John is the referent of the pronoun and the predicate 'to be' is interpreted in a non standard manner (i.e. *to be the owner of*). Applying the same strategy to (39)b results in *John is the hostess's husband coated with ice*, which might be fine (although a bit weird) on its own and does not trigger a non-standard interpretation. Therefore, this choice of words would not be very well suited to convey the intended reading. Let us now turn to agreement phenomena.

(40) *Agreement* (based on the Italian equivalent of 'I'm parked out back').

    a. *Said by a male Italian speaker*: Io sono parcheggiato$_{[1st.person,sing.,masc.]}$ dietro.

    b. *Said by a male Italian speaker*: *Io sono parcheggiata$_{[1st.person,sing.,fem.]}$ dietro.

Based on the contrast above, Nunberg (1995) proposes the following argument. If the pronoun were interpreted descriptively (similarly to the interpretation of 'la macchina'$_{fem.}$ (the car) in Italian), the predicate 'to be parked' would have to agree with it. This is invisible in languages with a poor inflectional system, like English,

but more visible in languages with a richer system, as in Italian. Since agreement patterns with the gender of the referent of 'I', the pronoun seems to be interpreted as an indexical, leaving the predicate to be interpreted in a non-standard manner (e.g. *to be the owner of a car parked out back*). Finally, consider the test where the pronoun is substituted by a description or name that determines the same object in the context it would normally pick out.

(41)   *Substitutions by descriptions and names* (examples and judgements in Nunberg 1995, 111-2).

    a.   *The key I'm holding is parked out back.

    b.   The man with the cigar (Mr. McDowell, etc.) is parked out back.

In this case, Nunberg argues that if the pronoun is interpreted in a non-standard way, that is, via the identification of an individual in the context, which is then connected to a related description, we could replace the pronoun by an expression which determines a contextually salient object (the car keys) that is related to the relevant interpretation (*the car*) without any problems. As (a) indicates, this is not possible. The only possible substitution is by an expression that determines the same referent as determined by 'I', as indicated by (b). This shows that the referent of the pronoun remains constant, ruling out the possibility of a descriptive pronominal interpretation.

At this point it is worth assessing the strategy presented so far. First, it is important to point out that is applies only to a small subset of the relevant expressions (only instances of first person singular pronouns), so even if plausible the range of the approach is quite limited. Second, the approach is motivated by a series of tests intended to rule out descriptive interpretations of the pronoun. In this respect, the tests face some methodological issues. If we consider more examples, we see that the results they give are inconclusive. Let us begin by considering the following example of co-ordination:

(42)   *Celebrity whose statue is at Madame Tussaud's:* I'm (located) in floor 2 and mad at the visitors/paparazzi (Barrios, 2011, sect. 6).

In this example, it could be argued that the pronoun is interpreted as *the statue of a certain celebrity* in the first conjunct and as the specific celebrity who uttered the sentence in the second, Madonna, say. The conjunction test is simply inconclusive. This also seems to be the case for the apposition test. Consider:

(43)   *Said by a guest to the parking valet*: I, the/an excellent tipper, am coated
with ice (Barrios, 2011, sect. 6).

Here, the appositive clause seems to describe the guest who uttered the sentence and is an excellent tipper, but 'coated with ice' cannot possibly be predicated of the same individual, which triggers some sort of non standard interpretation.

For presentation purposes, the remaining two tests will be discussed in other parts of the thesis. The reason for this is that they connect to broader issues in semantics and pragmatics. The agreement phenomenon mentioned by Nunberg seems to reflect certain perspectival elements of language use to be described in the next section (on the extended properties of descriptive pronouns). As we shall see there, the tests cannot be used in an argument in the way suggested by Nunberg. Finally, the substitution patterns that Nunberg employs are an instance of a broadly Gricean picture on the distinction between two levels of information: one governed by conventional (encoded) information, another by norms on cooperative communicative behaviour. As we shall see in the argument against the non-conversational status of descriptive uses, chapter 2 section 2.2.2, the way Nunberg draws this distinction (following Grice 1967) faces certain problems and hence cannot be used in his argument.

The tentative conclusions of this section is this. Descriptive pronouns seem to be quite similar to instances of metonymy, where the concept that is related to the intended interpretation is provided visually rather than linguistically. Moreover, there is no knock down argument against the idea that first person singular pronouns can convey descriptive interpretations. Note, however, that there is also no argument against the idea that these expressions are interpreted indexically, whilst the predicates they combine with are interpreted in some figurative way (Nunberg, 1995). These positions, therefore, stand on equal ground until further interpretative properties are brought to light. Building on the review of the pronominal data and the different ways to capture the descriptive interpretations mentioned so far, it is now time to present the extended data of descriptive pronouns.

1.8 Descriptive pronouns: the extended properties

So far in this chapter, the core features of identification-based descriptive uses of pronouns have been reviewed as well as a variety of other pronominal uses. Roughly, descriptive uses seem to exploit the conceptual identification of an individual in the context in order to express descriptive information related to it. Such

information does not appear to be regulated by any grammatically given structure and rather depends on extra-linguistic processes. Bound and donkey uses of pronouns, on the other hand, seem to depend on information that occurs in previous discourse (antecedents), whose availability may be subject to an explanation that is grammatical in nature (i.e. structural constraints on antecedent representations like *c-command* and similar notions). In this section, we will take a look at the possible interactions between information that emerges from extra-linguistic processes (visual identification) and linguistic structure. In doing so, a new set of properties of identification-based descriptive uses will emerge.

To begin with, let us look at the relations between pronominal binding and descriptive uses of pronouns. As mentioned in section 1.4, bound pronouns require accessible antecedents in the same utterance that occur at a specific position in a hierarchical structure (e.g. tree representation). Reflexives ('self') pronouns require antecedents that are 'closer' to the position of the pronoun when compared to non-reflexive (bare) pronouns. Now, given that in identification-based descriptive uses, content emerges extra-linguistically, we can ask two questions. First, can descriptive uses be antecedents to bound uses? Second, can descriptive uses provide content for pronouns that are bound by other antecedents in the clause? With regards to the first question, consider a modified version of the car scenario (10), below.

(44)   *Ann and Matt pass by a sports car with the sentence 'I'm the greatest' written on its windshield*:

    a.   *Ann says to Matt*: He [*pointing to the car*] adores himself.

    b.   *Matt to Ann:* Yeah, he [*pointing to the car*] believes he is better than other people.

In this scenario, the first occurrence of the pronoun 'he' seems to be interpreted in each utterance as *the owner of the car*, which binds the reflexive pronoun 'himself' in (44)a and the second occurrence of the pronoun 'he' in (44)b. In short, identification-based descriptive uses can provide antecedents for bound pronouns in the same way indexical uses can[18]. Regarding our second question things are more complicated. Consider:

(45)   *Ann and Matt are in a wedding shop browsing through the items*.

---

[18]It is important to observe that binding can target non standard interpretations, for example we could imagine contexts where utterances of 'Every flower$_4$ adores herself$_4$' and 'Every flower$_4$ thinks she$_4$ will win the contest', where 'every flower' can be metonymically linked to girls with flower t-shirts, say, or metaphorically interpreted as *every pretty girl*, for example.

a. *Matt to Ann:* Every groom thinks his bride is worth dying for.

b. *Matt to Ann:* Every groom thinks *she/ ?his [*pointing to a wedding dress*] is worth dying for.

In (45)a, the quantifier 'Every groom' binds the NP 'his bride'. We could imagine that a similar interpretation could arise via the gesture towards the wedding dress in (45)b. If the pronoun of choice is in nominative form (e.g. 'she') the bound interpretation conveyed by the descriptive use (i.e. *his bride*) does not seem to be possible. I tried to come up with similar examples where this form of binding is established, but did not succeed[19]. But if we change to the genitive form (from 'she' to 'his'), the bound reading becomes much more salient. This might be due to the fact the genitive can be re-analysed in terms of definite descriptions that need contextual completion (Partee and Borschev, 1998), that is, 'his' can be re-analysed in terms of *the N R x*, where *N* is a place holder for a nominal (e.g. *book*), *R* a place holder for a relation between the unique satisfier of that nominal and the value of the pronoun (e.g. *authorship*), and *x* is a place holder for a variable or an index that can be bound or assigned to an individual. Given this structure, a bound reading is possible (i.e. *each groom$_4$ thinks the bride of $_4$ is worth dying for*). Still, some native speakers suggest that (45)b (with 'his' instead of 'she'), is only marginally acceptable. As we shall see shortly, Jacobson (2000) has an example which shows that simpler cases of descriptive uses may be bound by antecedents.

Given this inconclusive result, it is worth taking a look at the relationship between identification-based descriptive uses and other forms of dependency; more specifically, donkey anaphora. According to section 1.5, unbound dependencies (e.g. 'Every farmer who owns a donkey feeds it') seem to be more structurally relaxed than pronominal binding as they do not require an antecedent that occurs in a particular configuration in the structure (such as one involving *c-command*, say). As before, let us ask two questions. First, can descriptive uses be antecedents to donkey interpretations? Second, can descriptive uses provide content for donkey pronouns that co-vary with other antecedents? Apparently, they can. For the sake of simplicity, let us consider an example which may provide an answer to the second question.

(46)   *Dependent descriptive pronouns*

---

[19]Thanks to Robyn Carston, Nicholas Allott, Dirk Bury and Matthew Reeve (personal communications) for feedback. For reasons of space, examples that confirm this trend will not be stated here.

Andy and Martha are talking about the killing of Osama Bin-Laden as they stroll around Washington DC. They walk past the White House and hop on a departing bus (relevant dependencies in bold-face).

a.　　　*Martha*: So, do you think a country has the right to kill an unarmed man?

b.　　　*Andy*: Well...**Every time there is a war**, **he** [*pointing to the White House*] has tough choices to make.

Intuitively, 'he' in (46)b is interpreted as *the American president*, which depends on 'Every time there is a war' (an adverbial modifier). The utterance communicates that for each time a war is being waged, the American president in that time has tough choices to make. This shows that the identification of an individual in the context can provide content, which establishes dependencies on other expressions in the utterance. This phenomenon has been first noticed by Jacobson (2000, 89): a speaker who utters 'Do most faculty members deposit it [*waving a paycheck*] in the credit union?' conveys an interpretation where the paychecks co-vary according to the faculty members. These cases (known as *deep anaphora*) differ from cases such as (46) in one important aspect though. In the former, the concept used in the identification of the individual (e.g. PAYCHECK) is the one that figures in the interpretation, whereas in the latter, this relation is indirect: the concept used in the identification (e.g. WHITE HOUSE) serves as a gateway to the concept that figures in the interpretation (e.g. PRESIDENT), making these uses similar to metonymy. Nevertheless, descriptive uses (and also donkey and deep anaphora) display *Co-variation*, reprinted below:

Co-variation: the interpretation of the pronoun co-varies with another expression in discourse.

If the intuitions are correct, what seems to be explaining the co-variation in this case is some sort of time (or eventuality) variable: for every war time *t*, the president in *t* has tough choices to make in *t*. Stanley and Szabó (2000) famously argued that in order for such bindings to occur, the relevant variable (e.g. *t*) would have to occur as part of the grammatical representation (or logical form, LF) of the utterance or as a result of mandatory processes of the language faculty. Without going into details, Stanley and Szabó seem to take *grammatical representations* as the outputs of (i) a rule based system that (ii) is obligatory in nature. If this strategy is pursued, pronouns would contribute at least two kinds of variables: an individual variable (bound by quantification over individuals, e.g. 'Every man')

and a time/event/situation variable (bound by quantification over times, locations, eventualities, etc; e.g. 'every war-time'). In chapter 5, I propose a grammar model that can account for binding of the relevant event variables via *optional* rules (see Martí, 2006) that manipulate extra-linguistic content.

Going back to the first question, in addition to the fact that descriptive pronouns can depend on other expressions in discourse, these uses can also provide antecedents for donkey pronouns. Consider:

(47)   *Descriptive pronouns as antecedents* (relevant dependencies in boldface)

    a.   *Someone at a jazz festival:* In many concerts, if **she** [*pointing to a saxophone*] does not have enough power in her lungs, she ruins **it**.

    b.   *As a reply to why in certain rural parts of the country people do not drink alcohol:* In most villages, if **he** [*pointing to a church*] says something is a sin, then townspeople will avoid **it** (Klaus Abels, p.c.).

In (47)a, 'she' (in boldface) seems to be interpreted as *the saxophone player*. The interpretation co-varies with the interpretation induced by another expression in the utterance, namely 'many concerts', (i.e. the concerts would determine the sax players in them) and also provides an interpretation for the pronoun 'it', namely, *the concert at which a specific sax player performs*[20]. Similarly, in (47)b, 'he' is interpreted as *the priest*, which co-varies with the adverbial 'in most villages' (each village determining priests in them) and also figures in the interpretation of the pronoun 'it', namely, *the thing that the priest in a particular village said is sinful* (see the comments about the incremental nature of donkey pronoun interpretation in section 1.5 due to Evans 1977).

These examples carry interesting consequences. Donkey pronouns seem to display *Accessibility*: certain expressions can provide antecedents for (donkey) pronouns, while others cannot. If descriptive uses can be (accessible) antecedents for donkey pronouns, then any principled explanation of what counts as an accessible antecedent must include interpretations that are not provided linguistically. This in fact runs against a trend in the literature since Postal (1969) (and followed by Evans 1977; Kamp 1981; Kadmon 1987; Heim 1990; Chierchia 1992; Elbourne 2001, 2005 among many others) that assumes that the antecedents of donkey pronouns can only be available *formally*, that is, in terms of linguistic structure (to be further

---

[20]Similarly to what was stated in the previous footnote and in section 1.7, pronominal dependencies can exploit non-literal interpretations of the antecedents. For example, we could replace the first occurrence of 'she' and the associated demonstration by the description 'the saxophone', which would be interpreted as *the saxophone player* (a case of metonymy) and provide the value of the second occurrence of 'it' in the utterance.

discussed in chapter 3). In section 1.5, this trend was motivated by the contrast in examples, such as (26), reprinted here as:

(48)   a.   Every man who has a wife should bring her to the party.
        b.   #?Every married man should bring her to the party (intended reading: as (a) above).

Postal explained this contrast in terms of a grammatical constraint: antecedents must be (i) overt NPs (ii) that are not word-parts. However, the uses in (47) are not the only counter-examples to this constraint. As argued by many since Anderson (1971) (Cooper, 1979; Breheny, 1999; Jacobson, 2000; Patel et al., 2009), there are cases that have the form of (48)b, but seem to be perfectly fine. Consider a few examples (boldface indicate dependencies):

(49)   a.   When the baby **threw-up**, did you find any pencil eraser in **it** (Anderson, 1971)?
        b.   Every **iphone owner** uses **it** for browsing (Patel et al., 2009).

In (49)a, we have an antecedent that is not an NP, and in (49)b one that is part of a word.In summary, donkey pronouns as well as descriptive pronouns may have similar *Accessibility* properties. I also leave a full assessment of a possible unification between such uses for chapters 3 and 5. For now, I believe we have data that motivates the following interpretative property.

Accessibility (identification-based): Certain mental representations can provide the right interpretations for (descriptive) pronouns, while others cannot.

The similarities between this formulation of *Accessibility* and *Connection* (an explanation of how the mental representations used to identify the contextually salient individual supply the right descriptive interpretation) are remarkable. This may suggest that deriving the only component in the explanation of how the descriptive interpretations emerge boils down to a description of how the 'antecedent' representation emerges though processes of visual identification (e.g. an IPHONE concept making a MOBILE PHONE USER concept available, say). The suggestion is misleading for the following reason. Constraints of a *grammatical* nature may be involved in the right account of the *Accessibility* property of descriptive uses. Consider the following examples, which as far as I know, have not been discussed anywhere in the literature.

(50)   *Anaphoric impossibilities of descriptive pronouns* (boldface indicates potential dependencies)

    a.   *Said by someone*: '**A/The condemned prisoner** is traditionally allowed a last meal. **He** is also allowed to invite close friends for the execution'.

    b.   *Said by a condemned prisoner*: '**I** am traditionally allowed a last meal. * '**He** [*said by same speaker*]/ **He** [*said by different speaker*] is also allowed to invite close friends for the execution'.

    c.   *Said by someone*: 'If the Democrats had won the last few presidential elections, **the American Supreme Court Justices** might have been liberals. **They** would guard public interest better'.

    d.   *Said by Supreme Court Justice O'Connor*: 'If the Democrats had won the last few presidential elections, **we** might have been liberals. ***They** [*said by same speaker*]/**They** [*said by different speaker*]/?**We** [*said by same speaker*] would guard public interest better'.

In (50)a, the pronoun 'he' can refer back to 'a/the condemned prisoner', but in (50)b where the speaker's use of 'I' conveys *the condemned prisoner*, the speaker's use of 'he' cannot refer back to the interpretation. Similarly, in (50)c, the pronoun 'they' can refer back to the description 'the US Supreme Court Justices', but if O'Connor's uses 'we' to convey *the US Supreme Court Justices*, she cannot use 'they' to refer to that interpretation. She could, possibly, use 'we' to do so, at best, and convey something similar to a conjunction of descriptive uses, as suggested by Andreas Stokke (p.c.). This is surprising, given the fact that third-person pronouns 'he' and 'they' allegedly can depend on virtually any kind of information that is sufficiently salient, as all the examples of descriptive uses and the ones in (49) illustrate so well. The pressing issue now is to determine what explains this asymmetry.

There seems to be two factors at play here. First, as indicated by (50), if we change the speaker of the dependent pronoun, the anaphoric relations are possible (this was suggested to me by Nathan Klinedinst, p.c.). It seems then that the matching of personal features in the example above captures some sort of perspective from which the interpretation is determined. This kind of perspectival shift also seems to be one of the hallmarks of multi-participant utterances. Consider (from Purver et al. 2010):

(51)   *Multi-party dialogue*

    a.   *Speaker A*: Did you give me back...

b. *Speaker B*: your penknife? It's on the table.

c. *Speaker C*: I heard a shout. Did you...

d. *Speaker D*: Burn myself? No, luckily.

These data shows uses where 'me' in (51)a binds 'your knife' in (51)b and 'you' in (51)c binds 'myself' in (51)d. The examples have peculiar features. First, it shows that anaphoric dependencies can be established across speakers, which suggests that operations of the grammar can target a common stock of representations available in the context. It also creates certain difficulties for a string based view of grammatical operations (Chomsky, 1986), for, in one of cases above, the string would be formed by the concatenations of (51)a and (51)b, resulting in 'Did you give me back your penknife?'. Now if grammar were to describe this string, the bound reading intended by the speakers would be lost, indicating that the dependencies are established at a richer representational level: one targeted by the entities determined by the uses of the pronouns above, perhaps.

The multi-party dialogue above shows that certain anaphoric relations can happen across language users. In contrast, the anaphoric dependencies in the previous example (50) *must* be established via multiple speakers. That is, even though the descriptive use of 'I' and 'we' contributes information (*the condemned prisoner*, *the US Supreme Court Justices*, respectively) it seems that the point of view according to which these representations were determined blocks the possibility of referring back to these interpretations via the same speaker's use of third person pronouns, which determines their contribution via a different perspective. In short, although descriptive pronouns convey interpretations that are different from those determined by indexical uses, they remain *perspective-bound* and this must factor in any description of the *Accessibility* of mental representations that constitute the communicated descriptive content[21].

There are other ways to make anaphoric relations between O'Connors use of 'we' to communicate *the US Supreme Court Justices* and her subsequent use of 'they'. We could either imagine a situation where there is a long silence (or unrelated conversation), before the use of the third person pronoun or a situation where another conversational participant uses another NP to describe the justices *before* O'Connor's use of the pronoun to refer to them. Note that in both cases, the per-

---

[21]In this respect, descriptive uses share some similarities with logophoric pronouns (or uses) in some African languages (Hagège, 1974; Clements, 1975), where the pronoun refers to a point of view within the narrative that might be distinct from the perspective of the narrator herself. In descriptive uses, however, the perspective is always determined relative to the identification of the salient object (explanations of this perspectival character would then depend on an account of *Identification Dependency*), see Sells (1987) for the notion of *pivot*.

spective bound nature of the descriptive use is neutralised, as it were. No other contextual change, at least to my mind, could make these anaphoric relations better in any way.

In conclusion, identification-based descriptive uses do seem to have interesting *Co-variation* and *Accessibility* properties. These will be considered here as the *extended* properties of the data and explaining them should also be considered as *desiderata* for accounts of such uses.

1.9 Conclusion

In this chapter, we have laid out the core as well as the extended interpretative properties of identification-based uses of descriptive pronouns. This was partially achieved by comparing such uses with other pronominal data as well as figurative uses of language that bear some striking similarities with descriptive uses, such as metonymy. We seem to have reached both negative and somewhat positive conclusions.

On the one hand, identification-based descriptive pronouns cannot be reduced as mere instances of: (i) indexicality, (ii) pronominal binding, or (iii) genericity. The first case is ruled out, because deictic uses convey information about specific entities in the context, and descriptive uses do not. The second is ruled out because binding dependencies are established within local environments within structured representations, and the dependencies exploited by descriptive uses are non-linguistic in nature. The association with the third type of phenomenon is ruled out, because descriptive readings do not have the associated characteristics usually associated with generic uses (e.g. interpretations that express generalisations that admit exceptions, are kind-denoting, and presuppose salient well established properties).

Second, identification-based descriptive uses seem to bear strong similarities with donkey anaphora (both convey co-varying interpretations and display unusual accessibility patterns) and referential metonymy (both make an interpretation salient via the saliency of another, related, object). It seems that pursuing treatments of descriptive pronouns as a special instance of one of these two types of interpretation seems promising. Accordingly, this sets up the following thesis plan.

In chapter 2, I examine how descriptive interpretations may result from some form of figurative use of language, like metonymy. The exposition begins with Grice's treatment of non-literalness and moves towards alternative positions, such as Relevance Theory. This chapter has the secondary aim of introducing general (pragmatic) principles that may guide utterance interpretation, which will be used

in many other parts of the thesis.

In chapter 3, I look at possible ways in which descriptive interpretations result from interpretative processes similar to the ones underlying donkey anaphora. We will assess whether visual experiences, say, of seeing a footprint, can make mental representations that are indirectly related to the object seen (e.g. *the footprint maker*) accessible for integration with linguistic content. This will be done from a conceptual as well as an empirical perspective. Thus, as a secondary aim, this chapter aims to give an empirically supported description of the kind of representations that are visually accessible for integration with linguistically governed information.

In chapter 4, I assess explanations of identification-based descriptive uses of pronouns in the literature and conclude that they cannot explain for all the properties discussed in this chapter.

In chapter 5, I put forward my own account of the data. In many parts of this chapter (sections 1.4, 1.5, and this page), the ultimate integrability of an account of identification-based descriptive with an account of other pronominal uses has been challenged. The main reason for this rests on a standard assumption about grammar that sets linguistic processes aside other cognitive processes, such as the identification of individuals in the environment or the attribution of intentions to the speaker on the basis of her actions. The proposal made in chapter 5 rejects this standard division of labour. I believe that the descriptive interpretations discussed in this thesis provide good reasons for doing so, but other phenomena that raise similar challenges will be mentioned as we go along. The grammar framework assumed (i.e. *Dynamic Syntax*, Kempson et al 2001, Cann et al 2005) allows pragmatic processes (described by the relevance-theoretic principles in chapters 2, Sperber and Wilson 1995; Carston 2002, and the object-identification processes discussed in chapter 3) to supply pronouns with values that can be singular or descriptive (bound, anaphoric). Identification-based descriptive uses are cases in which the relevant value of the pronoun is supplied rather indirectly in the context (i.e. a form of anaphora with implicit antecedents), but can nevertheless be explained by the same means. Although this account incorporates pragmatic reasoning in its core, it employs a grammar that is able to capture the non-standard forms of dependencies and the perspective-bound nature of pronominal person. This nicely captures the extended properties of the data. Moreover, pronouns of different person features give rise to different interpretations. Some of them might not allow identification-based descriptive interpretations to arise at the level of an utterance's assertoric content (what-is-said, explicature). Thus, we end up with a heterogeneous account: some descriptive interpretations arise from the full integration of salient representa-

tions with the encoded meaning of the pronoun (explicature), while others depend on a more basic representation that gets selected as the value of the pronoun (hence counting as an utterance's implicature). These elements come together in a proposal that can account for the data without the shortcomings raised to other proposals in chapter 4. Chapter 6 concludes.

Chapter 2

Descriptive pronouns, levels of communicated information, and pragmatics

2.1 Introduction

In the previous chapter, it was argued that any account of descriptive uses must explain the properties of *Identification Dependency, Descriptiveness, Connection, Co-variation, and Accessibility*. However, it is unlikely that all these properties would hold at the same level of linguistically expressed information. For example, reconsider the following cases:

(1)   a.   *Said by Justice O'Connor*: 'We might have been liberals'.
        b.   *Near the White House in Washington, D.C.:* 'Every time there is a war, he [*pointing to the white house*] has tough choices to make' .

    Identification Dependency states that an individual must be contextually identified in order for the relevant interpretations to be conveyed: O'Connor in (1)a and the White House in (1)b. *Descriptiveness* tells us that the pronouns convey descriptive information. In the cases above, they contribute *the American Supreme Court Justices* and *the American president* to the truth-conditions of the utterances, respectively. *Connection* tells us that the concept used to classify the identified individual provides the building blocks of the descriptive interpretations. Note, however, that this is done in an indirect way. The concepts about O'Connor and the White House do not figure in what is communicated, but rather the concepts AMERICAN JUSTICE and AMERICAN PRESIDENT. Therefore, there is not a single level of information

that can sufficiently account for both *Identification Dependency* and *Descriptiveness*. In short, these interpretative properties appear to be explained by distinct pieces of information.

The appeal to different, yet related, levels of information is the cornerstone of an approach championed by Paul Grice (1967). He developed a theory (a 'logic of conversation') that exploited different levels of communicated (or, in his terms, speaker-meant) content in order to explain a variety of philosophical and linguistic problems. In doing so, he became the founder of modern inferential pragmatics: the field that explains the principles of human communication and the ways in which linguistic and contextual information fully interact. Given that the representations that constitute the descriptive pronominal uses are not linguistically given, appealing to pragmatics in order to explain the data is crucial.

Grice's picture involves two sets of norms governing what is expressed by an utterance: norms that emerge from the conventions associated with the words uttered and norms that emerge from the presumption that communication is a cooperative endeavour. This allows a description of the data where *Identification Dependency* is governed by a set of norms that determine a level of interpretation known as *what-is-said* by an utterance (or speaker), while another set of norms explain the other interpretative properties, at a level of interpretation known as the utterance's (or speaker's) *conversational implicatures*. For Grice, all non-literal uses of words are not a matter of the conventions associated with them and hence are better captured as instances of conversational implicatures. Given the similarities between descriptive pronouns and metonymy (see section 1.7), this picture suggests that *Descriptiveness* may be captured at the level of conversational implicatures.

Assessing a Gricean explanation of the data is important for many reasons. First, it tests the plausibility of a very natural explanation of the phenomenon. Second, it also tests the adequacy of the principles, postulated by Grice, according to which linguistic and contextual information interact. If the principles are not well grounded, alternative pragmatic frameworks should be sought. My attempt at such an assessment runs as follows.

In section 2.2, I will introduce Grice's pragmatic framework, sketch how an explanation of descriptive pronouns in these terms would go, and relate it to other instances of non-literal interpretations, such as metaphor and metonymy. Three arguments against a Gricean inspired explanation of the phenomenon will be raised: one concerning the general distinction between what-is-said and what-is-implicated (*the argument from pragmatic intrusion*), the second concerning views on non-literal meaning (*the argument against Grice's picture of non-literal meaning*), and

the third concerning descriptive pronouns, more specifically (*the argument against the conversational status of descriptive uses*, Nunberg 1993).

As we shall see, although the last argument is not sound, the first two arguments seem so, motivating alternative pragmatic frameworks. In section 2.3, *Relevance Theory* (Sperber and Wilson, 1995) will be introduced. Its presentation follows the structure laid out in the previous section. The framework's core machinery will be presented and then we shall proceed to possible treatments of the descriptive pronouns, occasionally comparing them with cases of non-literal use, such as metaphor and metonymy. As it turns out, Relevance Theory does not inherit the problems raised for Gricean pragmatics. It draws upon very general cognitive principles and, as a result, is much more flexible when it comes to capturing the data.

2.2 Gricean pragmatics

2.2.1 Grice on utterance comprehension

Meaning, What-is-said, and Implicature

In the lectures entitled 'Logic and Conversation', Grice (1967, 1975) observes two opposing ways of solving the problems that natural languages raise for philosophical analysis. One is pursued by philosophers of the *formalist* tradition; another is proposed by philosophers from the *informalist* tradition. Broadly speaking, these views are similar, respectively, to the non-intentionalist and intentionalist proposals in my discussion of the descriptive data in chapter 1 (see section 1.2).

The formalist approach aimed to devise formal languages (a *Begriffsschrift*, or a perspicuous way of writing concepts, according to Frege 1967) which captured certain properties displayed by language or even thought itself[1]. In this project, the notion of a natural language sentence (as opposed to an utterance) played an important role, for it highlights two interesting facts. Sentences tokens (relativised to context à la Kaplan 1977) are both generated from basic units (lexical types) via a recursive mechanism and are the minimal units for bearing truth and carrying inferences. Early formalists such as Frege and Russell aimed to set out a recursively built formal language capable of capturing the entailment patterns present in a natural language fragment (in terms of truth transmission from premises to conclusion).

---

[1]Besides Gottlob Frege and Bertrand Russell (1905), this tradition has been voiced, perhaps most prominently, by Carnap (1947) and the Vienna Circle (see Coffa, 1993), and recently reflourished in the works of Montague (1970a,b, 1973) and his followers, who, differently from early formalists, saw natural languages as being no different from formal languages.

On the other hand, the informalists (or ordinary language philosophers, e.g. Moore 1918; Wittgenstein 1953; Austin 1962) saw the formalist project as partial at best. Capturing the notion of meaning at the sentence level is misleading, for it misses a crucial aspect of natural languages: what people *use* them for. Austin (1962), for example, stressed the similarities between language and action based on the idea that the same sentence can be used by a speaker to perform various different acts. For example, the question 'Have you eaten today?' may sometimes be used as a genuine question, say, by a doctor wondering whether she should proceed in taking a blood sample or as an offer, say, on an occasion in which one is invited for dinner. The task of the informalist was to elucidate the meanings of natural language utterances by identifying the conditions under which they can be used by speakers for particular purposes, hence the famous Wittgensteinean slogan: *meaning is use* (Wittgenstein, 1953). One of the most influential groups in the informalist tradition, namely, the *Speech Act* theorists (Austin, 1962; Searle, 1969), sought to distinguish between the content of an utterance (locutionary content), which, when uttered with a certain communicative force (e.g. an assertion, command, request, etc...), imparts (illocutionary) content that elicits (and explains) certain acts. For example, when issuing a command, the speaker performs a certain act by which she obliges the hearer to perform another act. As we shall see, this idea of acts having certain consequences for the audience had a major influence over Grice's work.

However, as the formalist argues, the broadly conceived informalist project sacrifices systematicity. Conditions on *use* are so varied that they cannot be explained via the specification of the meanings of basic units plus the way they are put together.

Grice aimed to strike a balance between the two positions. On the one hand, he follows the ordinary language philosopher in assuming that utterances (not sentences relativised to contexts) are the carriers of meaning. In his seminal 'Meaning', Grice (1957) took that *non-natural meaning* or *meaning$_{NN}$* is a product of intentional action (contrasting non-intentional events that also carry information, such as smoke meaning fire, i.e. cases of *natural meaning*). Thus, expressing information by the use of a sentence (an utterance) is governed by:

M-Intentionality: By producing an utterance (or gesture) *u*, the speaker means that *p*, if (and perhaps only if) by uttering *u* she *intends*: (i) to produce a certain response (e.g. the belief that p) in the audience, (ii) such that the intention in (i) is recognised by the audience, and (iii) that the recognition in (ii) acts as a reason for the response (e.g. the belief that p) in (i) (in Grice 1957, 217-9, formulation due to Strawson

1964).

Now the speaker's utterance or gesture sets a specific task for the audience, namely, that of retrieving what the speaker intended to impart on the basis of her utterance and the communicative context. According to Grice, this process is largely a matter of non-demonstrative inference: utterance and contextual information provide evidence used by the audience in the attempt to identify the communicative (or m-) intentions of the speaker. The word 'largely' here is no accident. By using conventional signs (words or any kind of code) the speaker facilitates immensely the audience's task of recognising the relevant intention, because the information they carry almost coincides with the information she m-intends to impart (indexicals, ambiguous words, and a few other cases create some turbulence for this picture). In short, the use of linguistic conventions provide an almost direct route to the central content m-intended by the speaker.

Therefore, Grice follows the formalist in assuming that the conventional meaning associated with the word (i.e. types) immediately determines some of the information the speaker m-intended. Grice called this level of information 'what is said' (what-is-said, henceforth) by an utterance, which can be roughly equated with the content the speaker asserts. This allowed Grice to preserve the systematicity sought by the informalist. What-is-said would be determined by the standing meaning of the words used and the way they are put together (the syntax of the language), which, in turn, explains the entailment patterns associated with (a use of) a sentence. In order to preserve some of the insights of the formalist position, Grice assumed that what-is-said was governed by two principles:

Linguistic Constraint on What-is-Said: What-is-said is closely determined by the conventional meaning of the words used and their mode of composition (i.e. the syntactic structure in which they occur).

Truth-Evaluability of What-is-Said: What-is-said corresponds to the truth-conditions of an utterance [2].

---

[2]*Linguistic Constraint on What-is-Said*, and *Truth-Evaluability of What-is-Said* are Gricean incarnations of Carston's *Linguistic Direction Principle* and *Minimal Truth-Evaluability Principle*, respectively. These principles aim at constraining pragmatic contributions to what-is-said to, and only to, processes that are necessary either for supplying a value to a linguistically given variable or for supplying constituents necessary in order to arrive at something minimally truth evaluable. Grice did not recognise that considerations of conversational norms (i.e. pragmatics) were involved in deriving what-is-said (though see Neale 1992). Thus, Carston's principles would be held by people who maintain a Gricean perspective but recognise the role of pragmatics only for these two cases. This position is one variety of Semantic Minimalism. For discussion, see Carston (1988, 2002,

The first of the two principles aims at making what-is-said a matter of retrieving conventional information (pronouns, tenses, and ambiguous words require extra, contextual, information). The second aims at connecting linguistically determined information with truth, necessary for establishing the notion of entailment between propositions whose contextual contribution is as *minimal* as possible. Such a minimising role is essential not only to preserve the gist of the formalist approach, but also to articulate what-is-said so it can be used to impart other kinds of propositions: a point that preserves the gist of the anti-formalist approach. This can be illustrated by the following example, originally due to Grice (1981):

(2)   a.   He took off his boots and got into bed.
         b.   He got into bed and took off his boots.
         c.   He performed the actions in the order presented by (2)a.
         d.   He performed the actions in the order presented by (2)b.

This example presents the following tension. On the one hand, it is desirable to follow the formalist philosopher in preserving a truth-functional account of the contribution of 'and', that is, as a connective that takes two truth-values as input and returns a single truth value as output, so logical rules like conjunction introduction and elimination can be properly defined. But, on the other hand, there is the intuition that the utterances in (2)a and (2)b convey that the event described by the first conjunct happened before the event described by the second; something that falls outside the truth-functional treatment just mentioned.

Grice dissolves this tension as follows: (2)a and (2)b express the same proposition (*modulo* an identical reference assignment to 'he') at the level of what-is-said (given that $p \land q$ and $q \land p$ are truth-conditionally equivalent). However, this does not mean that what-is-said *exhausts* what the speaker communicates. For example, the speaker might have *suggested* different things by the different utterances. More technically, by *uttering* (2)a the speaker might have *implicated* (2)c and by uttering (2)b the speaker might have *implicated* (2)d. Grice describes the contents in (2)c and (2)d as instances of *generalised conversational implicatures*. *Conversational* implicatures emerge as a result of the expectation that communication is a rational cooperative enterprise: an expectation that is captured by the *Cooperative Principle* (CP, henceforth) 'Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in

—————————————

195-9).

which you are engaged' Grice (1975, 26). The term 'generalised' indicates that the implicated content emerges in most contexts, where it is natural to expect that the order of events is reflected in the order of saying. Generalised conversational implicatures contrast *conventional implicatures*, on the one hand, and *particularised conversational implicatures*, on the other. The former corresponds to contents that are calculated on top of what-is-said in virtue of the conventional meaning of the word. For example, an utterance of 'P but Q' and an utterance of 'P and Q' *say* (in Grice's technical sense) the same thing, but the use of 'but' also conveys as part of its conventional meaning that *P contrasts Q*[3]. Particularised conversational implicatures correspond to contents that arise from the same principles governing generalised ones (like the CP and maxims), but hold in particular contexts (rather than most). Grice (1975) considers the example of an utterance of 'He is very punctual and has good handwriting' as a reply to whether a particular person is a good philosophy student. The utterance conveys that the student is not very good, but this does not seem to generalise across contexts. Some have argued that the distinction between generalised conversational implicatures and particularised ones is not interesting, as they draw from the very same set of principles (explained below) (see Sperber and Wilson 1995, and Carston 1995, 2002, 111 for arguments). Given these three candidates, identification-based descriptive interpretations seem to be best captured as instances of particularised conversational implicatures (occasionally referred to simply as implicatures, from now on).

Consonant with the CP, there are more specific conversational maxims whose exploitation allows the speaker to convey information at an implicit level of meaning. The specification of the maxims and the details of how they work was one of the major accomplishments in Grice (1967, 1975, 26). I will summarise them under the heading *Communicative Cooperativity*[4], below:

Communicative Cooperativity: Make your contribution as informative (Quantity), truthful (Quality), relevant (Relation), and perspicuous (Manner) as required by the talk exchange in which the communicator is engaged in.

Above, the parenthesis indicate the category of each maxim associated with the CP. Specific sub- or super-maxims will be mentioned when necessary. Going back to the previous example, *Communicative Cooperativity* derives the temporal

---

[3]For details about the relation of conventional implicatures and the communication of higher level speech acts, see Grice (1989b, 361-2).

[4]This does not mean that *Communicative Cooperativity* exhausts all the principles that regulate conversation. One that is not tied to communication is the maxim 'Be polite'.

implicatures associated with 'and' in (2) in the following way. First, the audience retrieves what-is-said by the utterances in (2)a and (2)b (e.g. *Max took off his boots ∧ got into bed*, *Max got into bed ∧ took off his boots*, respectively), observes that the speaker is paying attention to the CP and associated maxims[5], retrieves necessary contextual information, and *infers* the intended implicatures (based on the category of manner, more specifically, the maxim 'Be orderly'). Such a reasoning exploits the following facts: (i) the speaker says two conjoined propositions, (ii) she has to utter them in a given order, (iii) the maxim of orderliness regulates possible relations between the conjuncts. In the case of (2)a, these steps can be illustrated by the following train of thoughts [6]:

(3)   *Sketch of the Derivation of a Conversational Implicature*

    a.   The speaker *said* that Max took off his boots ∧ got into bed.

    b.   The speaker is observing *Communicative Cooperativity* and knows that I know this.

    c.   Events happen in a certain temporal order, so by observing the sub-maxim 'Be orderly', the speaker communicates the order of events by describing them in a certain order.

    d.   The event described by the first conjunct of (3)a happened before the event described by the second one.

    e.   The speaker knows (and knows that I know) that I will assume this and did not do anything to stop me from arriving at this conclusion.

    f.   Therefore the speaker implicated (2)c (e.g. that Max took off his boots and *then* got into bed).

The abductive argument in (3) is an instance of the general *working out schema* for the derivation of conversational implicatures. I would like to point out that there

---

[5]There is an interesting discussion on whether Grice thought that a substantial part of his maxims would be involved in the determination of what-is-said. For discussion, see Grice (1957, 222), Neale (1992, 530), and Carston (2002, 104-5).

[6]Here is how Grice himself describes the derivation of conversational implicatures:

> A man who, by (in, when) *saying* (or *making as if to say*) that *p* has implicated that *q*, may be said to have *conversationally* implicated that *q*, provided that [i] he is to be presumed to be observing the conversational maxims, or at least the Cooperative Principle; [ii] the supposition that he is aware that, or thinks that, *q* is required in order to make his *saying* or *making as if to say p* (or doing so in those those terms) consistent with this presumption; and [iii] the speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in [ii] is required. (Grice 1975, 24 square brackets and emphasis mine)

is some arbitrariness involved in fleshing out the exact premises in it[7]. I justify my particular choice by the resemblance it bears to Grice's own example of the derivation of a conversational implicature (see Grice, 1975, 50).

The working out schema above provide a window into the properties of conversational implicatures. For example, the schemas is an instance of a non-demonstrative inference: the content of the implicature being the argument's conclusion. From this, an important property follows: conversational implicatures are *calculable*, that is, the inferred conclusion depends on the inferential steps that lead to it[8]. From the inferential nature of the derivation, another property follows. If utterances of different sentences *say* the same thing on different occasions and contextually salient information is sufficiently similar, then the speaker would convey the same conversational implicatures. That is, for Grice, conversational implicatures are *non-detachable* (with regards to what-is-said)[9]. Finally, given the non-demonstrative character of the inference, the speaker could easily provide evidence capable of cancelling the relevant implicature *without contradicting herself.* For example, in the dialogue above, the speaker could cancel the conclusion of the argument, namely, (3)f, by saying '...but I do not mean that these events happened in the order I just said' without any problems. The cancellability of conversational implicatures has been used to distinguish them from entailments, which allegedly hold at the level of what-is-said. For example, if the speaker *says* 'John has a red car' and continues by saying 'but I do not mean he has a car', she would contradict herself.

---

[7]For example, if one allows for contextual information to be shaped freely by the theory, one could flesh out premises in a way such that conversational implicatures come out as *deduced* conclusions.

[8] There has been some discussion in the literature on whether articulating the working-out schema must be exercised by the audience (a position taken by Neale 1992, 527 and Recanati 1993, 245) or only by the pragmatic theorist who offers a 'rational reconstruction' of the audience's behaviour (a position taken by Grandy 1989, 519). For discussion, see Carston (2002, 108). Moreover, there is the further issue of whether the audience must be *capable* of consciously articulating the inferential steps to derive the relevant implicatures (see Recanati 1993, 2004) or just providing some *post hoc* reconstruction of how the implicatures were arrived at: a position assumed by relevance theorists (Sperber and Wilson, 1995; Carston, 2002).

[9] Grice acknowledges that non-detachability is not a regular property of some implicatures based on the maxim of manner, that is, even if two sentences say the same thing, the manner of saying may convey different contents. Here is the crucial passage (noted by Korta 1997):

> Insofar as the calculation that a particular conversational implicature is present requires, besides contextual and background information, only a knowledge of what has been said (or of the conventional commitment of the utterance), and insofar as the manner of expression plays no role in the calculation, it will not be possible to find another way of saying the same thing, which simply lacks the implicature in question, except where some special feature of the substituted version is itself relevant to the determination of an implicature (in virtue of one of the maxims of Manner) (Grice, 1975, 39) .

In sum, Grice's insights on meaning take the following form. *M-Intentionality* characterises every instance of speaker meant information. *Linguistic Constraint on What-is-Said* and *Truth-Evaluability of What-is-Said* further constrain speaker meaning, determining the level of information represented by what-is-said. Finally, *Communicative Cooperativity* provides a set of conversational norms that are used together with what-is-said to derive the conversational implicatures meant by the speaker. Grice's theory of conversation employed different categories of implicatures: conventional, generalised conversational and particularised conversational. Grice took calculability, non-detachability and cancellability as 'tests' for the presence of conversational as opposed to conventional implicatures.

At this point, we can evaluate the suggestion made at the beginning of this chapter, namely, that of accounting for the properties of descriptive uses of pronouns at different levels of information. For the sake of simplicity, I will focus here on only the *core* data: *Identification Dependency, Descriptiveness* and *Connection*. As an illustration, let us consider the case where the speaker points to a footprint on the ground and says 'He must be a giant'. The Gricean might say that the a proposition is expressed at the level of what-is-said, which implicates the proposition that *the footprint maker must be a giant*. In a nutshell, *Identification Dependency* would be captured at the level of what-is-said, *Descriptiveness* at the level of what-is-implicated, and the non-demonstrative inferential relation between the two levels would account for *Connection* (which requires that the identification of the object provides access to the building blocks of the descriptive proposition). The structure of this explanation, as it turns out, is identical to Grice's proposal for non-literal word uses. In the next section, we shall take a look at Grice's picture of non-literal meaning, starting with metaphor (which Grice explicitly discussed), moving to metonymy (which, as suggested in chapter 1 section 1.7, is closely associated to the phenomenon we are interested in), and finishing with descriptive pronouns. In the section on metaphor, an argument against Grice's picture of non-literal meaning will be raised. Moreover, after the presentation of a Grice inspired account of the core data, other types of objections will be levelled: one specifically concerning the status of descriptive uses as conversational implicatures, another, more fundamentally, concerning the Gricean division between what-is-said and what-is-implicated.

2.2.2 A Gricean inspired explanation of the core properties

The basic picture of non-literal meaning

According to Wilson and Sperber (2002, 587), a Gricean explanation of tropes and figures of speech assumes the following explanatory pattern:

(4)   a.   Speaker meant information - whether presented literally or non-literally - corresponds entirely to conceptually structured propositional content.

        b.   Non-literal meaning differs from literal meaning in the ways it is generated. Literal meaning is governed by linguistic rules, whereas non-literal meaning is not.

        c.   Non-literal meaning is explained by systematic departures from literal meaning.

The pattern in (4) can be traced back to classical rhetoricians and is present even today in the analysis of many philosophers and linguists. The assumptions in (4)a to (4)c are fairly standard, but, as Wilson and Sperber (2002) point out, there is some flexibility when it comes to cashing out what such 'systematic departures' are. For Grice, they are explained within a two-stage picture. First, what-is-said is determined by the conventional linguistic meaning of the sentence uttered and a few features of the context. Secondly, *Communicative Cooperativity* and what-is-said by the utterance offer the means by which non-literal information can be conveyed. In short, although non-literal uses of words contribute contents they do not encode, their conventional meaning is used by the speaker to violate norms of rational communicative behaviour, which then makes the audience search for an interpretation which satisfies them. Thus, non-literal uses are instances of implicatures[10]. It is worth pointing out that the systematicity in this explanation emerges entirely from *Communicative Cooperativity*. Let us take a look now at how this specific picture applies to the case of metaphor (which has been the subject of much attention in the relevance theoretic literature, see Sperber and Wilson 1986; Carston 1997; Moreno 2007; Rubio Fernandez 2005; Wilson and Carston 2006; Sperber and Wilson 2008; Carston 2010).

---

[10]Lewis (1983) also made sense of such 'departures' within a two-stage picture: first there is the specification of the literal meaning of a sentence, in a language called 'literal-£', which is then linked up to a distinct meaning of the sentence, in a language called 'nonliteral-£', via some sort of rule. Wilson and Sperber (2002) point out that, unlike Grice, Lewis took non-literal meaning as a matter of ambiguity resolution.

Metaphor and the argument against Grice's picture of non-literal meaning

According to the rhetoric tradition, metaphor is a figure of speech that invokes *resemblance* between things. For example, by uttering 'Jane is a flower' the speaker invites the audience to focus on similarities between a person and a flower, thereby conveying something in virtue of such similarity. On the basis of this utterance together with the contextual assumption that flowers are beautiful and the resemblance relation between Jane and flowers, the utterance might convey that *Jane is beautiful*, as well as some extra imagistic content, which, according to certain romantic poets and critics, explains much of the beauty of metaphorical uses. As suggested previously, Grice's insight was to capture non-literal discourse in virtue of the mechanisms that regulate *saying* and *implicating*. In order to illustrate how his account works, consider the following dialogue (Grice, 1975, 34):

(5)   a.   A: Are you fond of Jane?

        b.   B: Oh she's the cream in my coffee.

(6)   *A's reasoning might go as follows:*

        a.   B has said that Jane is the cream in B's coffee.[*what-is-said* ??? by (5)b].

        b.   By *saying* (5)a, B would blatantly violate the maxim of truthfulness 'Do not say what you believe to be false' [from (5)a and *Communicative Cooperativity*].

        c.   B gives no overt reason for supposing that he does believe (5)a to be the case, and I would be completely irrational to do so, so I must suppose that B thinks something else other than (5)a [from (5)b and *Communicative Cooperativity*] .

        d.   B thinks (or at least wants me to think) that Jane is a delightful person.

        e.   B knows (and knows that I know) that I can work out these steps and did not do anything to stop me from doing so [from (5)a to (5)d, *Communicative Cooperativity*].

        f.   Therefore, B implicates that Jane is a delightful person [from (5)a to (5)e].

Obviously, the (particularised) conversational implicature in (5)f could be used by the audience to derive further implicatures, such as an affirmative answer to A's question above. It is worth noting that this explanation involves what Grice calls a 'blatant' violation of conversational maxims (i.e. a flouting of a maxim). As (5)c illustrates, in (5)a the speaker would flout the maxims 'Do not say what you believe

to be false' (i.e. Quality's *maxim of truthfulness*) and 'Be relevant' (Maxim of Relation). Blatant violations differ from covert violations, such as *lying* or *deliberate omission*, and the temporary *opting out* of the maxims, as in cases of fiction, in the sense that both speaker and audience know (and know that they know) that a maxim is being violated, most usually the maxim of truthfulness (see Grice, 1975, 32-34)[11]. This triggers the search for an implicature that would preserve *Communicative Cooperativity* by conveying truthful information (i.e. (5)f above). However, a problem emerges for this account. Grice took both what-is-said and what-is-implicated as instances of speaker meaning (what she m-intended), therefore (5)a cannot figure as a premise in the schema because 'Jane is the cream in B's coffee' is simply not m-intended by B (this is the reason behind the question marks in (5)a, see Neale 1992, sec. 2 and Wilson and Sperber 2002, 590). Now, if (5)a is absent, the inference simply does not go through. For this reason, Grice moved from *saying* to the notion of *making as if to say* (henceforth, making-as-if-to-say), that is, in cases of metaphor and irony 'nothing may be said, though there is something that the speaker makes as if to say' (Grice, 1975, 41). Let us reconsider the same schema but assuming this new notion, to be explained immediately below.

(7)   *A's reasoning might go as follows*

    a.   Jane is the cream in B's coffee [what the speaker made-as-if-to-say by (5)b].

    b.   By *making as if to say* that Jane is the cream in B's coffee, B would violate some maxim, such as 'Be relevant' (Relation) or standards of informativeness (Quantity) [from (5)a and *Communicative Cooperativity*].

    c.   B gives no overt reason for supposing that he does believe (5)a to be the case, so I must suppose that B thinks something else other than (5)a [from (5)b and *Communicative Cooperativity*] .

    d.   B thinks (or at least wants me to think) that Jane is delightful.

    e.   B knows (and knows that I know) that I can work out these steps and did not make anything to stop me from doing so [from (5)a to (5)d, *Communicative Cooperativity*].

    f.   Therefore, B implicates that Jane is delightful [from (5)a to (5)e].

In order to get this explanation off the ground, Grice had to cast making-as-if-to-say outside the domain of speaker meaning. To be more precise, the proposition

---

[11]Grice also considered cases where two or more maxims clash. In such scenarios, a maxim could be violated in order to preserve another one.

in (7)a is simply not m-intended by the speaker. To me, Grice's use of the notion suggests that making-as-if-to-say is closely related to pretence or entertaining a thought without much commitment (see the discussion about irony in Grice 1978).

Although appealing to the notion of making-as-if-to-say avoids the problem raised by non-literal *sayings* and the maxim of truthfulness, the move carries its own consequences. As Wilson and Sperber (2002, 590) argue, (7)b cannot count as a blatant violation of the maxim of truthfulness, since the speaker is not *saying* anything. For the derivation of the implicature in (7)f to go through, a maxim must be violated *blatantly*, which requires overt recognition by both speaker and audience. This form of violation seems to require something to be *asserted*, rather than merely considered, entertained, or *made-as-if-to-be-said*. For example, if A makes-as-if-to-say that he wants to murder B, B would not act as if his life were at risk. In short, apart from revealing an awkward sense of humour, A's words would not carry many psychological consequences for B[12]. On these grounds, (7)a cannot violate *any* maxim blatantly. As a consequence, (7)b cannot figure as a premise in the schema above. Again, the inference does not go through[13].

In conclusion, Wilson and Sperber (2002) show that no matter which notion is employed - saying or making-as-if-to-say - Grice's working out schema fails to deliver the conversational implicatures necessary to capture the speaker's intended meaning: the conclusion of what may be presented as *the argument against Grice's picture of non-literal meaning*. In the next subsections I will show how metonymy and descriptive uses can be regarded as an instance of the pattern described in this section and that they, therefore, also fall under this argument.

Metonymy

As mentioned in chapter 1 section 1.7, metonymy is a figure of speech in which a word (e.g. 'Plato') contributes with something which is related to its literal denotation (e.g. *Thaetetus*, the book by Plato) rather than the literal denotation itself (e.g. Plato, the philosopher). Such 'substitutions' are mediated by a variety of relations

---

[12]This might not be entirely right. Timothy Chan (personal communication) suggested that at some level the audience may act a bit suspicious. Still, in this situation we would not count such psychological effects as communicated by the speaker, because she did not intend to convey them.

[13]Nicholas Allott (personal communication) suggested to me that a Gricean could hold that propositions could be non-literally conveyed without a blatant violation, just mere violations of maxims pertaining to the categories of manner and relation. I acknowledge the point, but the move carries problematic consequences. First, we lose a way to distinguish non-literal meaning from lying and other forms of maxim violation. Secondly, we still would have to admit that an absurd proposition, e.g. *that Jane is actually the cream in one's coffee* is *meant* by the speaker. These are consequences that most Griceans would not be willing to accept.

- contiguity, causal, or part-whole (synecdoche)[14] - between entities, such as event - cause, object - possessor, artefact - creator, concrete entities - abstract properties they exemplify. Reconsider Nunberg's 'ham sandwich' case, reprinted from chapter 1 section 1.7):

(8)   *In a restaurant, the waiter warns the cook*:

   a.   Waiter: 'The ham sandwich is getting restless.' (Nunberg, 1978, 186)
   b.   proposition communicated: *The ham sandwich orderer is getting restless*.

Above, 'ham sandwich' contributes *ham sandwich orderer*, rather than *ham sandwich* to what is communicated by the utterance, where the substitution between these two thought constituents is licensed by some form of contiguity or causal relation holding between the entities they denote. Although Grice did not treat metonymy explicitly, his assumptions about metaphor apply easily here. In a nutshell, the proposal is this: by *saying* (8)a, the speaker would blatantly violate a conversational maxim, triggering an inferential process that results in the implicature in (8)b, which preserves the CP and maxims under *Communicative Cooperativity*. Could this account be along the right tracks?

The argument against Grice's picture of non-literal meaning (Wilson and Sperber, 2002) suggests that it cannot be. In (8)a, the speaker does not mean that *a ham sandwich*, the culinary item, *is losing patience*. If the speaker does not intend to convey this proposition, she cannot mean it, and hence it cannot be *said*. A crucial premise in the schema that delivers the implicatures is missing, and the inference cannot go through.

This motivates the move to making-as-if-to-say. Thus, instead of *saying* the statement in (8)a, the speaker makes-as-if-to-say it. However, this notion does not carry enough commitment to *blatantly* violate a conversational norm. In this case, there is no trigger for the search for an implicature capable of satisfying *Communicative Cooperativity*. Another premise necessary for the derivation of the implicature corresponding to (8)b is missing. Again, the inference cannot go through.

In conclusion, the same pattern that emerged from metaphor applies to cases of metonymy. However, to be fair with Grice, the formulation of his maxims were a first approximation to describe complex pragmatic phenomena. If he had not formulated the maxim of truthfullness using the technical notion of *saying* or if the notion of making-as-if-to-say had been better connected to the notions of assertion

---

[14]This would unify metonymy and synecdoche, which I assume for the purposes of this thesis.

and maxim violation, the problems described here could have been avoided. The problem is that what begins as just a reformulation to avoid a technical problem ends in a quest for the true principles that govern pragmatic reasoning. As we shall see in due time, such principles might lie outside a Gricean framework. Before presenting this, I move to an assessment of a Gricean explanation for cases of descriptive pronouns, which in some respects, resemble metonymy.

Identification-based descriptive pronouns and the argument from circularity

Under the rhetorician's understanding of figures of speech, descriptive pronouns can be approximated to metonymy. In (8), above, the interpretation *the ham sandwich*, which would be decoded from the words 'the ham sandwich', is replaced by the interpretation *the ham sandwich orderer*, via a contiguity or causal relation. Given the similarities between metonymy and descriptive pronouns pointed out in chapter 1 section 1.7, we could imagine a context quite similar to the one in (8), but in which the speaker uses a pronoun rather than the words 'the ham sandwich' in order to establish the relevant reading. Consider:

(9)   *In a restaurant, the waiter warns the cook*:

    a.   Waiter: 'He [*pointing to a ham sandwich on the counter with an order number under it*] is getting restless'.

    b.   Proposition communicated: *The ham sandwich orderer is getting restless*.

Here, the alleged 'substitution' between the interpretations *the ham sandwich* and *the ham sandwich orderer* takes place as before, but the former (*the vehicle* in the rhetoric tradition) is induced on the basis of perceptual information and the encoded meaning of the pronoun, rather than the concepts associated to 'ham sandwich'. On this basis, there are two Gricean options for accounting for descriptive uses of *third person singular* pronouns.

The first proposal is structurally identical to the metonymy case discussed previously. By uttering (9)a, the speaker would *say* that *the/this ham sandwich* (the culinary item itself) *is getting restless*. This (singular) proposition would count as a *blatant* violation of the CP and trigger the search for an implicature capable of complying with it. The audience would reach, as a result of an inference to the best explanation, the (descriptive) proposition in (9)b. It is important to note that this explanation accounts for the core features of the data. The requirement that an individual must be identified by the audience as belonging to a certain kind (*Identifi-*

cation Dependency) is captured by the proposition at the level of what-is-said. The intuition that the information communicated by such pronominal uses is descriptive (*Descriptiveness*) is captured by the implicated proposition in (9)b. Finally, the fact that the kind used in the identification somehow provides the building blocks of the descriptive proposition (*Connection*) is captured by the (inferential) dependency between what-is-said and its associated implicature.

Despite its apparent success, the proposal faces shortcomings. First, in the same way as the Gricean explanation for metaphor and metonymy, it cannot avoid *the argument against Grice's picture of non-literal meaning* (Wilson and Sperber, 2002). By uttering (9)a, the speaker *does not intend* to convey that *the/this ham sandwich* (the culinary item itself) *is getting restless*. Therefore, this proposition cannot possibly be *said*. A crucial premise in the schema that delivers the implicature in (9)b is missing and the inference cannot go through. Moving the relevant singular proposition from the level of what-is-said to that of making-as-if-say faces the usual problem. This notion does not carry enough commitment for a *blatant* violation of conversational norms. Without such violations, the search for an appropriate implicature that complies with *Communicative Cooperativity* is not triggered. Another ingredient necessary for the derivation of the proposition in (9)b is missing, and the inference cannot go through.

The Gricean, however, could make a second move. As stated earlier, the level of *saying* is closely determined by the conventional (lexical) meaning of the words uttered (*Linguistic Constraint on What-is-Said*). Pronouns raise a problem for this picture, because they rely on contextual information in order to make their truth-conditional contribution. Still, no matter how such contextual elements factor in utterance comprehension, the ham sandwich itself cannot be the value assigned to the pronoun 'he', given the clash between lexical features of the expression (e.g. masculine, animate) and features of the object itself (i.e genderless, inanimate). In short, if the *Linguistic Constraint on What-is-Said* is indeed assumed, no object in the context of (9) counts as a referent for 'he' in (9)a and, as a result, *nothing is said*. Thus, something like *By uttering* 'He [*pointing to the ham sandwich*] is getting restless' *the speaker m-intended to say nothing* would figure as a premise in the working-out schema. The fact that nothing is said, in turn, blatantly violates conversational maxims that contain the notion of saying in their definitions (e.g. 'Do not say what you believe to be false', 'Do not believe that for which you lack adequate evidence', Quality), which triggers the search for an implicature, in this case, *that the ham sandwich orderer is getting restless* (i.e. (8)b), that observes *Communicative Cooperativity*. Note that this possibility is not standardly available for the cases

70

of lexical metaphor or metonymy, because their conventional (encoded) meaning fully determines what-is-said (no contextual completion is necessary).

Although this move is possible, it still faces a problem known as *the argument from circularity* (first presented by Hugly and Sayward 1979). Whenever the speaker does not say something truthful, she must at least imply something true, so the talk exchange conforms to *Communicative Cooperativity* (supermaxim of quality: 'Try to make your contribution one that is true'). However, determining whether the speaker *is* cooperative, and not merely opting out of norms on rational communicative behaviour, requires knowledge of the specific implicature that makes a truthful contribution. Since implicatures are calculable on the basis of the what-is-said and the assumptions under *Communicative Cooperativity*, the audience must know that the speaker is cooperative in order to derive the relevant implicatures. Hugly and Sayward (1979) point out an argumentative circle: knowledge of the specific truthful (or informative) implicature is required for determining that the speaker is cooperative and speaker cooperativeness is required for the derivation of the relevant implicature. This move ought to be rejected on pains of circularity.

Finally, let us consider another option. Suppose that 'he' contributes *the male* (Cooper, 1979; Elbourne, 2005), so what-is-said by the utterance is paraphrased as *the male is getting restless*. Although this proposition would not violate any maxim and therefore lacks the associated problems, it fails to capture the core properties of the data. Identifying the sandwich would be irrelevant for conveying the propositions which we take to be intuitively communicated in the context. As a result, *Identification Dependency* and *Connection* cannot be accounted for. Moreover, as we shall in the discussion on how to distinguish the two levels of meaning put forward by Grice, there are reasons for assuming that the description *the ham sandwich customer*, instead of *the male*, is the one that belongs to the proposition expressed by the utterance.

In summary, it seems that a Gricean kind of explanation does not work out neatly for the cases of metaphor, metonymy and descriptive uses of third person pronouns. But does this point generalise to other pronominal forms? Apparently not. Let us consider the cases of first person pronouns (singular and plural) with the relevant singular propositions at the level of what-is-said and the associated communicated descriptive propositions at the level of what-is-implicated.

(10)   *Bill Clinton saying:* 'The founders invested me with the sole responsibility for appointing Supreme Court Justices' (Nunberg, 1993, 20).

        a.   What-is-said by (10):   *The founders invested Clinton with the sole*

> *responsibility for appointing Supreme Court Justices.*

    b.   What-is-implicated by (10):   *The founders invested the President of the USA with the sole responsibility for appointing Supreme Court Justices.*

(11)   *Said by US Supreme Court Justice O' Connor*: 'We might have been liberals' (Nunberg, 1993, 14-15).

    a.   What-is-said by (11): *The actual group of Justices that include O'Connor might have been liberals.*

    b.   *The/A group of US Justices (not necessarily including O'Connor) might have been liberals.*

Differently from the third person cases, the propositions that appears at the level of what-is-said in (10)a and (11)a make sense, could be meant by the speaker and do not violate standards of truthfulness. It might be true that Clinton, in virtue of being president, was invested with some powers by the writers of the American constitution and that there is a possible way in which the world might have turned out to be such that O'Connor and the other Justices around the early 90s are liberals. However, a familiar problem emerges for some of these cases, more specifically (11). The singular proposition in (11)a is not m-intended by O'Connor in a context where she describes a group of people that does not necessarily includes her, and hence cannot be used to blatantly violate a maxim in order to trigger a search for the relevant implicature. Moving to making-as-if-to-say would be of no help, because the notion does not carry enough speaker commitment to blatantly violate a maxim and thus no implicatures can be derived. On the other hand, the singular proposition in (10)a may be intended by Clinton, which renders the implicature based account of first person singular descriptive pronouns relatively acceptable. This may be extra evidence for the heterogeneity of the phenomenon of descriptive interpretations, mentioned towards the end of chapter 1.

Although a Gricean implicature-based explanation of descriptive uses might be viable for some cases of descriptive uses of pronouns (e.g. first person singular), it does not explain the other cases. This, by itself, motivates looking at alternative proposals. Before doing so, it is worthy taking a look at two other arguments against the Gricean account: one specifically targeting an implicature account of descriptive uses; another threatening the general Gricean distinction between what-is-said and what-is-implicated. The analysis of these arguments will bring extra motivation for the introduction of alternative pragmatic frameworks, which will be done immediately after.

The argument against the conversational status of descriptive pronouns

Nunberg (1993, 20-24,30) pointed out that descriptive interpretations do not seem to have the usual properties associated with *conversational implicatures*. His main argument concentrates on non-detachability, but it would also be worth considering the hypothesis of whether descriptive uses of pronouns are cancellable and/or calculable. Let us reconsider the Clinton case in (10), but suppose Clinton or another speaker used the proper name 'Clinton' instead of the pronoun 'me'.

(12)   *Bill Clinton or another speaker saying:* 'The founders invested Clinton with the sole responsibility for appointing Supreme Court Justices'.

a.   What-is-said by (12): *The founders invested Clinton with the sole responsibility for appointing Supreme Court Justices*.

b.   What-is-implicated by (12): *The founders invested the President of the USA with the sole responsibility for appointing Supreme Court Justices*.

Non-detachability, as a test for the presence of conversational implicatures, states that if two sentences can be used in similar contexts to determine the same information at the level of what-is-said, then the same conversational implicatures would arise. The implicatures depend on what-is-said (not the linguistic form used) and hence are non-detachable from it. Against non-detachability, Nunberg argues that, in this case, the level of information at (12)a does not implicate (12)b. Since conversational implicatures are non-detachable, descriptive interpretations of pronouns cannot be implicatures.

Calculability, as a test, states that what-is-implicated must be calculable on the basis of what-is-said. Against calculability, one could argue, based on the discussion in the previous subsection, that in some cases it is hard to know what the level of what-is-said would look like according to Grice's account. Consider again the example in (9), where the speaker points to the ham sandwich while uttering 'He is getting restless'. To say that the pronoun refers to the actual sandwich is at least problematic because of the feature clash between the gender and animacy features encoded by the word and the genderless inanimate nature of the referent. On the other hand, if we maintain that nothing is said, it is difficult to see on what basis the dependent implicatures would be derived. Without what-is-said, implicatures cannot be calculated. Therefore, descriptive uses do not seem to be instances of conversational implicatures (see the argument against Grice's picture of non-literal meaning and argument from circularity in the previous subsection).

Related to this point, cancellability, as a test, states that implicatures can be cancelled explicitly by the speaker without resulting in contradictions [15]. In the ham sandwich scenario just considered, imagine the speaker continuing (9)a by saying '... but I don't mean that the ham sandwich orderer, whoever he is, is getting restless'. In this situation it is hard to see what would be communicated by 'He [*pointing to the ham sandwich*] is getting restless'. Since implicatures are cancellable and this descriptive reading cannot be, otherwise nothing would be communicated, descriptive pronouns are not implicatures.

The conclusion of this partially hypothetical argument, initiated by Nunberg (1993), is the following. Since descriptive interpretations fail the non-detachability, calculability, and cancellability 'tests', they are not implicatures. Nunberg, based solely on the failure of non-detachability, actually goes further than that. Since implicatures are the model by which all pragmatic phenomena must be understood, descriptive interpretations must be treated 'semantically' and by this Nunberg means that the phenomenon must be accounted for by some sort of linguistic control. But, is this conclusion warranted?

There are two main reasons for a negative answer. First, Grice's 'tests' are neither sufficient nor necessary for detecting the presence of a conversational implicature. Secondly, they were also misused by Nunberg. Let us begin with the first point.

The status of non-detachability is unclear even for Grice, since for him the manner of *saying* can carry different implicatures (see footnote 9 page 62). For example, the utterances 'He took his boots off and got into bed' and 'He got into bed and took his boots off' (Grice, 1981), *say* the same thing, but imply different temporal orders of the events in the conjuncts. Thus, the maxim of manner is one principle which would guide interpretations that obviously violate the non-detachability test (the ordering in this case, in fact, might not be an implicature, as the *argument from pragmatic intrusion* suggests). In short, to expect that *for every* two different utterances that determine the same level of what-is-said in their respective contexts convey *exactly* the same implicatures is unrealistic. Arguably, non-detachability seems to work for some cases. But the fact that it does not work for others makes this test, by itself, an unreliable method for detecting the presence of an implicature.

As for cancellability, the observations are simply not true. If we consider the ham sandwich example again, it seems that the descriptive interpretation can be

---

[15]Relevance theorists argue that every pragmatic aspect of utterance interpretation can be cancellable. Some argue that there is a distinction between cancelling something so something else can be expressed and cancelling something so that nothing is meant. However, it is hard to draw this distinction in a theoretically unbiased was and therefore I will not discuss it here.

cancelled in favour of another interpretation, that is, the speaker of 'he [*pointing to the ham sandwich*] is getting restless' could continue her utterance, without contradiction, by saying 'I don't mean that the ham sandwich orderer is getting restless, I mean that the waiter, who needs the fries that accompany the sandwich, is getting restless' (see Stokke 2010, ch. 3 for more cases of cancellation). Thus, it seems that in some cases, descriptive uses can be cancelled. Given that cancellability is intertwined with calculability, it would be natural to expect that descriptive uses would also be calculable. In order to show this, I must make a small digression.

Descriptive interpretations are cancellable and calculable because they might be conveyed by (non-demonstrative) inferential processes. Thus the 'tests' not only determine the presence of conversational implicatures but all inferential aspects of utterance interpretation which might contribute to what-is-said. Consider reference fixing, which Grice took to contribute to the level of what-is-said, but did not talk much about. Now suppose a scenario in which Matt and Ann, office co-workers, talk about a party last week at the office. Ann says to Matt 'Bill got too drunk', to which Matt replies 'I thought your husband was away last week'. Now, it is perfectly fine for Ann to reply: 'Ohh I don't mean that Bill, my husband, was too drunk at last week's party, I meant that Bill, from the 4th floor, was'.This dialogue seems totally natural. Therefore, *contra* Grice, there might be processes that contribute to the level of what he took to be what-is-said but are distinct from the mere retrieval of the conventional meaning of words (as the argument from pragmatic intrusion, spelled out in the next subsection, suggests). *Contra* Nunberg, descriptive interpretations can be *calculable*[16]. But that does not mean that they are implicatures, because calculability and cancellability could underlie pragmatic processes involved in the delivery of what-is-said.

In conclusion, the argument from the non-conversational status of descriptive pronouns fails for the following reason. Some of the alleged tests for implicatures simply do not hold and the others that do hold not only characterise implicature derivation but also (pragmatic) processes that might contribute to what-is-said. Therefore, the tests cannot be used to support that a given phenomenon should (or should not) be treated as an instance of conversational implicatures. This conclusion paves the way to a more general objection against Gricean pragmatics: a challenge to the very way in which he distinguished *saying* from *implicating*. This is the topic of the next subsection.

---

[16]For an initial reaction to Nunberg's argument see Recanati (1993, ch. 16), for exposition and criticism of Nunberg's position from Relevance Theory, see Grimberg (1994); Powell (1998, 2003); Galery (2008).

2.2.3 The argument from pragmatic intrusion

For Grice, with the exception of reference fixing and ambiguity resolution, what-is-said is by an utterance is essentially determined by its conventional meaning, that is, what-is-said is *not* a matter of non-demonstrative inferences to speaker meaning. This claim, however, faces some problems. There seem to be cases where the proposition intuitively taken as *said* by the speaker contains elements that are not controlled by any linguistic item in the sentence that conveys it, thus violating the *Linguistic Constraint on What-is-Said*. To use a term from Perry (1986), such constituents would be 'unarticulated', because they cannot be traced back to any aspect of linguistic meanings and yet are present in the interpretation (thoughts) they induce. Consider:

(13)   a.   *Said by someone in Oslo:* It's snowing [in Oslo].

           b.   *Said by a car salesperson:* [That car has] Only 10,000 km!

           c.   *Said by Sam, a high school teacher:* Every student [in Sam's class] passed the exam.

           d.   *Said by Jane, after mentioning how hard she has studied for an exam:* I'm ready [to take the exam].

In these examples, the constituents in square brackets are part of the proposition the speaker intended to get across, yet these do not seem to be controlled by any overt linguistic expression in the utterance. An explanation of these cases in terms of particularised conversational implicatures seems unlikely. Take (13)b, for example; we cannot say that what-is-said by the speaker is *Only 10,000 km*, which implicates in the context *Car x has only 10,000 km*, because many other things have 10,000 km (i.e. the approximate distance between the equator and the north pole), rendering what-is-said communicatively inert, that is, what-is-said would systematically fail to capture what is specifically intended by the speaker. Some propose the existence of covert linguistic structure for some of the cases above (Stanley, 2000; Stanley and Szabó, 2000). This approach usually invokes pronominal-like variables of the appropriate type. I will not go into the details, as I believe that much of this debate depends on case by case analysis (see Carston, 2004; Martí, 2011). Rather, I would like to point out that this approach certainly would not be able to generalise to every instance. To illustrate, consider Cohen's (1971) classical examples:

(14)   a.   The old king has died of a heart attack and a republic has been declared.

b. A republic has been declared and the old king has died of a heart attack.

According to the Gricean explanation, the two utterances above have the same truth-conditions (they express the same what-is-said), and any information suggesting that the second conjunct is a consequence of the first emerges as a conversational implicature based on the maxim of 'orderliness' (whether it is a generalised or a particularised one does not matter here). However, if these sentences are embedded in conditionals, things look rather different. Consider other examples by Cohen (1971):

(15)　a. If the old king has died of a heart attack and a republic has been declared, then Tom will be happy.
　　　b. If a republic has been declared and the old king has died of a heart attack, then Tom will be happy.

As Cohen (1971) pointed out (see Carston, 2002, 109), these two utterances simply do not have the same intuitive truth-conditions (i.e. what-is-said). Grice (1989a) recognised this in his retrospective epilogue, but he did not recognise its implications for his way of distinguishing what-is-said (i.e. semantic content) from implicatures (i.e. pragmatic content). In addition to the intrusion of temporal ordering, utterances containing 'and' can be used to convey a cause-consequence relation. Consider:

(16)　She shot him in the head and he died instantly (Carston, 2002, 223).

In this example, the death would be understood as not only happening before the shooting, but as being a consequence of it. The word 'and' would not be interpreted as '...and then...', but as '...and, as a result...'. These examples are instances of *pragmatic intrusion into what-is-said* (a term due to Levinson 2000[17]). They represent a problem for the Gricean distinction between semantics and pragmatics because constituents that usually become available through appeal to conversational norms seem to enter the truth-conditions expressed by an utterance. There are two main ways to react to the problem: (i) reject it by postulating richer linguistic structure or

---

[17]The word 'intrusion' is sometimes used with a negative connotation, as if life would be better if such intrusions did not happen. A different approach, such as the one advocated in Relevance Theory (Sperber and Wilson, 1995; Carston, 2002) and by other truth-conditional pragmaticists (Recanati, 1993), have a different take on this. Pragmatic contributions to the proposition expressed are theoretically desirable because they capture the communicative intentions of the speaker in the right way.

lexical semantics, or (ii) reject it by insisting that the phenomenon is pragmatic but does not intrude into truth-conditional content proper. As an example of the first approach, Cohen (1971) associated temporal and causal features to the linguistic meaning of 'and' that would regulate the contents that enter the scope of the conditional. However, it is not difficult to come up with cases that cannot be accounted for by such features. For example, a use of 'Jane married and got pregnant' conveys a particular temporal ordering of the events, but 'Jane got pregnant and married' might convey some form of explanatory reading: the pregnancy was a reason for the marriage. Cohen's proposal does not account for this reading. Postulating extra linguistic features associated with conjunction (or the relevant expression) seriously risks overgeneration, as one could come up with examples that might not be properly explained by them.

Alternatively, one could argue that the alleged 'intrusive' constructions do not intrude the level of what-is-said but rather are a matter of implicature retrieval. Carston (2002, ch. 3) makes a series of objections against such possibility. First, as Cohen (1971) pointed out, it remains obscure how the alleged implicatures (defined as non-truth-conditional components of utterance meaning) can fall in the scope of logical operators such as conditionals and negation[18]. In fact, Cohen's use of conditionals and other operators became one of the most reliable tests for deciding whether a given constituent belongs to the truth-conditional content of an utterance (section 2.3.2 discusses this topic thoroughly). In conclusion, the association of what-is-said with a level of information entirely governed by linguistic rules does not seem to be empirically adequate. Conversational norms do seem to provide constituents that enter truth-conditions. As a result, the very notion of saying, as defined by the *Linguistic Constraint on What is Said* and the *Minimal Truth-Evaluability of What-is-Said*, must be revisited (see section 2.3.2).

2.2.4 Tentative conclusions

It is helpful to briefly list what has been achieved in this section. First, Grice's pragmatic framework was presented. It was initially motivated by a series of philosophical problems raised by the clash between the formalist and the informalist positions. This conflict was solved by a division between two levels of content associated with an utterance: what-is-said and what is-communicated. The former was defined by

---

[18]For an exhaustive argumentation against a Gricean explanation for the different readings associated with 'and' and in favour of pragmatic contributions to truth-conditional content more generally, see Sperber and Wilson (1995); Carston (1988, 2002); Recanati (1989, 1993); Green (1998), among many many others.

three constraints, namely, *M-Intentionality, Linguistic Constraint on What-is-Said*, and *Minimal Truth-Evaluability of What is Said* and the latter by norms of cooperative communicative behaviour under *Communicative Cooperativity* (the CP and maxims).

The main aim of outlining Grice's framework was to assess how it might explain the data raised by descriptive pronouns. Given the close similarities between these cases and metonymy (a non-literal or figurative use), the Gricean picture of non-literal meaning was introduced. Under a certain understanding of the Gricean position on descriptive pronouns, these can be seen as almost identical to metonymy, and thus the analysis is subject to the *argument against Grice's picture of non-literal meaning* (Wilson and Sperber, 2002). Under another understanding of the Gricean position, descriptive uses would say *nothing* and imply certain descriptive propositions in virtue of that. This possibility, however, is subject to *the argument from circularity* (Hugly and Sayward, 1979). Both objections apply to uses of third person pronouns. Some descriptive uses of first person pronouns, on the other hand, do not seem to involve the same difficulties. In addition to the two arguments just mentioned, we took a look at Nunberg's argument against the non-conversational status of descriptive pronouns. It simply stated that if implicatures are defined by Grice's 'tests' - non-detachability, calculability and cancelability - and descriptive pronouns fail them, then such uses cannot be a matter of implicature derivation. The argument, however, is not sound, which means that some descriptive uses of first person pronouns might be instances of implicatures.

From a more general perspective, we then took a look at Grice's distinction between what-is-said and what-is-implicated. As it was originally drawn, the distinction simply cannot accommodate interpretations that emerge from norms on rational communicative behaviour, but are part and parcel of the truth-conditional content of the utterance: the so-called 'intrusive constructions'. This motivates a serious revision of the way we carve the distinction between the truth-conditional content of an utterance and its implicit content. In the next section, we will present an alternative to the Gricean (or Gricean-inspired) proposal made here. The framework of choice is known as *Relevance Theory* (Sperber and Wilson, 1995) and with it we can re-assess the pragmatic possibilities concerning descriptive uses of pronouns.

2.3 Relevance Theoretic pragmatics

2.3.1 Relevance Theory and ostensive-inferential communication

In this section, I will present a theory that improves on Grice's by providing a model of utterance comprehension that draws on fewer but general and cognitively grounded principles and that satisfactorily account for non-literal uses, namely, *Relevance Theory* (RT) (Sperber and Wilson, 1995). Schematically, it departs from Grice's framework along the following lines: the notions of *saying* and *making-as-if-to-say* and the correlated *Linguistic Constraint on What-is-Said* and *Truth-Evaluability of What-is-Said* are utterly rejected, and *Communicative Cooperativity* is replaced by *Communicative Relevance*.

This brief presentation makes clear that RT preserves one of the central aspects of the picture of meaning sketched in Grice (1957): meaning *that p* can be described as making manifest an intention to convey *that p*. That is, like Grice, Relevance Theory assumes that much, if not all, of communication comes down to figuring out what the communicative intentions of the speaker are. Now, the challenge is to define the relation between the intentions of the speaker and the utterance by which they are revealed. RT assumes that this relation is explained by the following definitions:

(17)   *Ostensive - Inferential Communication*:

Communicating something (i.e. a proposition *p*) involves the speaker's production of a stimulus *u* with the intention:

a.   to inform an audience of something (Sperber and Wilson, 1995, 29), or more technically speaking, (Sperber and Wilson, 1995, 58) to make *manifest* a set of assumptions *I* (which includes *that p*) (*Informative Intention*).

b.   to inform an audience of one's informative intention (Sperber and Wilson, 1995, 29), or, more technically speaking, to make one's informative intention *mutually manifest* (Sperber and Wilson, 1995, 61) (*Communicative Intention*)[19].

---

[19]The technical definitions of manifest and mutually manifest information can be unpacked into the definition of the Ostensive-Inferential model, resulting in:

Ostensive-inferential communication: the communicator produces a stimulus which makes it mutually manifest to communicator and audience that the communicator intends, by means of this stimulus, to make manifest or more manifest to the audience a set of assumptions **I** (Sperber and Wilson, 1995, 63).

The distinction between communicative and informative intentions permits desirable moves in pragmatic theory. The speaker's *communicative intention* is fulfilled when her *informative intention* is *recognised* by the audience. The speaker's *informative intention* is fulfilled when the proposition she intends to convey (i.e. her informative intention) is not only *recognised* but also *believed* by the audience. Although the fulfilment of one's communicative intention normally acts as a reason for the fulfilment of one's informative intention, there is no necessary connection between the two. The fulfilment of the intentions is more or less independent from each other.

This double dissociation can be illustrated easily. When someone utters on the street 'Jesus is our saviour', one could get the message that is being put across, but one does not necessarily need to accept or believe what was just said. We recognise something is intended to be communicated, we just do not believe in it, that is, information is not transmitted. Conversely, the informative intention can be fulfilled without the fulfilment of the corresponding communicative intention. Suppose Kasha leaves the newspaper on the kitchen table, so John can be properly informed about the repercussions of the latest financial crisis, but by the time John sees the newspaper, Kasha had already left the house. He is indeed informed, but does not realise Kasha intended that. This situation is not a communicative one, because Kasha did not make her intentions overt, even though the information she intended to convey was transmitted.

The formulation of intentions according to the ostensive-inferential model in (17) brings another advantage over the Gricean: communication and overt information are connected in the right way, that is, communication *is* making information overt in a mutually recognisable way. This is achieved by using the notion of *mutually manifest* information in the formulation of the communicative intention [20].

Sperber and Wilson (1995, 39-46) define manifest information as assumptions that an individual (i) is able to represent at a given time and (ii) takes as true or probably true. A set of manifest assumptions constitutes an individual's *cognitive environment* (CE). Now, among the information represented in an individual's cognitive environment, there is information concerning who might have access to it. Sperber and Wilson (1995, 41) exemplify as follows. Suppose an individual gained

[20]Strawson (1964, 445-7) and Schiffer (1972) have pointed out cases in which the three clauses - (i) to (iii) - of Grice's *M-Intentionality* are satisfied and yet no information is intuitively communicated, because not all of the intentions of the speaker were overt to the audience (i.e. cases of 'sneaky' intentions, see also Neale 1992). These cases are properly ruled out by the way RT defines the informative and communicative intentions as well as the notions of manifestness and mutual manifestness.

knowledge *that p* by becoming a Freemason. By coming to know *p*, the individual also came to know that *all Freemasons have access to p*. This property permits the characterisation of *mutually manifest information*. First, the intersection of two (or more) cognitive environments constitutes a *shared* cognitive environment, as in a situation where two people, unaware of each other, look at a water fountain in a park. Second, shared cognitive environments containing information about who has access to it constitute *mutual cognitive environments*, as in a situation similar to the one just mentioned, but where two people become aware of each other by sitting on the same bench, say [21].

These notions nicely characterise communicative situations and the transmission of information that accompanies it along the following lines. The production of ostensive stimuli makes the information it carries manifest. Such information provides evidence of what the speaker wishes to convey, feeding a two-stage recognition model. The speaker's communicative intention is fulfilled, if the audience recognises the embedded informative intention. In this case, there is communication, but that does not necessarily mean that the speaker's informative intention is fulfilled. For this to happen, the audience not only needs to recognise the speaker's intention, but also to believe in it. In this case, information is transmitted through communication[22].

According to this model, the stimulus produced by the speaker has an evidential status: it provides a set of clues of what the speaker's intentions might be. Thus, the recognition of the relevant intentions depends on a process of non-demonstrative inference. Since this is also one of the fundamental tenets of Gricean pragmatics, it is worth noting potential differences between Grice's model and RT. According to him, the use of a linguistic code more or less determines the content of one of the propositions falling under the speaker's m-intention. Grasping what-is-said, in many cases, does not require inference. However, the argument from intrusion, the argument against the Gricean picture of non-literal meaning, and the argument

---

[21]The notion of mutual manifestedness emerged as an alternative to that of *common* or *mutual knowledge* (Lewis, 1969; Schiffer, 1972). By devising the notion of mutual manifestedness, Sperber and Wilson (1995) intended to capture communicative overtness in a way such that (i) it comes in degrees, and (ii) it does not carry any commitment to the fact that a given piece of information is mutually known for sure. When it comes to the latter point, the proponent of the mutual knowledge hypothesis might object by saying that much of the epistemology after Wittgenstein (1969) aimed to free the notion of knowledge from that of certainty.

[22]Epistemic vigilance is a cognitive ability aimed at selecting good information sources. Thus, although conditions on the fulfilment of the speaker's communicative intention are quite relaxed, conditions on the fulfilment of the informative intention highly depend on the information source or what is at stake. Since humans massively depend on communication to acquire information, they would be open to the risk of being often misinformed. Sperber et al. (2010) argue that hearers develop such ability to counter-balance this drawback.

from circularity raise serious objections to this picture. For RT, no such difficulties arise, as the notion of *saying* and the specific way in which it provides a basis for *implicating* do not exist in the theory. As a consequence, non-demonstrative inferences can contribute constituents at the level which Grice took to be what-is-said as well as to that of implicatures. But if this is the case, what then would guide these inferences, if not the CP and the maxims under *Communicative Cooperativity*? The answer is a particular notion of communicative *relevance*.

According to Sperber and Wilson (1995), *relevance* is a property of cognitive inputs (both external and internal stimuli). It is usually conceived in a comparative way: the greater the input's positive effects on an individual's cognitive system the greater its relevance; the lower the input's cognitive costs the greater its relevance. In short, the relevance of a stimulus increases with an increase in the effects/efforts ratio that it brings to the state the cognitive system is in. With this definition in mind, the following principle seems to hold:

(18)   *Cognitive Principle of Relevance*: Human cognition is geared towards the maximisation of Relevance, that is, to the achievement of as many positive contextual (cognitive) effects for as little processing effort as possible (see Sperber and Wilson, 1995, 260-266).

Note that the *Cognitive Principle of Relevance* is overarching in nature. Every aspect of cognitive life would be regulated by it. Sperber and Wilson also assume another principle; one that is more specifically tied to communication:

(19)   *Communicative Principle of Relevance*: Every ostensive act *communicates* (i.e. makes mutually manifest) the presumption of its own optimal relevance (see Sperber and Wilson, 1995, 260).

A stimulus is *optimally relevant* if it is (i) relevant enough to warrant the audience's processing effort and (ii) the most relevant one compatible with the communicator's abilities and preferences (see Sperber and Wilson 1995, 269, Carston 2002, 379). The Communicative Principle of Relevance warrants what has come to be known as the *relevance theoretic comprehension procedure*, stated as:

(20)   *Relevance Theoretic Comprehension Procedure*: Test interpretative hypotheses based on their order of accessibility, that is, follow a path of least effort, until an interpretation that satisfies the expectation of relevance is found.

An audience's expectation of Relevance comprises (i) a presumption of optimal Relevance *modulo* (ii) a specific set of mutually manifest assumptions. Thus, in many communicative situations, the turn taking amongst communicators creates more specific expectations about the relevance of incoming stimuli. For example, if Ann asks Matt a question, she will presume that Matt will not only utter something that is worth paying attention to (i.e. expectation of optimal relevance), but also that it provides the means to an answer (or conveys more positive cognitive effects than an answer would). This specific instance of the comprehension procedure makes two things mutually manifest: explicitly, the information the question linguistically communicates and, implicitly, the move from the presumption of optimal relevance to the current and more specific expectations about the relevance of the current ostensive stimulus[23]. In short, the difference between Gricean pragmatics and RT can be schematically represented by differences in the principles that govern interpretation. For Grice, *Communicative Cooperativity* guided only the derivation of conversational implicatures, whereas for RT every inferential aspect of utterance comprehension boils down to following the relevance theoretic comprehension procedure above.

By placing the relevant theoretic comprehension procedure in the *lieu* of contextualised appeals to the CP and the maxims under *Communicative Cooperativity*, Sperber and Wilson (1995) offer a way to reduce disconnected norms on rational communicative behaviour to a very compact and incredibly general set of principles. The procedure is uniquely involved in the many processes, such as reference assignment, disambiguation, and the supply of unarticulated constituents, that determine the proposition expressed by an utterance in addition to the associated implicatures. This raises an important issue: if the theory allows for pragmatic contributions to both levels of meaning, how can the proposition expressed be distinguished from its associated implicatures?

2.3.2   The Explicit-Implicit distinction

By rejecting the notion of what-is-said (and the associated *Linguistic Constraint on What-is-Said* and *Truth-Evaluability of What-is-Said*) while also accepting a Gricean-inspired intentional characterisation of speaker meaning, RT assumes a broadly underdeterministic picture of encoded information. This assumption, known in the literature, as the *Underdeterminacy Thesis* (Carston, 2002, 19), can be for-

---

[23]Note that the talk about expectations here plays a similar role to Grice's talk about the audience's expectation that the speaker obeys the CP and the conversational maxims.

mulated in its most general form as:

Underdeterminacy Thesis: The linguistic meaning of an utterance underdetermines every level of speaker meant information, even the proposition expressed (i.e. what-is-said in Gricean terms).

RT assumes that the proposition expressed corresponds to the *explicature* of an utterance (what the speaker *explicates* by uttering something) and, crucially, this is a communicated (speaker-meant) proposition. This level of utterance meaning is defined as follows (Sperber and Wilson, 1995, 182):

Explicitness: An assumption communicated by an utterance U is explicit if and only if it is a development of a logical form encoded by U.

Propositions that are speaker meant but not explicit belong to the implicit level: the utterance's implicatures. This allows the implementation of a machinery some-what similar to Grice's: the derivation of implicatures would be driven by the expectation that the stimulus made by the speaker is (at least, optimally) relevant and hence subject to the *relevance theoretic comprehension procedure*. Expectations of relevance would be essentially dynamic: information made mutually manifest by previous ostensive stimuli can make room for more finely calibrated presumptions, rather than the general expectation of optimal relevance. However, the problem with *Explicitness* is to understand what 'development of a logical form' means[24].

It is important to observe that RT assumes that human minds represent information in a language like medium or language of thought (LOT), as proposed by Fodor (1975). The syntax of this language must have the means to represent quantification and variable binding, as well as predicates and constants. Its basic building blocks are atomic concepts: concepts that are individuated by their reference relation to properties or individuals in the world and by formal attributes (i.e. a neural pattern that the concept supervenes on) necessary for distinguishing states constituted by co-referring terms, like 'horse' and 'steed'. Concept acquisition would be a matter of an individual mind 'locking on' to the relevant property. Although concepts are structureless, they relate to other pieces of information in the mind. RT assumes that there are three main entries that a given concept is associated to: the lexical entry, the encyclopaedic entry, and the logical entry. The first associates the concept, at a minimum, with a word form (spoken or written) used to express it and

---

[24]Sperber and Wilson (1995, 294) themselves consider some difficulties for the definition.

a syntactic form (e.g. a noun phrase or the lexical actions described in chapter 5). The second associates the concept with knowledge about the entities that it refers to (e.g. The concept DOG might activate one's knowledge that dogs are playful, good pets, affectionate, etc...). The third relates the concept to certain deductive rules, for example from the concept BACHELOR one may infer the concepts UNMARRIED and MALE [25]. Thus, the theory assumes that developing the logical form encoded by an utterance would result in a representation stated in something like Fodorian LOT.

In the next subsection, I will take a look at the criteria proposed in the literature for distinguishing explicit and implicit content. This would provide tools for determining whether that data raised by descriptive pronouns should be accommodated at one level or another. First, ways of distinguishing the explicit from the implicit that draw from conceptual analysis, hence referred to as 'off-line tests' for explicit content, will be discussed. Second, tests that are more empirically based will be considered. Since much of the next sections deals with arguments used to support some revamped notion of what-is-said, I will assume the more neutral terminology of 'proposition expressed' (which pragmatists take as speaker-meant). Once we have this distinction in place, we will be able to look at different proposals for non-literal cases and descriptive pronouns within RT.

Off-line Tests

Four proposals for delineating the distinction between explicit and implicit content will be discussed here: *Functional Independence, Availability, Embedding (Scope)* and *Agreement-based* tests. Following Carston (2002, ch.2), I assume that the intuition that these should not be interpreted as principles that guide interpretation like the principles of relevance, rather they are theoretical tools that might suggest one kind of semantic-pragmatic analysis over another. Let us begin with the first:

Functional Independence: 'The proposition expressed by an utterance should have a role to play, distinct from and independent of its implicatures, in the hearer's inferential processing, specifically, it should function independently as a premise in arguments' (Carston, 2002, 190).

The Functional Independence test is motivated by the thought that, on economy of effort grounds, the proposition expressed, construed as a communicated assump-

---

[25] There is some debate on whether all these entries need to be assumed by the theory. Kjøll (2010) is sceptical about the existence of logical entries, as the inference from BACHELOR to UNMARRIED may be licensed by encyclopaedic entries of special status.

tion, should have a role to play in the audience's cognitive life. The principle is deeply rooted on the cognitive principle of relevance. If an implicature contextually entails a proposition that allegedly corresponds to the explicature, processing the latter independently would bring processing costs without extra positive effects. Independence here means that the proposition expressed cannot be (i) entailed by any implicature of the utterance and (ii) should be used together with other contextual assumptions, some of which are *implicated premises*, to generate further implicatures. Implicated premises are contextual assumptions that are used in the derivation of further implicatures [26], for example (Sperber and Wilson, 1995, 34):

(21)   *Implicated premises and conclusions*

    a.   Peter: Do you want some coffee?

    b.   Mary: Coffee would keep me awake.

    c.   *implicated premise*: Mary does not want to be awake.

    d.   *implicated conclusion*: Mary does not want coffee.

According to this characterisation, Functional Independence would be wrong if one could find an implicated premise that entails the proposition expressed. Something that Carston (2002, 191) claims to have never encountered[27]. Although the test preserves the intuition that the proposition expressed is the *basis* on which implicatures are derived, something which is also present in Grice's working-out schema, it does not give a rock-solid tool for distinguishing the proposition expressed from implicatures, for the following reason. Suppose we wonder whether p corresponds to the proposition expressed or an implicature. One could go either way: p could be the proposition expressed, provided that independent implicatures are derived, or p could be actually implicated by another independent proposition, provided that this proposition plays some role in the agent's cognitive life. For example, consider an utterance of 'Tigers are striped' in a context where two people are travelling through the Siberian south-west during winter time. The proposition *that Caspian tigers are striped* could be regarded as the utterance's explicature provided that independent implicatures are generated (e.g. *Caspian tigers are easy to spot in such a white landscape*) or it could be regarded as the independent implicature of the more basic

---

[26] Implicated premises and conclusions may shape the audience's hypotheses about the speaker's explicatures. The development of the speaker's explicature as the utterance is processed on-line may, in turn, be used to revise the assumed implicated premises. This is known as the *mutual* adjustment of explicatures and implicatures, see the first table in Wilson and Sperber (2002) for details.

[27] The formulation here is different from the one given by Recanati (1989, 316) and avoids a counter example he raised to his own formulation of Functional Independence, see Carston (2002, 190-1) for discussion.

proposition *Tigers are stripped* (which does not entail that every sub-type of tiger is)[28]. Now, let us move to the Availability Test.

Availability Test: In deciding whether a pragmatically determined aspect of utterance meaning is part of [the proposition expressed], that is, in making a decision concerning [the proposition expressed], we should always try to preserve our pretheoretic intuitions on the matter (Recanati 1989, 310; Recanati 1993, 245; my square brackets replace 'what is said' in the original).

According to Recanati, availability to communicator's consciousness is a property that regulates a two-fold distinction: 'we have *distinct conscious representations* for [the proposition expressed] and for what is implicated by a given utterance: both are consciously accessible, and are consciously accessible *as distinct*' (Recanati 1993, 245). Thus, although the availability test appeals to intuitions about what has been communicated (i.e. an off-line test), it seems to make claims about human cognitive life: a property that is also found in Carston's Functional Independence. Recanati (2004, sect. 1.5) elucidates his position by saying that conscious accessibility to the proposition expressed might be underpinned by communicators' ability to pair utterances with the situation in which they are true. According to him, this ability is more basic than the ability to say or describe the conditions under which an utterance is true or false. Now, can this test do the job of distinguishing the proposition expressed by an utterance from its implicatures?

There does not seem to be a knock down argument against the availability test. In fact, as Carston (2002, 166-170) pointed out, intuitions about information communicated at the level of the proposition expressed or at the implicit level have always been present in theory building in linguistics and philosophy. The intuitions of communicators do offer a rough guide to theorising, but it is hard to see how they would provide a decisive method for classifying *p* as a proposition expressed or as an implicature. Experimental results have shown that scalar implicatures, e.g. 'some' implying 'not all', enter speaker judgements about the truth or falsity of the propositions expressed quite frequently, but not decisively (see Noveck 2004; Noveck and Sperber 2007[29]; Carston 2002, 166-170 and Carston and Hall 2011).

---

[28]This generic statement does not involve the usual entailments that hold in most contexts, and therefore, it could be argued that it is a bit artificial. I confess that finding an example suitable to my purpose was difficult, which suggests that the functional independence test may be aiming at an intuitive, or psychologically real, way to carve the explicit/implicit distinction.

[29]Noveck (2004) used scalar utterances like 'Some of the turtles are in the box' that should be evaluated against a range of scenarios: one in which some turtles were inside and some others outside, another in which all the turtles were inside the box. Subjects were asked to judge the utterances

This means that the proposition expressed and implicatures are not present to the subject as categorically distinct entities. Indeed, one should try to preserve one's pre-theoretical intuitions, but that is just good theory building. It seems then that the availability test does not provide a clear-cut method for distinguishing implicatures from the proposition expressed. Let us move to the next test: the Embedding (Scope) test.

Embedding (Scope) Test: A pragmatically derived aspect of meaning is part of the [proposition expressed] (and therefore, not a conversational implicature) if - and, perhaps, only if - it falls within the scope of logical operators such as negation and conditionals (Recanati 1989, my square brackets replace 'what is said' in the original).

The examples in Cohen (1971), used to motivate the argument from pragmatic intrusion, constitute an important precursor of the embedding test. As discussed in section on this argument, pragmatically derived aspects of meaning might contribute to the explicature of utterance's containing natural language logical connectives, such as 'and', or 'if..., then...'. The test also shares similarities with Kripke's 1972 rigidity test, which inquires into the type of content induced by an expression by embedding it under modal operators (e.g. 'if..., then...', 'necessarily', 'must', see chapter 1 section 1.3). Both Krikpe's and the embedding tests assume that the content which operators operate on corresponds to the proposition expressed. So far, I have not seen any decisive argument against the scope test in the literature (see Carston 2002, 191-7 for discussion) and it has been guiding many theoretical decisions regarding whether a contribution should be traced to the proposition expressed or to an implicature of the utterance. The scope test seems a promising method for testing the behaviour of descriptive pronouns. With this remark in mind, let us move on to other possibilities.

The final method for carving the explicit/implicit distinction out is the Agreement-based test. It has been more recently voiced by Hawthorne (2006). He proposed that the semantics of certain clauses, in this case, 'agrees that', can be used to specify the content of the utterances they embed: a form of content specifying operator[30].

---

true or false. In the scenario in which all the turtles were in the box, 'Some of the turtles are in the box' was judged true by 53% of the participants, indicating that for them, 'some' contributed its lexical meaning only, and false by 47% of the participants, indicating that, for these people, the scalar inference from 'some' to 'not all' entered the proposition expressed. For experimental precursors that aimed at testing the empirical plausibility of the availability test, see Gibbs and Moise (1997); Nicolle and Clark (1999).

[30]The indirect quotation test (Bach, 1994, 1999, 2001), the Inter-Contextual Indirect report test

Consider the predicate 'ready'. If Bill utters 'Magda is ready' thinking that *she is ready to leave* and John utters 'Magda is ready' thinking that *she is ready to get married*, one could easily imagine different contexts in which John and Bill could report (or say) that Magda is ready, but it seems difficult, if not impossible, for Bill and John to be in agreement.

I suppose Hawthorne takes his test to be a semantic one, that is, it inquires into the lexically determined content that is the objects of agreement. In this case, the predicate 'ready' comes out as context dependent (or relative to some other circumstance or parameter). However, one could consider the test as a possible tool for drawing the distinction between the implicatures and the proposition (explicature) expressed by an utterance, which would certainly admit pragmatically provided constituents that are not linguistically supplied. Roughly, this takes the following form:

Agreement Based Test: If two communicators A and B utter two sentences and one can report that, by the use of these sentences, *A and B agree that p*, then *p* captures the content expressed by the utterances of A and B[31].

Although agreement based tests seem to presuppose commonality of content, it remains dubious *which* level of content this is. As von Fintel and Gillies (2008) point out, agreement or disagreement can target a variety of levels, such as, presuppositions and maybe even implicatures conveyed by an utterance. For example, the following dialogue, adapted from Grice, seems quite natural to me:

(22)   *Implicature disagreement*

   a.   *Professor Wilson:* Mr. Green is applying for a position in the department. Professor Smith, you supervised him, right? Is he good?

   b.   *Professor Smith:* Hum... All I can say is that he has good handwriting and is very punctual.

(Cappelen and Lepore, 2005, 88) and the Collective Descriptions test (Cappelen and Lepore, 2005, 99) would employ either 'says that' (e.g. John said that *Bill is ready*) or any verb predicated of a collective subject (e.g. Both *John and Bill are ready*) to specify the relevant content (in italics here). Although these tests can be extended to cover cases of pragmatic intrusion in addition to context sensitivity of the indexical form, they are subject to the same criticism levelled against agreement-based tests. For these reasons, they will not be discussed separately here.

[31]Agreement/disagreement is one of the topics that sparked a recent debate on the nature of assertion, truth and structure of content, in addition to the notion of proposition expressed, there might be *thinner* notions of content, mere properties, for example, that play a role in understanding agreement related phenomena. For discussion, see MacFarlane (2003, 2005, 2007, 2009); Lasersohn (2005); Recanati (2007); Cappelen and Hawthorne (2009) .

c.   *Professor Jones:* I disagree, Professor Smith. He is very hard-working and has great intellectual potential.

It is difficult to see how the object of disagreement, namely, the proposition that *Mr. Green is not very good at philosophy*, would constitute the proposition expressed by the utterance(s) in (22)b. Like the Availability test, the Agreement-based test seems to be at the mercy of speaker intuitions. To the extent speaker intuitions vary, no decisive method for distinguishing implicatures from the proposition expressed is provided. To the extent that they converge, speaker intuitions seem to be very permissive with regards to what can be targeted. Therefore, this test, at least as formulated above, should not be relied upon.

In this subsection, I have looked at various tests for distinguishing explicit from implicit content, including: (i) Functional-Independence, (ii) Embedding (Scope), (iii) Availability, and (iv) Agreement-based tests. Intuitions seem to converge more evenly under the Embedding test and for this reason I consider it to be the best methodological off-line tool for drawing the distinction we are after. Now, I will take a look at some possibilities for looking at the explicit/implicit distinction from a more experimental perspective.

Empirically-oriented Tests

When it comes to the relations between experimental predictions and pragmatics, Grice's framework is a good starting point. According to it, the literal/non-literal distinction is explained in terms of the explicit (what-is-said)/implicit distinction, hence any predictions entailed by the latter automatically carry over to the former. Since one of the distinctive traits of implicatures is their dependence on the explicit level (what-is-said), it follows that they would be accessed only after the proposition expressed is articulated and so would be more costly to process. This prediction can be tested, but is it shared by other pragmatic frameworks, more specifically RT?

Not really. As mentioned in the previous section, ostensive stimuli might make implicit premises available before any hypotheses about the proposition expressed are made. Moreover, implicit premises can shape the on-line development of hypotheses about the proposition expressed (i.e. as part of the process of *mutual adjustment* see example (21), footnote 26). Consider the following dialogue (Wilson and Sperber, 2006):

(23)   *Mutual adjustment of implicatures and explicatures*

a.   *Peter*: Did John pay back the money he owed you?

91

b.   *Mary*: No. He forgot to go to the bank.

Peter's question in (23)a indicates his expectations: Mary must provide means to an answer. The partial decoding of Mary's utterance gives rise to the expectation that it will be relevant if it explains why John has not repaid her. As a consequence, *that forgetting to go to the bank, the financial institution, is a good reason for not repaying someone* is conveyed as an implicated premise. This assumption, in turn, is used to disambiguate the word 'bank' deriving an unambiguous proposition as explicature. This illustrates how *implicated* premises can be accessed before the proposition expressed. But what about *implicated conclusions*?

Imagine that Peter continues the conversation by asking if he could borrow some money from Mary. She answers with some stuttering 'Well... the thing... is...', at which point Peter could retrieve a negative answer to his question as an *implicated conclusion*, before hypotheses about the proposition expressed are fully formed. In sum, Gricean pragmatics predicts more effort in processing implicatures than the proposition expressed (what-is-said), while RT does not (for some empirical evidence for the relevance-theoretic alternative, see Matsui 1998, 2000 and Wilson and Matsui 2000).

In conclusion, RT cannot distinguish between these two levels in terms of differences in processing effort, at least not without bringing extra assumptions or detailed information about the specific communicative situation into closer scrutiny.

2.3.3   Relevance Theory, literal vs non-literal meaning, and descriptive pronouns

Literal and non-literal Meaning: an Overview

As suggested by Wilson and Sperber (2002, 587), Grice followed three aspects of the rhetorician's description of non-literal meaning (see (4)): (a) like literal meaning, non-literal meaning corresponds to conceptually structure propositional content, (b) non-literal meaning differs from literal meaning in the ways it is generated, and (c) non-literal meaning is explained by systematic departures from literal meaning. Grice fleshed out these assumptions in a very specific way. An utterance's literal meaning coincides with what-is-said, whereas what-is-said by non-literal uses blatantly violates conversational norms, thus triggering the search for implicatures that capture the intended interpretation and are capable of preserving such norms. In RT, options (b) and (c) (and more recently (a) as well, see Carston 2010) are rejected.

Before analysing how RT deals with non-literal used, I would like to mention yet another two ways in which they have been overlooked by Gricean pragmatics. The first of them concerns the fact that *every* non-literal interpretation results from a *blatant* violation of a maxim (truthfulness). If this is the case, utterances like the ones below seem hard to explain (see Sperber and Wilson, 1986):

(24)   *After a moderately tiring Friday, Bill says:*

    a.   I'm pretty tired. I deserve a beer (approximation).

    b.   I'm exhausted. I deserve a beer (hyperbole).

    c.   I'm a walking corpse. I deserve a beer (metaphor).

Strictly speaking, Bill's utterance in (24)a is not true, but it also does not seem to intuitively count as a *blatant* violation of the maxim of truthfulness ('Do not say what you believe to be false'). Without a blatant violation, the Gricean explanation cannot get off the ground. Moreover, such deviations from literalness seem to occur in *degrees* (see Wilson and Carston 2006, 2007, for a recent statement). The hyperbole in (24)b count as a greater deviation in comparison to (24)a and the metaphor in (24)c an even greater one in comparison to (24)b. Since Sperber and Wilson (1986), it has been pointed out that there does not seem to be a clear cut off point between approximations and hyperboles, and hyperboles and metaphors. This raises an important question. How would the maxim-flouting account deal with cases that only slightly deviate from what is literally expressed and the continuum formed among uses that deviate from literalness in different degrees? It seems that Grice simply cannot account for these facts[32].

Moreover, Grice's analysis faces one last shortcoming. When it comes to more poetic uses, such as an utterance of 'Joanna is a flower' that figuratively conveys that JOANNA IS BEAUTIFUL, the maxim-flouting account characterises the speaker's choice of words as a very clumsy and uneconomical way of conveying the relevant proposition. But if that is the case, something very important about non-literal uses is lost. Most tropes and figures of speech are used for a purpose. Many of them evoke powerful images and emotions, and are also useful tools for persuasion. According to the romantic critique of the classical rhetoric tradition, if the figurative interpretation (e.g. *Joanna is beautiful*) exhausts what is communicated, the ornamental or stylistic effects created by that specific choice of words are left

---

[32]One possibility is to deny that slight deviations such as approximations and hyperboles are cases of non-literalness. The loose use of language here would be a form of relaxation of the maxim of truthfulness: the speaker aims to say something true but does not set her standards of accuracy very high.

unexplained. In this spirit, Coleridge argues that a property of great style in speech or writing is:

> its untranslatableness in words of the same language without injury to the meaning. Be it observed, however, that I include in the meaning of a word not only its correspondent object but likewise all the associations which it recalls. (Coleridge 2007, XXII, cited in Sperber and Wilson 1986, 155)

We will now take a look at the pragmatic options to deal with non-literal uses within RT in the hope they will provide the grounding for possible understanding of descriptive uses of pronouns. So far, accounts of non-literal uses must: (i) capture their intuitive truth-conditions, (ii) explain the cognitive effects (images, empathy, etc.) that they give rise to, and (iii) explain the continuum of departures from literalness. As in the section on Grice, my comments on non-literal uses begin with metaphor, which paves the way to a treatment of metonymy and descriptive pronouns. It is important to note that accounts that try to accommodate non-literal contributions at the explicit level face the challenge of allocating inferential aspects of communication within the range of possible developments of the logical form encoded by the utterance (i.e. *Explicitness*).

Metaphor

RT rejects an important aspect of the Gricean tradition. Both decoding and inferring play a role in determining even the *literal* interpretation of an utterance. Thus, *non-literal* interpretations would not be generated by any different mechanism. On the assumption that the principles that guide such inferences are not violable (*contra* Grice's maxim-flouting account) and pragmatic constituents can figure at any level of meaning, RT could explain the cognitive effects imparted by metaphorical utterances at the explicit or implicit levels of meaning (or both). But, does anything motivate one explanation over the other? In subsection 2.3.2, we reached the conclusion that the embedding (scope) principle is the best diagnostic tool for making this decision. Let us then embed a metaphorical word use under a conditional and see how it behaves. I shall use Grice's example of 'cream in one's coffee' uttered in a context to convey that a person is delightful. Consider:

(25)   If Jane is the cream in your coffee, she deserves a nice birthday present.

Intuitively, the utterance above expresses that Jane deserves a nice birthday present, only if she is delightful or nice (as opposed to being actual cream). The embed-

ding test suggests that some metaphors contribute to the explicature of the utterance. Other metaphors that are conveyed by the lengthy development of a series of thoughts and images are not so easily testable[33] and hence will not be discussed here. This move, however, faces a problem. How can we conciliate the fact that a metaphorical interpretation (i.e. 'cream in one's coffee' conveying that a person is delightful) contributes to the explicit content *and*, at the same time, counts as a development of a logical form of an utterance (see *Explicitness*)? RT's current answer rests on the notion of *ad hoc concepts*.

The RT account of ad hoc concepts was inspired on Barsalou's (1983; 1987; 1993) work on categorisation. At the time, Barsalou reacted against the *Prototype* theory of concepts. This view was motivated by experiments suggesting that human categorisation skills, for example, deciding whether an individual (say, Fido) falls under a given category (say, of dogs), are subject to prototypicality effects. Smith and Medin (1981), for instance, argue that people do not treat all the individuals falling under the same concept equally. A robin would be classified as a bird faster and at a higher ranking in a classificatory scale than a duck would, even though both fall under the concept BIRD: a clear prototypicality effect. *Contra* Fodor, the prototype view assumes that concepts have internal structure, namely, a set of typical properties of the members that fall under them.

Barsalou's studies suggest that prototypicality effects clearly exist, but they cannot be explained by the prototype view, because categorising is a *context-dependent activity* and hence cannot be accounted for by a set of *stable* features. He found that prototypes vary according to different: (i) individuals, (ii) contexts (e.g. in a discussion about milking, cows and goats would be judged more prototypical mammals, whereas in a discussion about farming, horses and donkeys would), (iii) perspectives (i.e. swans would be judged prototypical birds from a Chinese stand point, but eagles would be judged prototypical birds from an American stand point), and, quite surprisingly, (iv) prototypes can be constructed on the fly (e.g. a wig, as an exemplar of things that could warm your head in a functional way). Barsalou argued that these (reviewed) prototypicality effects are best explained by an *occasion* specific mental representation: an *ad hoc category*.

Borrowing from Barsalou, RT devised the notion of ad hoc concepts (represented by an asterisk attached to conceptual representations) that are built along the following lines: words (e.g. 'coffee in cream') activate atomic (Fodorian) concepts (COFFEE CREAM), which are strongly associated with information stored in ency-

---

[33]Here is an example: 'The fog comes on little cat feet. It sits looking over harbour and city on silent haunches and then moves on' (*Fog* in Sandburg 1950 recently quoted in Carston 2010) .

clopaedic entries (e.g. CREAM IS CONSIDERED A NICE TREAT IN ONE'S COFFEE; HAVING COFFEE WITH CREAM IS A DELIGHTFUL EXPERIENCE), which are used in inferential processes that create an occasion specific sense for this word use: COFFEE CREAM*. Thus, a metaphorical use of 'cream in one's coffee' communicates, in a situation similar to Grice's example, the ad hoc concept COFFEE IN CREAM*; roughly meaning DELIGHTFUL IN A WAY THAT ONLY COFFEE WITH CREAM IS: WARM AND INVIGORATING, and so on. Such an explanation assumes a pattern like this (activation spreading from the word flows from left to right)[34]:

Table 2.1: Ad Hoc Concept Construction

Word	Lexical Concept	Assumptions	Ad Hoc Concept
'cream in coffee'	COFFEE CREAM	X BEING DELIGHTFUL IS A GOOD REASON FOR BEING FOND OF X; CREAM IS CONSIDERED A NICE TREAT IN ONE'S COFFEE; HAVING COFFEE WITH CREAM IS A DELIGHTFUL EXPERIENCE; etc...	COFFEE CREAM*

Note that this is not a sequentialist view. Given the evidential status of stimuli and the (relevance-driven) inferential nature of the comprehension process, many of the assumptions above would be made available by the addressee's current expectation of relevance in the specific context: one expects reasons for Jane deserving a present in the context of (25). This kind of effect-oriented top-down flow of activation is crucial to RT, setting it apart from purely associationistic (Lakoff

---

[34] In early RT, metaphors and other loose uses were analysed as relevance-driven implicatures. In the case of a metaphorical utterance of 'Jane is a bulldozer' that communicates that JANE IS TOUGH, the proposal assumed that a nonsensical proposition (i.e. that JANE IS (*actually*) A BULLDOZER) was *expressed* but not *communicated*. The expressed proposition, besides conveying the metaphorical interpretation as an implicature, would also control a range of implicatures, such as JANE IS ABLE TO ENDURE ANY SORT OF HARDSHIP *or* JANE IS RESILIENT AND NON-SENTIMENTAL at a low degree of manifestness and account for many of the extra cognitive effects conveyed by metaphors (e.g. imagistic content) as weak implicatures; see Sperber and Wilson (1986, 1987, 1995); Wilson and Sperber (1988) for details. Despite the fact that all the desiderata of non-literal uses was met, a strong asymmetry between two processes - narrowing (that contributes to the proposition expressed) and loosening - was predicted. This was not theoretically motivated: a point made by Carston (1997) and more recently by Wilson and Carston (2007). Moreover, the proposal has a strong resemblance to Grice's. In particular, the idea of expressing a proposition without communicating it could be regarded as a form of making-as-if-to-explicate, which could be subject to criticism similar to that against the notion of making-as-if-to-say.

and Johnson, 1980; Lakoff and Turner, 1989) and mixed inferential-associationistic frameworks (Recanati, 1993, 1995, 2004)[35]. Accordingly, the derivation of the interpretation of (25) has the following form (following Wilson and Sperber 2002).

(26)   *A's reasoning: metaphors as ad hoc concepts*

    a.   B has said to A, 'Jane is the cream in your coffee' [embedding of the decoded (incomplete) logical form of B's utterance into a description of B's ostensive behaviour].

    b.   B's utterance will be optimally relevant to A. [expectation raised by recognition of A's ostensive behaviour and acceptance of the presumption of relevance it conveys].

    c.   CREAM IS CONSIDERED A NICE TREAT IN ONE'S COFFEE; HAVING COFFEE WITH CREAM IS A DELIGHTFUL EXPERIENCE; HANGING OUT WITH JANE IS A DELIGHFUL EXPERIENCE, etc... [assumptions made salient by the decoding of the logical form in (26)a and the (specific) expectations of relevance in (26)b].

    d.   JANE IS THE CREAM* IN A'S COFFEE [since being delightful is a good reason for deserving a gift, inferred from (26)b and (26)c, and accepted as an explicature of B's utterance][36].

Schematically, the use of 'cream in one's coffee' activates the conceptual cluster COFFEE CREAM, which, in turn, makes assumptions stored in the associated

[35] Since the availability of representations is mediated by considerations of relevance, assumptions that are not contained in the entries associated with a concept are usually available for the construction of the relevant ad hoc concept. This amounts to a good explanation of the *emergent property* problem for metaphors, that is, the fact that, in many cases, the intended metaphorical interpretation (e.g. INSENSITIVE) is not represented in the encyclopaedic entry activated by the figurative use of the word 'cold' or 'bulldozer'. I will not discuss emergence issues here. For recent discussion, see Wilson and Carston (2007); Sperber and Wilson (2008).

[36] Another possibility is simply to say that the explicature simply is JANE IS DELIGHTFUL. I call this the substitution approach. It faces two problems. First, how would this proposition be a development of the logical form of 'Jane is the cream in my coffee'. Second, how would the ornamental value of non-literal uses be accounted for, given that the sentence uttered is a clumsy way of saying something quite simple. As proposals for metonymy will illustrate shortly, the first point may not turn out to be that problematic. With regards to the second point, the proponent of the simple account could argue that the ornamental value of non-literal uses are conveyed at a *nonconceptual* level of content. For example, by uttering the sentence above the speaker intends to *cause* feelings or images that usually accompany the event of drinking coffee with cream on the audience, as an illustration of Jane's delightfulness (as recently suggested by Carston 2010). This move requires a serious revision on assumption (a) of the basic scheme in (4), namely, *that communicated content is exhausted by conceptually structured propositional content.* Of course, much more would have to be said in order for the assumption that the conveyed imagistic content should be captured at the level of non conceptual content to do any serious theoretical work. Note that this could be regarded as a treatment of metaphor *qua* metonymy: one concept would simply be replaced by another.

encyclopaedic entries available for inferring the speaker's meaning. The overall accessibility of the assumptions in (26)c derives from the lexical meaning of the words as well as the demand for an interpretation that satisfies the expectations of relevance held in the particular context. They provide the means for the construction of an ad hoc concept - CREAM IN ONE'S COFFEE*. Now, since this process depends partially on the lexical concept activated by the word, ad hoc interpretations can be regarded as developments of the logical form of utterance (see *Explicitness*).

Although there might be other proposals for metaphor within RT (see footnotes 34 and 36) the ad hoc concept construction approach has good explanatory properties. For starters, it is in consonance with the results given by the embedding test and explains how explicatures containing ad hoc concepts can convey further relevant information based on the assumptions used to shape the content of the relevant concept (e.g. *Jane should get a present because she is delightful*). Moreover, the fact that the ad hoc concepts usually denote a very complex property (e.g. *cozy delightfulness*) explains why non-literal uses have ornamental value (e.g. imagistic content that may be considered as extra cognitive effects). Some would say that, as a result, ad hoc concepts are ineffable (see Carston, 2002, ch. 5): the properties they denote are hardly specifiable (i.e. not paraphraseable). Finally, a similar explanation can be devised for narrowing (or strengthening) and loosening (or broadening). Narrowing would be considered as the building up of an ad hoc concept whose denotation is a subset of the set denoted by the lexical concept. For example, suppose that Mary expresses her desire to get married by uttering 'I want to meet some bachelors' and by it she means that she wants to meet *not only* unmarried males, but also men with additional characteristics, such as being emotionally stable, heterosexual, trustworthy, and committed. Understanding Mary's utterance can be described as a process by which the concept BACHELOR is used as a starting point for the construction of an ad hoc concept BACHELOR* that denotes entities that have just these properties. Entities that fall under BACHELOR* automatically fall under BACHELOR[37]. Broadening is the converse operation. An approximation like 'France is hexagonal' can be taken as conveying an interpretation containing the concept HEXAGONAL*, which denotes entities that are not only strictly hexagonal but also sufficiently approximate to this shape. Entities that fall under HEXAGONAL automatically fall under HEXAGONAL*. Metaphors are just an extreme version of broadening (CREAM IN COFFEE* denotes cream and also individuals that are delightful in a very specific way) or a combination of broadening and narrowing (the concept is broadened to cover both

---

[37]This form of narrowing slightly differs from cases of lexical narrowing characterised by the addition of extra conceptual material, say, a use of 'P and Q' that conveys *P and as a result Q*.

kinds of entities and then narrowed to denote only delightful people, see Wilson and Carston 2006, 2007, for details). In this way, the ad hoc concepts approach also explains the continuum amongst different kinds of departures from literal meaning.

In conclusion, the relevance-theoretic proposal based on ad hoc concepts is not subject to the same problems raised by Grice's account and preserves all the intuitive properties associated with non-literal interpretations. In what follows, I will assess whether its nice explanatory properties carry over to cases of metonymy and descriptive pronouns. Like my discussion on metaphor, I will mention possible moves within RT for such cases.

Metonymy and Descriptive pronouns

Let us take a look at metonymy and identification-based descriptive pronouns in the same way as we did with metaphors. First, we will assess the diagnostics given by the embedding test and then move to ways of accounting the data within RT. Consider:

(27)   *Two waiters closing up the restaurant after work:* If the ham sandwich/he [*pointing to a ham sandwich*] left without paying, we won't get as much money as we thought.

Above, a use of 'the ham sandwhich' or 'he' while pointing to a ham sandwich falls within the scope of the conditional. The waiters will not get as much money as they thought, if *the ham sandwich orderer* leaves without paying. The embedding test motivates treating uses of third person singular pronouns (and possibly 2nd person singular and all the plural pronouns) as contributing to the explicit content of the utterance. This approach is also motivated conceptually. The absurd and non-sensical proposition THE HAM SANDWICH (*the culinary item*) LEFT WITHOUT PAYING does not seem to play any role besides conveying the relevant figurative interpretation. This point, originally due to Sag (1981), is captured in RT under the idea that the nonsensical proposition simply cannot meet the audience's presumption of relevance in this context and hence, on grounds of economy, it should not be even considered as speaker-meant (e.g. it does not play an independent cognitive role in the sense of Carston's Functional Independence test). Following this lead, Papafragou (1996) and Falkum (2010) argue that metonymy exploits minimisation of effort, whereas metaphor imparts more cognitive effects[38]. If such uses contribute

---

[38] It is possible to object to this point. Consider the utterance 'You should avoid marrying a sheep at all costs' (Gerrig, 1989), where 'sheep' refers to someone born in the year of the sheep. It could be argued that the utterance imparts both humour (via the literal interpretation of the words) and

to the explicit level of content, how then could *the ham sandwich orderer* count as a development of the logical form of a use of 'the ham sandwich' or 'he', above?

Given the unificatory basis of lexical adjustments laid out in Wilson and Carston (2007), one could try to treat metonymy and identification-based descriptive pronouns in terms of ad hoc concept construction, like the proposal for metaphors in the section above. Understanding the utterance above would proceed as follows: the words 'ham sandwich' or the gesture towards the ham sandwich activate the complex concept HAM SANDWICH, which, in turn, makes certain assumptions contextually salient, including: PEOPLE ORDER HAM SANDWICHES FOR LUNCH; ORDERING A HAM SANDWICH COSTS MONEY, etc. The availability of such assumptions also stems from the addressee's expectations of relevance in the particular context (suppose the utterance above is an answer to the question 'Did we get many tips today?'). Both the high activation status of these assumptions together with the fact that the HAM SANDWICH concept cannot combine coherently with the predicate 'to leave without paying' are as cues for the construction of an ad hoc concept. Following the proposal for metaphor, the metonymic use of 'ham sandwich' would convey the ad hoc concept HAM SANDWICH*. Its construction would involve an initial broadening of the concept so it denotes both certain culinary items and the people who ordered them, and then narrowed to denote only the latter.

This proposal, however, is subject to a number of objections. First, ad hoc concepts are created by cognitive operations that broaden or narrow the set denoted by the (original) lexical concept. This is what allowed the placement of metaphor within the continuum of departures from literalness mentioned earlier. As Falkum (2010, ch.6), following Wilson and Carston (2006, 2007), points out, the modulation of lexical meaning (broadening or narrowing) is warranted by implications that may feed a process of backwards inference that shapes the content of the occasion specific concept (as in the mutual adjustment of implicatures and explicatures). With regards to this specific case, broadening the ad hoc concept HAM SANDWICH* so it denotes both culinary items and individuals who ordered them does not seem to be licensed by any cognitive effects captured in terms of implicated premises or conclusions. Conversely, in the metaphor case, the modulation of FLOWER in 'Joanna is a flower' to convey that Joanna is beautiful, is underpinned by a similarity relation: both flowers and Joanna are beautiful things and thus both would be denoted under the initial modulation of the concept (which gets narrowed afterwards). The similarity relation explains a range of implications. Joanna may have a delicate beauty (as opposed to a feral kind of beauty), because flowers are

---

information about a person (via the metonymical interpretation).

beautiful in that way. In metonymy, there are no such cognitive effects and the contiguity relation does not play the pragmatic role played by the resemblance relation in metaphor comprehension. As a result, broadening the initial metonymically used word is not pragmatically licensed (huge effort involved and without any apparent gain in effects, although see footnote 38). Since the broadening is necessary for a further narrowing of the HAM SANDWICH* so it denotes only certain customers, the system is unable to create an occasion specific sense that captures the intuitive contribution of the words 'ham sandwich' in the utterance above.

There is yet another difficulty for the approach. The interpretations conveyed by metaphor and many other instances of broadening or narrowing have an ineffable nature, which is nicely captured by these occasion specific senses. The properties they denote are hard to specify (e.g. *beautiful in a flower-like delicate kind of way*). This does not seem to be the case for metonymies. The concept HAM SANDWICH* is exhausted by the description *the ham sandwich orderer*.

As for descriptive uses of pronouns, an explanation in terms of ad hoc concepts seems to be even more difficult. The reason being the fact that pronouns are not lexical words. Their contributions are determined in an essentially context dependent way. Still, one could argue that, in the utterance above, 'he' contributes an ad hoc concept MALE* (in virtue of the lexical material of the pronoun) or HAM SANDWICH* (in virtue of the gesture) to the proposition expressed by the utterance. Regardless of the form taken, the occasion-specific concept is roughly equivalent to the description *the ham sandwich orderer*.

This proposal in fact inherits the same shortcomings as in cases of metonymy. The conceptual representation would have to be first broadened to denote both culinary items and their customer and then just the latter, but since the first operation cannot kick off the ground (no cognitive effects are conveyed *modulo* huge processing effort), the concept cannot be shaped so it corresponds to the intuitive contribution the pronoun makes to the proposition expressed. In short, the ad hoc concept account does not seem to be a good way to explain how interpretations resulting from metonymical uses of words nor descriptive uses of pronouns. Alternative proposals must be sought[39].

---

[39]In the case of metonymy, Papafragou (1996) pioneered explicature based treatments based on metarepresentations. Borrowing from Kaplan (1989, 558ff), she explored the idea that metonymy involves a specific form of naming, that is, the explicature of an utterance of 'The ham sandwich left without paying' would be paraphrased as THE INDIVIDUAL APPROPRIATELY CALLED/CONCEIVED 'THE HAM SANDWICH' LEFT WITHOUT PAYING. The main problem I see with this account is that it builds into the explanation what it should explain. We want to know the circumstances under which such namings are appropriate. Falkum (2010, ch. 6) builds on this proposal in a way such that it does not inherit this shortcoming (as well as other difficulties). As Falkum's proposal does not carry

At this point it is worth considering *substitution-based accounts*. The idea is simple. In the right contextual setting, the lexical material of the word would be entirely replaced by a contextually salient mental representation. This proposal follows much of the spirit, if not the letter, of Nunberg's (1993) talk of *transfer* (after Fauconnier 1985, see also Recanati 1993, 2004) and the idea of certain mental representations serving as *triggers* for other *target* representations.

In the case of lexical metonymy, the substitution-based account would proceed as follows. The use of 'ham sandwich' activates the concept HAM SANDWICH, which in turn makes a range of assumptions contextually salient; including PEOPLE ORDER HAM SANDWICHES FOR LUNCH; ORDERING A HAM SANDWICH COSTS MONEY, etc. The availability of such assumptions also stems from the addressee's expectations of relevance in the particular context: the means for an answer to 'Did we get many tips today?' must be provided. By replacing the lexical concept HAM SANDWICH with the (phrasal) concept HAM SANDWICH ORDERED the audience is able to arrive at the explicature *the ham sandwich orderer left without paying*.

Note that the proposal does not inherit the problems raised by the ad hoc concept account, but it does have some problems of its own. The main issue is this. How can a *substitution* between conceptual material count as a development of the logical form of an utterance? If they indeed can count as so, the way encoded information constrains inferences to the explicature(s) meant by the speaker is trivialised. The main challenge is one of overgeneration. We explain the intuitive truth conditions of the metonymical use of 'ham sandwich' above, but fail to explain why the use of 'Jane' is a bad name for the dish she prepared in the example below:

(28)   a.   *Bill:* Who is London's best cook?
        b.   *John:* ?Mary won the cooking contest, although Jane was very tasty as well (judgement in Papafragou 1996, 143).

Let us abstract away from this difficulty for cases of lexical metonymy for the moment and consider the substitution-based proposal for identification-based descriptive uses of pronouns. Suppose that the use of 'he' introduces a variable-like entity $\mathbf{x}$[40] whose value is supplied contextually and that the gesture towards the ham sandwich makes a range of assumptions, including, PEOPLE ORDER HAM SANDWICHES FOR LUNCH; THE PERSON WHO ORDERED THE HAM SANDWICH IS NOT AT HER TABLE, etc., contextually salient. Replacing $\mathbf{x}$ by the concept HAM SANDWICH does not provide an overall intelligible interpretation that is compatible with

---

over to descriptive uses of pronouns straightforwardly, I cannot assess it fully here.

[40]My use of 'variable-like' here contrasts proper variables that are bound by quantifiers.

the audience's current expectations of relevance and that can satisfy the predicate 'leaves without paying'. Therefore, such substitution is not pragmatically licensed. However, these assumptions provide other concepts that can replace **x** in a way such that the resulting interpretation satisfies the current expectations of relevance. More specifically, the pronominal variable can be replaced by the value THE PERSON WHO ORDERED THE HAM SANDWICH (as an implicit antecedent, as it were), capturing the descriptive truth-conditions of the utterance in terms of the proposition *the person who ordered the ham sandwich left without paying*. In (some) descriptive uses of pronouns, the substitution operation is not as problematic as in the metonymy case, since pronouns may require this type of operation in virtue of their logical form. Indexical cases illustrate this well. The only difference is that 'substitution' in this case targets a singular (as opposed to a descriptive) mental representation.

Although the *substitution-based proposal* is much more plausible in the case of descriptive uses of (third person singular) pronouns than cases of metonymy, it still faces two challenges. The first concerns the precise specification of the lexical information encoded by pronouns. Above, we assumed that they encode a variable-like entity and, crucially, that it could be replaced by the description *the ham sandwich orderer*, which some (Montague, 1973; Barwise and Cooper, 1981) take to be a quantified term. The challenge then is to come up with a lexical entry for such expressions that allows for 'substitutions' that target both general and singular terms (as in indexical cases) and that accounts for the binding and the cross sentential dependencies. In short we need a grammar for these expressions so we can have a better idea of how the 'developments' of the logical form encoded by pronouns (see *Explicitness*) can derive all the types of pronominal interpretation reviewed in chapter 1, but nothing more. This will be crucial for an explanation of the extended properties of descriptive uses.

The second point concerns some accessibility puzzles that metonymy and also identification-based descriptive pronouns face. If in both cases, the alleged substitution would target an indirectly or non-immediately available concept (e.g. the interpretation *the ham sandwich orderer* is accessed through or mediated by the HAM SANDWICH concept), then it is mysterious why some cases, where the relevant interpretations are also indirectly available, cannot be interpreted properly. Re-consider the instance of donkey anaphora below:

(29)   *Accessibility patterns of indirectly available concepts*

      a.   Every man who has **a wife** should bring **her** to the party.

b. #?Every **married** man should bring **her** to the party. (intended reading: as (a) above)

Here the word 'married' would activate the concept MARRIED, which in turn would make the assumption A HUSBAND IS MARRIED TO A WIFE salient enough so it could provide a value for the pronoun 'her'. However, this does not seem to be possible. Note also that there are identification-based descriptive uses that display a similar accessibility pattern. Consider:

(30)   Every groom thinks *she/ ?his [*pointing to a wedding dress*] is worth dying for.

Here the gesture towards the wedding dress does not make *the bride* salient enough to become the selected value of the pronominal expression.

Finally, it is worth noting that the assumptions made here are also compatible with an heterogeneous treatment of the data. For example, as in the section on the Gricean-inspired treatment of descriptive uses, some cases may be instances of implicatures. Consider the embedding of a use of the first person singular pronoun under a conditional.

(31)   *Melvin, a condemned prisoner says*: If the incarceration system in this country changes, **I** will be traditionally allowed to order whatever I like for my last meal. (based on Nunberg 1993, 20)

According to Nunberg, when the proposition in the consequent is unembedded, 'I' contributes *the condemned prisoner*. However, we do not get this reading here. The utterance conveys that if the system changes, there would be a tradition according to which Melvin himself is able to chose whatever he wants for his last meal. There might be many factors at play here. For example, even in cases where a descriptive interpretation of 'I' gets under the scope of a conditional, it is impossible to rule out a reading where the pronoun is interpreted indexically (i.e. referring rigidly to Melvin) but the predicate receives a non-standard interpretation (see the discussion in chapter 1 section 1.7). Moreover, as in the section on Grice, some of the descriptive interpretations of the first person singular pronoun may be captured as implicatures. For example, Clinton's utterance of 'The founders invested me with the sole responsibility for appointing Supreme Court Justices', may convey that *The founders invested Clinton with the sole responsibility for appointing Supreme Court Justices* at the explicit level and implicate that *The founders invested the American President with the sole responsibility for appointing Supreme Court Justices*. These

propositions could be generated without any problem in RT.

The strategy that seems to have emerged is this. According to the Embedding (scope) test, some descriptive uses of pronouns contribute to the utterance's explicature. These are best explained under the substitution-based account mentioned above. Some other uses (e.g. first person singular) are diagnosed as contributing to the utterance's implicatures. RT (as well as a Gricean inspired account) account for these cases without any problems. Such proposal mixed both implicit and explicit contributions depending on the particular use and will be fully developed in chapter 5.

2.4 Conclusion

At this point, let us take stock and look at what we have accomplished in this chapter. Its aim was to introduce pragmatic frameworks that describe how extra-linguistic information interacts with encoded information in utterance comprehension. Given that descriptive interpretations are not provided by a linguistic antecedent or coincide with a immediate entity in the environment (like cases of deixis), appeal to pragmatics in order to explain the data is essential.

We begun by assessing Grice's framework, which laid out the general foundations for pragmatics. A Gricean-inspired account of descriptive uses of pronouns explains the core properties of descriptive uses of pronouns in two levels of communicated information, one determined by the linguistic meaning of words (what-is-said), the other by norms of rational communicative behaviour (what-is-implicated). Thus, the individual identified in the context, e.g. a footprint in the ground, contributes a certain propositional constituent at the level of what-is-said, i.e. FOOTPRINT. Since what-is-said by descriptive uses of pronouns violate conversational norms (the footprint cannot be a gigantic person), the search for an implicature that preserves them (i.e. *the footprint-maker is a giant*) is triggered. This explains why the identification of certain individuals in the context (*Identification Dependency*) is necessary for conveying (*Connection*) the descriptive truth-conditions that are intuitively communicated (*Descriptiveness*).

Various problems were raised to this approach. To begin with, Grice took what-is-said to be closely determined by the linguist meaning of words. Under this assumption, it is difficult to see how the personal pronoun 'he' could contribute THE FOOTPRINT to the level of what-is-said. Moreover, what-is-said is also determined by the communicative intentions of the speaker. Thus, the absurd proposition that *the footprint is a gigantic person* cannot be possibly meant, as it is hardly intended.

Without what-is-said, the relevant implicatures cannot be conveyed. Moving to the notion of making-as-if-to-say is not helpful, because it does not carry enough speaker commitment to blatantly violate a maxim. Again, the relevant implicatures cannot be generated. These and other shortcomings motivated the examination of alternative frameworks.

We then moved to a presentation of Relevance Theory. As one of its many departures from Grice, it assumes that inferential (pragmatic) processes can contribute the explicit or implicit levels of meaning. Such contributions are regulated *not* by norms of (idealised) rational communicative behaviour, but by overarching principles that regulate the processing of information and communicative overtness. In this framework, descriptive pronouns can induce readings that fall at the explicit or the implicit levels of meaning. The use of the scope test as a diagnostics for deciding the level to which descriptive pronouns contribute gave mixed results. The descriptive readings of some first person singular pronominal uses may be captured as implicatures, while other types of descriptive uses appear to contribute to the utterance's explicatures. The first approach (even in Gricean terms) explains the core data straightforwardly and does not carry any problems. With regards to the explanation of the data in terms of relevance-driven explicatures, two sub-types of proposals were analysed: an account based in ad hoc concept construction (borrowing from RT's recent treatment of metaphor) and the substitution-based account. The first faced serious conceptual problems. The second appears very promissory, but it faces two main challenges.

First, the account presupposes that a variable-like entity, allegedly encoded by the pronoun, can be replaced by descriptive information (a quantificational constituent) available in the context. The challenge then is to specify the linguistic meaning of pronouns in a way such that not only the identification-based descriptive interpretations are predicted, but also the other pronominal interpretations reviewed in chapter 1. That is, this is the challenge of describing what counts as a legitimate development of the pronoun's logical form (see *Explicitness*).

Secondly, the substitution-based account faces the challenge of explaining how visual processing establishes certain 'antecedent' representations that can provide the pronouns with the relevant value. The description of such processes should also aim to account for some accessibility puzzles regarding the unavailability of some interpretations. As the parallels between metonymy and descriptive uses make clear, the concept that is immediately available in the context (e.g. FOOTPRINT, upon seeing a footprint in the ground) is not the one that is selected as the semantic value of the expression, but rather provides a gateway to the relevant interpretation (e.g.

the footprint-maker). However, there are other cases where the relevant conceptual representations are also indirectly available, but the relevant readings cannot be established. As for example, certain cases of cross sentential anaphora (see chapter 1 section 1.5). For example 'Every married man should bring her to the party' does not make the concept WIFE salient enough to be selected as the pronoun's interpretation, even though it is closely associated with the MARRIAGE concept. A similar form of unavailability can also emerge for some descriptive cases. This is particularly puzzling. On the one hand, the fact that some identification-based descriptive interpretations are possible suggest that they can be generated as positive cognitive effects in some contexts. On the other, the fact that interpretations that rely on indirectly available concepts are sometimes unavailable suggest that such positive cognitive effects cannot be generated in some other contexts. How then is this duality possible, given that the principles underlying both types of interpretative attempts are the same and the encoded meaning of the pronoun remains stable across contexts?

The first of these challenges will be dealt in chapter 5, where a grammar formalism that plays a key role in explaining not only the core but also the extended data will be laid out. In the next chapter, I will take care of the second one. In order to overcome it, we need to supplement our pragmatic account with a description of specific visual processes by which certain mental representations can integrate with encoded information. Note that this makes sense only within pragmatic frameworks that are grounded on cognitive principles[41] as opposed to pragmatic frameworks that were motivated as solutions to certain philosophical problems (e.g. Grice's programme).

---

[41]This strategy may presuppose something akin to a competence/performance distinction within pragmatics. The cognitive principle of relevance (similarly to competence) is an overarching principle for information-flow regulation. Specific representational abilities (similarly to performance) recruit such a principle in the delivery of certain representational states.

Chapter 3

Descriptive pronouns and the representation of individuals

3.1 The Representation of individuals and linguistic understanding

At the end of the previous chapter, I motivated the idea that explaining utterance comprehension not only involves an account of overarching pragmatic principles, but also a description of the specific cognitive mechanisms employed by humans in comprehension. In this chapter, I engage with a debate about how certain views on the nature of visual representations bear on linguistic theorising, more specifically accounts of descriptive pronouns. According to chapter 1, these as well as other pronominal uses depend on the identification of a specific individual in the environment as belonging to a certain kind (*Identification Dependency*). Below, there are examples that illustrate that the reliance on the identification of an individual in context is not restricted to these expressions and consists in a more widespread phenomenon.

(1)   *Linguistic and conceptual inter-dependencies*

      a.   *A new faculty member picks up her first pay check from the mailbox. Waving it, she asks a colleague*: Do most faculty members deposit it [~~their paycheck~~] in the Credit Union? (deep anaphora, Jacobson 2000, 89).

      b.   *A visitor is leapt on by his host's dog and utters:* Mine [~~my dog~~] does the same (noun phrase ellipsis, Hankamer and Sag 1976).

      c.   *Chris utters to Mandy, who is looking for a box of cereal around the*

kitchen: [The cereal box is] On the top shelf! (subsententials, Carston 2002, 17)

d. *Johnny talking about a friend who just finished his PhD and making a circular motion to the side of his head*: He worked so hard he went (gesture) [crazy] (cases of 'showing', Wharton 2003).

e. *Pointing to the White House in Washington DC*: Every time there is a war, he [the president] has tough choices to make (identification-dependent descriptive pronouns).

The bracketed expressions in strikethrough are not uttered, but seem to be extracted from the context and become the constituents of the interpretations of the utterances above. They indicate conceptual representations that emerge from the identification of certain individuals as falling under a kind: the paycheck in (1)a, the dog in (1)b, the box of cereal in (1)c, the gesture in (1)d, and the White House in (1)e. In the face of such variety of linguistic phenomena, descriptive pronouns are just another case in which words require a suitable conceptual constituent, extra-linguistically provided, for communicating some information.

In order to illuminate the relevant analogies and dissimilarities between the use in (1)e and the other examples above, let us compare it to the deep anaphora case in (1)a. It is clear that both display some form of *Identification Dependency* in the sense of chapter 1, that is, in order for the relevant propositions to be communicated, the contextually salient individuals must be identified as belonging to a kind. In (1)a this is achieved under the concept PAYCHECK. In the case of in (1)e, this is achieved by the concept WHITE HOUSE. However, there is an important difference between the two. In the deep anaphora case, the concept that is employed in object identification is the one that is selected as a constituent of the proposition expressed. In the case of descriptive pronouns, the concept that is employed in object identification serves as a gateway for another concept, namely, AMERICAN PRESIDENT, that is selected as the value of the pronominal expression. This is what makes descriptive pronouns similar to metonymy (see chapter 1 section 1.7 and chapter 2 section 2.3.3).

Let us elaborate on the intuitive vocabulary used towards the end of previous chapters and assume that any concept that is *not* involved in the identification of the contextually salient object is an *indirectly salient or available* concept. Following this assumption, the concept AMERICAN PRESIDENT is not directly available, because it is not involved in the identification of the White House in (1)e, whereas the concept PAYCHECK is directly available in the context of (1)a, because it is involved

in the identification of the object waved by the speaker. This form of direct involvement can be conceived as concept tokens that emerge as a result of the detection of instances of certain properties in the environment. This terminology allows us to formulate a hypothesis about how non-linguistically available information and encoded information interact. More specifically, it allows the formulation of the *Cross Modal Integration* (CMI) question, below, which this chapter aims to answer[1]:

(*CMI Question*): Can indirectly salient concepts contribute to the proposition expressed by an utterance (i.e. count as a development of the logical form of an utterance)?

On the basis of the intuitive truth-conditions of (1)e, it seems that the CMI question deserves a positive answer. However, there are cases that raise difficulties for it, as, for example, the examples that motivated formal treatments of cross-sentential anaphora in chapter 1 section 1.5 and exemplify the *accessibility puzzles* mentioned in the subsection on descriptive pronouns in chapter 2 section 2.3.3. Consider again:

(2)   *Accessibility patterns of indirectly available concepts*

    a.   Every man who has **a wife** should bring **her** to the party.

    b.   *Every **married** man should bring **her** to the party. (intended reading: as (a) above)

Above, the utterance in (2)b contains a linguistic expression (e.g. 'married') that would be able to *indirectly* single out an individual to establish the intended interpretation of the pronoun (e.g. *the wife*). However, it is only with an antecedent that *directly* singles out the relevant discourse entities, as in (2)a, that the anaphoric link can be established. Note that in these cases, identification would be achieved via a linguistic expression, rather than via visual processes (the ways in which linguistic and visual identification can be unified will be taken up in chapter 5). Taking this point for granted for the moment, one could argue, based on the examples in (2), that indirectly available concepts cannot contribute to the proposition expressed. Let us develop this suggestion more clearly, relating it to some background discussion on anaphora.

---

[1] Elugardo and Stainton (2003), for example, are among the few in the philosophical and linguistic communities who detail how actual mechanisms of visual representation deliver the content that integrates with encoded information in substentential utterances.

In chapter 1 section 1.5, we looked at the first proposal in the literature for accounting for the asymmetry in (2): a condition on anaphoric dependencies known as the *anaphoric island constraint* (Postal, 1969). It states that only (i) overt NPs that (ii) are not part of words can serve as antecedents for anaphoric pronouns. The constraint would rule out cases like (2)b and many others (e.g. cases where '**McCarthy**ites' cannot serve as an antecedent for interpreting 'his' as *McCarthy's*). This shaped what Kadmon (1987) later called the 'formal link condition': the imperative according to which the relationship between pronoun and antecedent must be described formally, that is, in terms of the architecture of grammar. The move, then, can be summarised as follows. On the basis of certain intuitions held by communicators, a condition on human cognitive architecture was postulated: certain dependencies cannot be represented because they go beyond what the grammar faculty is able to establish.

There are two inter-related problems with the trend initiated by Postal. First, something like the anaphoric island constraint would rule out the deep anaphora cases in (1)a. This simply would get the data wrong and go counter speaker intuitions (which motivated the constraint in the first place). Second, the status of the anaphoric island constraint as a condition imposed by grammar is a bit mysterious. For example, Chomsky (1986) famously argued for the idea that grammar is a set of rules (comprising an I-language) induced from the stimuli in a specific natural language *modulo* a universal set of constraints (a *Universal Grammar*). Under this conception, the rules of an I-language basically amount to certain conditions on well formedness of strings (competence). Since donkey anaphora exploits relations across strings (see chapter 1 section 1.5), Postal's anaphoric island constraint cannot be stated at the level of grammar, according to this narrow view. That is, his constraint seems to capture conditions on dependencies between discourse 'referents', a level of representation much richer than the string-bound forms of dependency that a narrow conception of grammar is concerned with[2]. However, if we depart from such narrow view, as Postal's constraint seems to do, it is mysterious why the saliency of the paycheck in (1)a and the dog in (1)b cannot provide discourse antecedents.

In order to overcome these difficulties, Elbourne (2001, 2005), a recent exponent of the tradition initiated by Postal, refines the early conditions on donkey anaphora. In order to capture the deep anaphora data in (1)a in the right way, he postulates that

---

[2]The narrow view is not the only conception of grammar. Discourse Representational Theory (Kamp, 1981), Buring's Discourse Trees (Büring, 2001) and *Dynamic Syntax* (Kempson et al., 2001; Cann et al., 2005) assume that grammar can describe structure that is not string bound.

donkey cases involve NP ellipsis, that is, it involves the deletion of a constituent of a noun phrase at the level of phonological form (PF, i.e. the pronounced elements of a sentence) which nevertheless contributes to the sentence's logical form (LF, i.e. the structured representations that corresponds to truth-conditional content). For Elbourne, such an operation is defined as follows:

Elbourne's condition on NP ellipsis: NP deletion at PF requires (i) an overt NP as an antecedent (ii) that is not part of a word *or* (iii) a sufficiently salient and pragmatically accessible representation that serves this role.

Conditions (i) and (ii) are basically the same as Postal's and effectively explain the antecedent accessibility pattern in (2). Condition (iii) explains the deep anaphora cases where PAYCHECK in (1)a is recovered at the level of LF and the elided noun DOG in (1)b is recovered in the same way. With regards to such licensing *via* pragmatic saliency, Elbourne (2005, 45) observes that:

> NP deletion in the absence of a linguistic antecedent would rely on some extralinguistic reconstruction by the hearer of what must be meant by the speaker; this explains the fact that it seems limited to cases where there is some *immediate cue* in the physical environment, which is indicated by some physical gesture for the greatest felicity to result. *Any harder task, presumably, would produce the feeling of mental stretching that one has on hearing* [e.g. (b) in (2)] (italics and square brackets mine).

It seems that an immediate salient cue corresponds to a feature in the environment that makes a concept directly available for integration with encoded information, in the intuitive sense used earlier[3]. Based on these claims, condition (iii) of the constraint would rule out cases where the cue does not make the 'discourse entity' immediately salient, as in (2)b ('Every married man should bring her to the party', where 'her' is interpreted as a particular man's wife). That is, concepts that are only indirectly available in the context cannot license the ellipsis operation (e.g. the reconstruction of the elided noun by the audience) and hence cannot

---

[3]Elbourne mentions that retrieving the LF of the sentence uttered corresponds to 'a reconstruction' of speaker's meaning, but he does not specify by which mechanisms such reconstruction takes place. As discussed in chapter 2, Grice did not describe how contextually available information figures in what-is-said and allegedly precluded conversational maxims from shaping this level of content. It seems that Grice's position cannot cast light on Elbourne's appeal to pragmatic saliency. Adopting a pragmatic framework like RT could make sense of this process, but with it a whole lot of processes that shape truth-conditional content would be allowed, threatening the strictness of conditions (i) and (ii) of the NP ellipsis constraint. See further discussion in this section.

contribute to the propositions expressed. This carries some interesting explanatory consequences. For example, some cases of descriptive pronouns, presented in chapter 2 and reprinted below, would be ruled out. Consider:

(3)   *Unavailability of descriptive interpretations*

    a.   *Ann and Matt are in a wedding shop browsing through the items, Matt to Ann*: Every groom thinks his bride is worth dying for.

    b.   *In the same context as (a), Matt utters to Ann:* Every groom thinks *she/ ?his [pointing to a wedding dress] is worth dying for.

Above, the gesture towards the wedding dress does not count as a immediate cue for the emergence of a BRIDE concept, hence accounting for the infelicity or ungrammaticality of the use in (3)b. However, by the same token, the speaker's gesture towards the White House in (1)e would *not* count as an immediate cue for interpreting 'He has tough choices to make' as *the American President has tough choices to make*. Condition (iii) of Elbourne's constraint rules out all identification-based descriptive pronouns, thus providing a negative answer to the CMI question: indirectly available concepts cannot integrate with encoded information in order to express a proposition (count as an development of the utterance's logical form). This is undesirable, for these interpretations seem available for most language users and, in some cases even fall within the scope of logical operators (in (1)e, the interpretation of the pronoun covaries with the (quantificational) adverbial 'Every time there is a war'). Moreover, as mentioned in chapter 1 section 1.8, Elbourne's constraint on NP deletion would rule out cases of donkey anaphora where neither an overt antecedent that is not a word-part (conditions (i) and (ii)) nor an immediate cue for the intended interpretation (condition (iii)) are present in the context. However, uses in these circumstances can be felicitous, like the following and previously mentioned examples.

(4)   a.   When the baby **threw-up**, did you find any pencil eraser in **it** (Anderson, 1971)?

    b.   Every **iphone owner** uses **it** for browsing (Patel et al., 2009).

Here, the VP 'threw-up' allows the pronoun to be interpreted as *the vomit* in (4)a and the complex 'iphone owner' in (4)b provides 'it' with the interpretation *his iphone*. If such dependencies as well as identification-based pronouns are to be explained, they should be explained by a mechanism different from that of NP

ellipsis licensing [4], but it is not clear how this alternative mechanism would be able to allocate the relevant content at a level of meaning distinct from the proposition expressed, on the assumption that constraints such as Postal's or Elbourne's exist.

Before raising the dilemma central to this section, I would like to make some comments on the transition from Postal's constraint to the one proposed by Elbourne. As previously mentioned, the definition given by Elbourne is able to account for cases of 'deep anaphora': an advantage over Postal's proposal. However, if the status of Postal's constraint *vis-à-vis* the architecture of grammar (narrowly conceived) is unclear, the status of Elbourne's is even more so. If condition (iii), which incorporates issues concerning the saliency of mental representations, is part of a constraint on NP-ellipsis, then either the operation is not licensed by grammar (in the narrow sense) or the operation is grammatically licensed, but grammar here may draw on other representational resources (e.g. the ability to visually single out individuals as pertaining to a kind; i.e. conceptual system, more generally). The fact that the two positions seem virtually indistinguishable from each other presses an interesting point. Elbourne's disjunctive condition does not explain what immediate saliency amounts to. But, if we try capture this notion in a theoretically interesting way, we could reduce Elbourne's disjunctive constraint into a single one that covers both linguistic and extra-linguistic licensors for the relevant felicitous dependencies (in (1), say) and, at the same time, rules out the relevant infelicitous cases (in (2)b and (3)b, say). This could achieved by the following:

The Individuative-Representational constraint: A dependent use of a linguistic expression must draw its interpretation from an individuative representation available from the context.

The use of 'dependent' above unifies dependencies that are linguistic in nature, like the standard cases of donkey anaphora, and cases that rely on the visual identification of objects in the environment, like the uses in (1)[5]. The notion of an *individuative representation* might be characterised by three roles: (i) it emerges cross-modally (perception and decoding could give rise to the same representations), (ii) it singles out individuals or discourse 'referents' in thought, and (iii) it captures a

---

[4]As an alternative, Elbourne (2008) writes a different semantics for identification dependent descriptive pronouns.

[5]The notion of identification I am interested in potentially covers other cases, like singling out an individual in memory and in testimony or communication (see Evans, 1982, ch. 5). Since this chapter aims to answer the CMI question based on empirical evidence, the literature review is restricted to identification in vision only.

level of representation without which such singling out would not be possible.

In the linguistic case, these roles would be played by the concept that is encoded by the antecedent. For example, in 'Every man who has a wife should bring her to the party' (i.e. (2)a), the concept WIFE emerges via the decoding of 'wife' and singles out certain discourse entities that can be picked out by the pronoun 'she'. In 'Every married man should bring her to the party' (i.e. (2)b), however, the concept MARRIED, which emerges via the decoding of 'married', does not discriminate between husbands and wives and hence does to play an individuative role, that is, it does not single out discourse entities that can be picked up by the pronoun. As a result, utterances of this type are ruled out by the *Individuative-Representational* constraint.

In the perceptual case, the roles (i) to (iii) above would be played by concepts that emerge visually. For example, in 'Do most faculty members deposit it [*waving a paycheck*] in the Credit Union ?' (i.e. (1)a), the concept PAYCHECK emerges via the visual identification of the object waved by the speaker and singles out discourse referents in a way such that they can be picked out by the pronoun 'it'. In 'Every groom thinks *she/?his [*pointing to a wedding dress*] is worth dying for' (i.e. (3)b), the concept WEDDING DRESS emerges through the visual experience, but it does not partition the world into discourse entities, in this case, *brides*, that can be picked out by the pronoun. This explains the infelicity of the pronominal use.

In terms of predictions, it is very difficult to see how Elbourne's NP deletion proposal and the *Individuative-Representational* constraint, sketched above, differ. Moving from the former to the latter highlights an important shift, however. The notion of an immediately salient concept, present in the NP deletion account, has been formulated in terms of concepts that, among other things, *single out* individuals in thought and *capture a level of representation without which such singling out would not be possible*. Although it is easy to see how that can be met by concepts that are encoded by certain lexical items, it is more difficult to see how that would come about in the perceptual case. Why would my seeing of an object waved by someone, require the concept PAYCHECK, or any concept at all? Thus, in order to play the role required by the *Individuative-Representational* constraint, we would need to find the counterpart, in vision, of concepts that emerge through decoding. Perhaps the best candidate for establishing the relation between the visual identification of a salient individual, on the one hand, and the availability of certain conceptual representations, on the other, is the notion of a *sortal*.

In the philosophical and psychological literature (Strawson 1959; Quine 1960; Wiggins 1967, 1980, 2001; Macnamara 1972, 1982, 1986; Spelke 1990, to name

just a few), it has been proposed that perceiving particulars, that is, individuals that move as a whole, have a bound contour and so on, *requires* a specific kind of conceptual representation, otherwise such entities would not be presented to us as the individuals they are. This representation is known as a *sortal* concept. Intuitively, sortals are the kinds of concepts that answer 'What is it?', where 'it' designates a given object, and complete the question 'How many ... are there?'. For example, in a situation where one sees a dog and is asked 'What is it?, the answer would be mediated by a DOG concept. If the question 'How many are there?' is further asked, the same concept (e.g. DOG) would determine the answer, namely, one (and not four, or five). In short, sortals seem to outline certain conditions on identity and individuation of particulars (I will be more explicit about these notions shortly), and hence are good candidates for playing the role required by the *Individuative-Representational* constraint. In short, sortals would be the most accessible conceptual representations in a given perceptual experience (i.e. seeing a ball moving), because they are required in order to characterise the experience as such (i.e. as the experience of a ball moving). More specifically, *Sortalism* is captured by two claims:

Sortal Individuation: Sortal concepts are necessary for singling out individuals in thought.

Sortal Identity: Sortal concepts are necessary for capturing the identity conditions of objects, that is, the conditions under which they remain the same through time and space and the set of transformations they can undergo while remaining the same objects [6].

It is important to notice that sortals are necessary but not sufficient for characterising a visual experience of solid, bound, three-dimensional material objects as such. A causal relation between the object and the visual system is also needed. I will introduce, motivate, and properly discuss the notion of a sortal in the next section. Now, if such representations do exist, then it seems that the *Individuative-Representational* constraint and the notion of directly available concepts, used by Elbourne's NP ellipsis constraint, can be sensibly formulated. The appeal to any

---

[6]This definition may lump together two notions of identity: one related to spatio-temporal continuity and other related to the changes an object can undergo. These two notions come apart in many cases. For example, when a person dies, she is spatio-temporally continuous with her body, but the individual that once existed ceases to be at the moment of death. In my assessment of the sortal theses these two notions of identity will be distinguished, even though they are represented here in a unified manner.

of such constraints explains some of the accessibility patterns present in donkey anaphora as well as cases of deep anaphora, but it would inevitably cast identification-based descriptive pronouns aside. As these exploit concepts that are indirectly available (like metonymy), the relevant representations would not be able to integrate with the linguistic meaning of pronouns (the CMI question would be answered negatively). How then would descriptive uses of pronouns be explained? It is not clear, but one possibility is this. Advocates of such constraints could appeal to other levels of information, such as implicatures (although chapter 2 discusses the limits of such appeal), or presupposition (i.e. imagine a scenario where the immediately salient concept WHITE HOUSE is presupposed to be a shorthand for talking about the presidents whose office is based in it). In short, alternative means to explain the data surrounding descriptive pronouns would have to be sought.

On the other hand, if sortals do not exist, then the Individuative-Representational constraint and the notion of 'directly' available concepts, used by Elbourne's NP ellipsis constraint, *cannot* be sensibly formulated. As a result, there would be no significant distinction between descriptive uses of pronouns and deep anaphora. Concepts that are 'indirectly' available can integrate with the linguistic meaning of pronouns in the same way as 'directly' available concepts would (the CMI question would be answered positively). As a result, the accessibility pattern present in donkey anaphora, some descriptive cases (and even metonymy) would have to be explained by mechanisms other than the constraints reviewed above. This clears the way for the development of a *substitution-based account* for some descriptive uses in the sense of chapter 2 (section 2.3.3). This will be fully developed in chapter 5.

In this chapter, we will assess the existence of sortal concepts as individuative representations that can make sense of the direct availability of certain interpretations, from both a conceptual and an empirical standpoint. The latter aspect has been motivated by comments made towards the end of chapter 2. If the pragmatic principles that guide interpretation are the same, it is mysterious why certain 'indirectly' available representations can be selected as constituents for the proposition expressed in some contexts but not in others. The next section, 3.2, provides motivations and evidence in support of Sortalism (captured under the Individuation and Identification theses mentioned above). Section 3.3 provides evidence against it. In section 3.4, I suggest that the empirical evidence reviewed raises a dilemma and I aim to resolve it by assuming an incremental notion of processing that is based on the cognitive principle outlined in chapter 2. In section 3.5, I suggest how the notion of incremental processing developed in the previous section helps us to answer the CMI question and understand some of the issues outlined here. More specifi-

117

cally, I argue against Sortalism and for the idea that the so called indirectly available concepts can be integrated with linguistic meaning. The accessibility patterns mentioned in this introduction and in previous parts of the thesis can be partially explained by the incremental notion of processing put forward in this chapter.

3.2 Sortalism and direct availability

3.2.1 Sortalism introduced

Previously, I mentioned that sortals are motivated to capture conditions on individuation and identity of particulars. The first type of condition is a quantitative or numeric notion (*Sortal Individuation*). The second is a qualitative one (*Sortal Identity*). In the literature, arguments for these two roles have been provided by different sources. In what follows, I will present some of them.

The idea that sortals are somehow necessary (but not sufficient: a causal relation between object and thinker is another necessary component) to individuate an object (or to capture its identity in the numeric or quantitative sense) is nicely illustrated by Frege's *counting argument* (Frege, 1950) and Quine's *argument from divided reference*.

The counting argument supposes a scenario where someone is looking at a deck of cards. Now, the same perceptual experience could be taken as eliciting either a representation about a single object – the deck – or a mental representation about fifty-two numerically distinct objects – the cards. So, how is it that a single perceptual episode can elicit two different thoughts? Frege's answer relies on the idea that concepts are necessary in order to count (i.e. to numerically individuate) objects. In one case, object individuation would proceed via the DECK concept; in the other, it would proceed via the CARD concept.

The argument from divided reference (Quine, 1960) seems to be a radical version of the counting argument. Consider a person looking at the waters of a river and thinking that the water looks refreshing. It seems natural to ask the following question: is there a specific entity this person's thought is about? It is certainly about something, but there does not seem to be a particular entity that is 'singled out' in thinking. In the words of Quine (1960, 61), sortals[7] are necessary for 'dividing

---

[7]Instead of postulating mental representations like concepts, whose existence is obscure, Quine actually held that the role played by sortals would be captured by linguistic entities: predicates. The idea here is that individuation reflects ontological commitments held by a community, as, for example, the social commitments involved in fixing the reference of the term 'couple'. For the sake of simplicity, I will unify the discussion of this chapter under *concepts*.

reference', that is, carving the environment into sets of discrete entities, something which mass concepts do not do. This idea exploits, among other things, situations in which objects may have dimensions that go beyond the visual field. Even though one does not see the object's boundaries when visualising mountains, lakes, rivers, etc., one can still think about one particular object (e.g. *this river*) as opposed to other objects. According to Quine, the possibility of having such thoughts (or referring to such entities) requires sortals. The underlying idea can also be read off the example used in his thesis of the inscrutability of reference (Quine, 1960, ch. 2). Consider a linguist in fieldwork trying to master an entirely alien language and a native speaker who says 'Gavagai!' while pointing to a white rabbit that just emerged in the scene. Apparently, there are various hypotheses that are consistent with the evidence available to the linguist: 'gavagai' could be taken as synonymous with 'fur', 'whiteness', or 'undetached rabbit-part', and so forth. Similarly, the perceptual experience of seeing the rabbit can be construed as being about a certain rabbit, but it also could be about the rabbit's colour, shape, parts and so forth. In order to rule out relevant alternatives, Quine argues that sortal concepts must be applied to the stimuli. What makes an experience one about rabbits is the fact that it has been organised under the concept RABBIT or ANIMAL[8].

In addition to their role in capturing the conditions of individuation on objects, sortals arguably have the role of capturing (some of) an object's identity conditions (the qualitative notion mentioned previously). In order to support this claim (it can also be found in some of Quine's remarks), Strawson (1959) argues that humans represent objects by some form of conceptual structure (scheme) that captures their persistence conditions (in a minimal sense: spatiotemporal continuity). If objects were ever changing in nature: multiplying, disintegrating and re-assembling themselves as seconds go by, we could not make sense of an agent's thinking about a specific entity in the environment. We simply would not be able to identify objects and re-identify them (i.e. identify something as the previously identified entity) in space-time. For example, I can identify the chair that I am now sitting on as the same chair I sat on yesterday and also identify a rolling ball as the same ball in adjacent spatial positions at successive times. Note that the sensory stimuli could be different: the chair might have acquired a coffee stain in between and the ball could reflect light differently as a cloud passes by.

The ability to (re-)identify requires that material objects are perceived as endur-

---

[8]Sortal concepts were of great interest to psychologists working on language acquisition, for they would offer a restriction on the entities that a given linguistic category (nouns) may refer to. For ground-breaking work, see Macnamara (1972, 1982, 1986).

ing, at least briefly, in space-time. Since the sensory stimulations are ever changing in nature (the retinal image of the object can change dramatically), information that hits the retina would not be able to capture the conditions under which the object remains the same (e.g. if a chair liquefies, there would also be a dramatic change in retinal image, but in this case, we would say that the chair that once existed does not exist any more). If the identity of an object cannot be traced back to the original stimulus, it must be captured by a basic conceptual structure, under which three-dimensional, bound, persistent objects fall. This roughly corresponds to the sortal MATERIAL OBJECT or BODY (see Strawson 1959, 31-40, 168)[9]. I will call this, *the argument from re-identification*. Schematically, it moves from the premises (i) that the ability to identify and re-identify objects in space-time underpins object representations, and (ii) that the ability to identify and re-identify objects requires the sortal concept MATERIAL OBJECT, to the conclusion (iii) that object representation involves the sortal concept MATERIAL OBJECT. Note that this concept is the most basic one in our conceptual scheme and more specific conditions of identity can be captured by more specific concepts, such as ARTEFACT or PERSON.

Finally, although I have reviewed here some arguments that motivate the idea that the necessary role played by sortals in characterising visual experience can be recruited as an explanation of the direct availability of some conceptual representations, it is rather unclear *which* concepts are the necessary ones. For example, for Quine, representing an object's individuation and identity conditions would be achieved by very specific concepts, such as DOG or BALL. Strawson, on the other hand, believed that such roles can be played by the less specific concept MATERIAL OBJECT (plus the concept PERSON to make the distinction between representations about the self and those about the objective world). Therefore, Sortalism, the position according to which sortal concepts are necessary for capturing the identity and individuation of objects, can assume different degrees of strength. A strong version of Sortalism requires more specific concepts to specify the individuation and identity conditions (Quine, 1960; Wiggins, 1967, 1980, 2001) of particular objects, a weak version states that something like MATERIAL OBJECT would do (Strawson 1959; Spelke 1990, among many others).

---

[9]In effect, this argument fits in with Strawson's search for the conditions of objectivity of experience, which for him require entities that are distinct from the self, hence the importance of MATERIAL OBJECT. Much of vision science (see Palmer (1999) for a textbook) had similar goals: how can a rich representation about the distal stimulus be constructed from an informationally poorer 2D proximal stimulation? The difference is that Strawson postulates that only concepts could make sense of human's re-identification of objects, whereas for many constructivist psychologist this could be explained by a system of constrained information processing units. Burge (2010) takes concepts to be not necessary for re-identification nor for rendering experience *objective*.

This has particular consequences for this chapter. A weak version of sortalism would not be able to account for the data reviewed in the introduction, as the interpretations meant the speaker contain more specific concepts, such as, PAYCHECK, DOG, etc. Thus, it seems that only strong versions of Sortalism, i.e. more specific concepts are necessary for individuating objects and capturing their persistence conditions, would do interesting theoretical work. However, since I have opted to assess sortalism from an empirical standpoint and much of the evidence discussed in this chapter concerns the role played by MATERIAL OBJECT (specially in the first year of human life), I will confine the initial discussion to this much more basic level of representation before moving towards evidence for the idea that more specific conceptual representations may be used for individuating and identifying objects in the environment. I urge the reader to bear with me through the next two sections.

3.2.2 Evidence for sortalism

A series of experiments conducted in the late '80s and early '90s brought evidence that was taken as supporting the sortalist position. These experiments used the method of violation of expectancy in infants from 2.5 months to 36 months old. The idea behind this methodology is the following. Children were presented with events which had two possible outcomes: an expected one and an unexpected one (according to adults' understanding of the world). If children look more at the unexpected outcome, it implies that their expectations about the way the world works are adult-like. The theoretical task then would be to speculate about what kind of cognitive mechanism generates such expectations (for ground-breaking work, see Baillargeon 1987; Spelke 1990).

Spelke et al. (1995) conducted the following study. First, two spatially separated screens were introduced. An object (rubber duck) was removed from behind screen one and shown to the infant. After that, the object was put back in its original place. The same process happened with screen two: another occluded object (rubber duck) was removed and shown to the infant and then replaced behind the screen. Finally, both screens were removed revealing either one of two possible outcomes: the expected outcome displaying two rubber ducks or the unexpected outcome displaying one rubber duck (see the dynamics in the left of figure 3.1). The authors found that the looking times of 4.5 month-olds was significantly longer when the outcome was unexpected, suggesting that they used the spatial gaps between screens to determine how many rubber ducks were occluded. In order to control for the possibility of children's sensitivity to the matter that constitute the objects, Huntley-Fenner

et al. (2002) devised a similar experiment with two conditions: one in which sand was poured behind two (separated) screens and another in which two objects made of sand were placed behind the screens. In the sand condition, infants could be presented with either one (unexpected) or two (expected) sand piles. In the sand object condition, infants could be presented with either one (unexpected) sand object or two (expected). Infants were at random in the sand condition, but looked longer when the outcome was one sand object, suggesting that they expected that two sand objects (one per screen) were occluded and thus reduplicating the results of Spelke et al. (1995). These findings were used to motivate what Spelke calls 'principle of cohesion', which states that an object is a bounded chunk of matter that preserves its connectedness and boundaries as it moves through space [10].
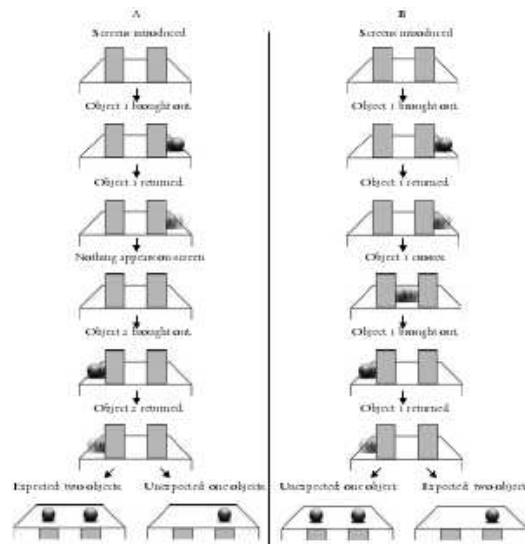
Figure 3.1: Diagram of the split-screen spatio-temporal continuity paradigm in Spelke et al. (1995). Permission to reproduce this image has been granted by Prof. E. Spelke

Going beyond perceptual processes explainable by a principle of cohesion, Baillargeon et al. (1985) investigated children's expectations regarding the behaviour of objects in motion. Four-month-olds saw an object being placed behind a drawbridge-like screen. As the screen rotates backwards, two outcomes are possible. The screen could either stop roughly at the position where the object is (expected outcome), or it could come all the way down to the ground, as if the object were ethereal (unexpected outcome, see figure 3.2). Like before, four-month-olds looked more at

---

[10]In addition to spatial gaps, Needham and Baillargeon (1998) show that 8 month-olds start to use Gestalt principles (e.g. similarity, good continuation, figure-ground segregation) to individuate objects. For example, although a horseman and his horse seem to move as a whole, differences in their shapes and colours may be used to distinguish one from the other.

the unexpected outcome. The finding supports two principles concerning infants' conception of physical objects. First, infants do not expect objects to disappear at one point and reappear at another (principle of continuity). Second, infants expect objects to constitute physical barriers to the movement of other objects (principle of solidity).
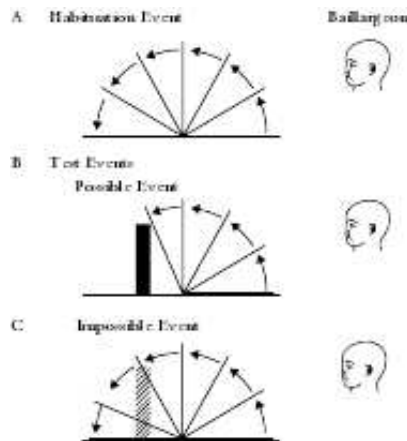
Figure 3.2: Diagram of the rotating screen paradigm in Baillargeon et al. (1985). Permission to reproduce this image has been granted by Prof. E. Spelke

In order to test whether children's expectations were based only on visual stimuli, Spelke (1990, 44-5) blocked the infants' visual field with a screen while letting them play with two rings that moved together, as if connected by a bar. After experiencing the rings as if they were parts of the same object, the screen would be removed and two outcomes were possible: either the rings were connected by a bar (expected) or they were disconnected (unexpected). Spelke found that children were surprised by the unexpected result, suggesting that infants' expectations applies to stimulation in multiple modalities (visual and haptic, at least). The conjunction of these principles (cohesion, continuity, solidity) is taken by psychologists to determine the concept SPELKE-OBJECT. Since it is essentially the same as Strawson's MATERIAL OBJECT, I will follow Strawson's terminology but take it to be also motivated by the empirical evidence discussed here. It is worth noting that, in addition to these principles, Spelke and collaborators postulate a principle of contact, necessary to capture children's expectations that only animate objects can move themselves (and which determines an ANIMAL *or* ANIMATE OBJECT concept).

The Spelke-concepts and their underlying principles[11] account for the evidence

---

[11]There seems to be a difference between the way Spelke and Kinzler (2007) and Carey (2009) see the status of conceptual representations that explain children's behaviour. Spelke takes the relevant

as follows. As the objects are shown to the infant, the MATERIAL OBJECT concept provides a criterion for individuation: each entity that falls under the concept would be numerically distinct from the others. Suppose there are two material objects at locations l_1 and l_2 in the subject's visual field. Given that the stimuli are such that the two corresponding entities fall under the MATERIAL OBJECT concept *and* that these instances of MATERIAL OBJECT occur at different spatial locations, the system of representation would give rise to two numerically distinct object representations. The system also captures the object's identity conditions. According to the set up in Spelke et al. (1995), the fact that first rubber duck falls under MATERIAL OBJECT enables the child to have certain expectations about the conditions under which these objects persist in time-space. As the duck is put behind the screen, the infant is required to maintain that specific object representation in working memory. As there are no other events happening, the MATERIAL OBJECT enables the child to assume that the object remains behind the screen, for material objects preserve their spatial properties if there is no interaction with other material objects (e.g. crushing by another heavier object). The same process would apply for the second duck. Metaphorically, one could describe this operation as performing a primitive form of adding 1 to 1, which creates an expectation of 2 objects behind the screen. If only one duck is seen after the screens are removed, then at least one of the Spelke principles, in this case, continuity or cohesion, was violated, which explains the longer looking times. Similar explanations apply to the other experiments, (solidity and cohesion for Baillargeon et al. 1985).

According to this type of explanation, children's expectations would emerge from a system of representation concerning the behaviour of medium-sized objects in the world (proto-physics). This system - alongside systems for representing purpose-oriented self-propelled objects and causality (via the principle of contact), basic numeric operations, and spatial geometry - constitute core systems of representations that are taken to be innate in humans and some other animals (see Kinzler and Spelke 2007; Spelke and Kinzler 2007; Baillargeon 2008 on how these systems are inter-related).

Note that the concept MATERIAL OBJECT does not capture the intended readings of many of the identification dependent uses reviewed previously. The more specific

---

concepts to be intrinsic to a system of core knowledge: the child looks more at the unexpected outcome because she knows it should not have happened. Carey tones it down. The relevant concepts are the outputs of systems of core cognition and do not necessarily correspond to propositional knowledge stored in the mind of infants. Children look more at the unexpected outcomes, because they go against what they are cognitively geared towards. Unfortunately, a proper assessment of these differences transcends the purposes of this thesis.

conceptual representations that are part of such readings have not been mastered by the children who participated in the studies reviewed here. Most of them were under the first year of life. Still, the argument for the necessity of MATERIAL OBJECT can be implemented for more specific concepts (representations that do bear on the availability of the readings mentioned at the beginning of the chapter). This will be done in the next section.

3.2.3 Object representation and conceptuality

It is intuitive to say that humans represent objects as bound chunks of matter that move as whole through time-space, but one may wonder whether the robust findings by Spelke, Baillargeon, Carey and colleagues really necessitate the postulation of conceptual representations it order to explain the empirical findings supporting this intuition. Carey and Xu (2001) propose three arguments for the idea that the systems for representing medium-sized object representations are *conceptual* in nature.

The first of them basically consists in evidence for Strawson's argument from re-identification. The fact that infants are surprised when they see one object (in the unexpected outcome condition) instead of two suggests that they were expecting the missing object to be behind the screen. Carey and Xu (2001) argue that such expectations require children to distinguish an object's *occlusion* from the *cessation of its existence*. If material objects were *non-persisting*, infants could assume that the missing object somehow got destroyed behind the screen and find the outcome containing a single object to be normal. Similarly to Strawson, they argue that only sortal concepts capture the conditions under which a given object persists (and hence can be re-identified by the subject). Therefore, concepts (e.g. MATERIAL OBJECT) become necessary to explain the behavioural data.

The second argument emerges from the *integrative* nature of the representations governing infants' behaviour, a terminology from Carey and Xu (2001, 208). The basic idea is that information that governs object representation seems to be available for other cognitive tasks, suggesting that the underlying representations have a common format. Under this assumption, it would not be far off to claim that concepts, the representations manipulated by central systems, play such a unifying role (see Fodor, 1983). In order to back up this claim, the argument invokes the following empirical evidence. Spelke (1990, 44-5) found that infants' expectations hold across modalities (visual and haptic, see also Gordon and Irwin 1996, 2000, for a similar finding regarding visual and 'linguistic' modalities). A similar point could be made about the infants' attitude towards content that emerges perceptually. The

output representations guide human volitional action, thus exploiting interactions between beliefs and desires. Given that these two mental states target the same content, one could assume that the relevant representations are structured by concepts (i.e representations in a common format). Also, Carey (2009) shows at length that the expectations of children concerning the behaviour of material objects and self moving agents are quite complex and sensitive to multiple cues. The basic idea is that the behaviour governed by the infant's representation of objects in the environment enters complex inferential relations. This has been shown previously: on the basis of certain stimuli children *infer* that two objects, rather than one, lie behind the screens. As more complex forms of inductive inferences will be discussed in the third argument, below, I will postpone a more substantial discussion about this point until then.

The third argument is based on evidence relating a harder object-permanence task to the availability of more specific sortal concepts (e.g. BALL, DOG, etc.). Xu and Carey (1996) tested the ability of 10- and 12-month-olds to use information concerning an object's identity-conditions as a criterion of individuation. They introduced a slight modification to the designs discussed previously. Subjects were presented with a single screen, instead of two separate ones. Initially, a ball was removed from behind the screen, shown to the participant, and placed back where it came from. Then, a duck was removed from behind the screen, shown to the child and then placed in its original position. As usual, the screen was removed and participants would be facing either the expected outcome of a display containing a ball and a duck, or the unexpected display containing just a ball (or just a duck, see figure 3.3). Now, there is one important difference between this experiment and previous ones. Since only one screen was introduced (no spatial gap between the locations the objects are placed at), the child not only has to represent the objects under occlusion, but also consider possible interactions between them while occluded. For example, if a chest and a ball were occluded and only the chest emerged as outcome, the child may infer that the ball is inside the chest, and hence infer that only one object will be displayed. The task requires no simple effort.

Xu and Carey (1996) found that 12-month-olds looked more at the unexpected display. 10-month-olds, on the other hand, looked at both displays for the same amount of time. Given the evidence that the MATERIAL OBJECT concept is available since 4.5 months of age, it seems that this concept is not enough for children to make the 1 [MATERIAL OBJECT] + 1 [MATERIAL OBJECT] = 2 inference. 12-month-olds' longer exposure to the occluded objects suggest that their success in the task is related to the availability of more specific sortal concepts, such as, DUCK or BALL.

The task is schematically represented below:
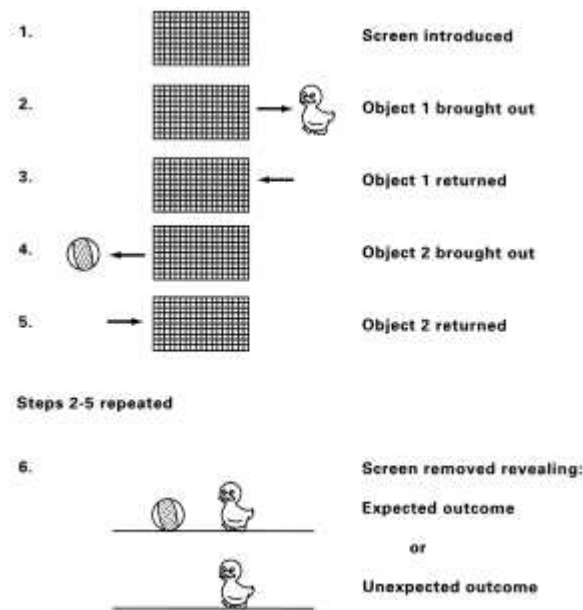
Figure 3.3: Diagram of the single screen paradigm in Xu and Carey (1996). Permission to reproduce this image has been granted by Elsevier.

As a follow-up to Xu and Carey (1996), Xu (1999, 2002, 2005) and Dewar and Xu (2007) tested 9-month-olds using the same design with the exception that, as objects were shown to the baby, linguistic cues were also given. At the time that the baby looked at the duck, someone would say 'Look [baby's name], a duck'. The same thing was performed for the ball. In this set up, 9-month-olds looked for a longer period of time at the unexpected outcome, similarly to the 12-month-olds in the Xu and Carey (1996) study. Moreover, Xu introduced a condition where both the displayed objects were followed by 'Look [baby's name] a toy'. The looking times were not affected by these linguistic stimuli. In order to control for other variables influencing the correlation between the saliency of distinct kinds, on the one hand, and the different objects that instantiate such kinds, on the other, Xu and colleagues included conditions where the objects were introduced while different non-linguistic sounds ('blink' vs. 'blonk') and even communicative sounds (approval 'yay' vs rejection 'yuck') were played to the child. None, with the exception of words similar to nouns in children's first language, improved performance. The only other factor that helped the individuation task was when the two objects moved, during the habituation phase, in a way such that it became clear to the child

127

that they belonged to different kinds. For example, Surian et al. (2004) introduced a condition where a toy train would move in and out from behind the screen making train-like mechanical noises and whistling, whereas the other object, an agent-like creature, would walk in and out from behind the screen accompanied by footstep sounds. In this condition, 10-month-olds behaved like 12-month-olds.

The explanation of this finding by Xu, Surian and colleagues runs along the following lines. There are important boot-strapping relations between concepts and the word categories (e.g. common nouns) they might be related to (Bonatti et al., 2002). It is likely that children around 8-12 months of age are learning the mappings between certain words, predominantly common nouns, and concepts. In the case of 9-month-olds, the mapping is not well established. The experimenter's use of different common nouns could make more salient to the child that there are two kinds classifying the objects behind the screen, and hence they could infer that 2 objects (falling under the respective kinds) are occluded. In the Surian et al. (2004) study, a similar effect was achieved by using stimuli that made clear that 10-month-olds were dealing with two different kinds of objects. The extra experience of 12-month-olds allows the relevant kinds (and the concepts that capture them) more easily accessible, hence the adult-like performance[12].

In conclusion, the studies reviewed in this section suggest that (i) there is a robust set of data concerning very young infants which needs to be accounted for and (ii) the interpretations given here display support for certain incrementality in the system of object representation. Children seem to be equipped with a MATERIAL OBJECT concept, whose availability matures in the first 6 months of life. With the acquisition of more concepts, more specific sortals can perhaps be used for object individuation and the representation of an object's identity conditions. Such concepts, such as PAYCHECK, DOG, would then be the ones that become directly available in visual experiences of (mature) language users and can be recruited by the constraints reviewed in this chapter's introduction in an explanation of some linguistic data.

---

[12]Carey (2009, ch.4) discusses whether these findings suggest that when children acquire more specific concepts, like TOY TRUCK, DOG, TABLE, they acquire new and improved mechanisms for individuating objects and representing their identity conditions. This possibility would involve a certain reconceptualisation of the world by the child (that is, a representational system's induction of another, more powerful, representational system, see the discussion of Quinian bootstrapping in Carey 2009). Carey assumes *contra* Fodor that this operation exists and underlies certain transitions in human ontogeny. If this is the case, then it seems to me that concepts lose much of their explanatory power as a common representational format, since there are interesting cognitive relations between systems whose representations may be stated in different formats. Since the data relevant for the discussion in this chapter could be explained simple by infants' acquisition of concepts along Fodorian lines, I will abstract away from the issue.

3.2.4 Interim conclusions: sortal concepts and linguistic integration

In the first section of this chapter we raised a question about cross-modal integration as well as the need for empirically supported answers to it. This (CMI) question inquires into whether indirectly available concepts can be constituents of the proposition expressed by an utterance. Among other things, this question was motivated by certain *accessibility puzzles* raised by the availability of identification-based descriptive uses of pronoun, on the one hand, and the necessity to explain the impossibility of certain cases of donkey anaphora (i.e. the asymmetry between the felicitous 'every man who has a wife should bring her to the party' and the infelicitous 'every married man should bring her to the party'). In order to explain such conditions, many constraints on the availability of such antecedents have been proposed. We looked earlier at Postal's anaphoric island constraint, which states that antecedents must be overt NPs that are not word-parts. However, this constraint rules out cases of deep anaphora (e.g. 'Do most faculty members deposit it [*waving a paycheck*] in the credit union?'), where perceptual information provides the right donkey interpretation. We then moved to a constraint that accounts for such cases, namely, Elbourne's. It assumes that antecedents for donkey dependencies are available either via Postal's condition *or* by an immediate cue in the environment that provides the relevant concept. Given Elbourne's appeal to extra-linguistic representations, his proposal does not differ much from the Individuative-Representational constraint, suggested immediately afterwards, which proposes that a directly available (individuative) concept is required by any form of dependency (linguistic or perceptual).

This section provided an empirically grounded understanding of the terms 'immediate cue' and 'directly available' concept, used in such constraints. An immediate cue comprises stimuli that require the presence of certain (sortal) concepts, otherwise the experience would not be the experience that it is. Similarly, directly available concepts are the ones necessary in order to characterise the nature of a particular mental episode. According to the studies reviewed, the sortal concepts MATERIAL OBJECT and PERSON (via the principle of contact) are available to humans from the age of 4 months and enables them to individuate certain objects in the environment and also capture objects' persistence conditions. By the completion of the first year of life, more specific sortal concepts, like DOG, BALL or TOY become available to the agent and seem to be involved in these cognitive tasks. In short, the notion of a sortal concept appears to be empirically well-grounded and

could be recruited in order to make sense of Elbourne's constraint on NP deletion or the Individuative-Representational constraint. A negative answer the CMI question would be motivated and, thus, identification-based descriptive pronouns cannot contribute to the proposition expressed and should be explained by alternative means. This may be a counter intuitive result, but is a consequence that follows from more empirically supported claims.

The next section presents evidence against the the idea that sortal concepts play a role in the individuation and the representation of an object's identity conditions, challenging the fundamental status of such representations in our cognitive architecture. If this is the case, the grounding for understanding the level of representation appealed to by constraints such as Elbourne's of the Individuative-Representational constraint is gone. This will pave the way to a rejection of these conditions on empirical grounds: 'indirectly' available concepts emerge for the agent in the same way as 'directly' available ones do.

3.3 Anti-Sortalism: representing bare individuals

3.3.1 Bare object representation introduced

In this section, I review some experimental data against the idea that the individuation and the identity (in a minimal sense: spatiotemporal continuity) of objects requires a sortal concept. The review is divided into two subsections. One deals with object individuation: the quantitative notion exploited by the counting argument (Frege) and the argument from divided reference (Quine). The crucial data in this section stems from the difference between estimating the quantity of a small number of objects (known as subitizing) and larger quantities. The second body of work reviewed deals with the processes involved in tracking simple 2-dimensional objects through visual displays (known as multiple object tracking, or MOT).

3.3.2 The subitizing data

When subjects are presented with a small number of items in a visual display and asked how many objects there are, they give fast, precise and confident judgements. Alternatively, when the number of items is not so small, the speed, accuracy and confidence of the judgements decrease to a much greater degree than is proportional to the increase in items. Given the sharp contrast in performance, psychologists coined the term 'subitizing' to describe the former type of behaviour, re-

serving 'enumerating' or 'counting' as descriptions of the latter. This distinction is supported by the following findings.

If the individuation of the objects displayed is achieved by attending to each one in turn, the difference in confidence, accuracy and speed in judgements would be correlated linearly with the number of displayed items. However, Trick and Pylyshyn (1994a,b) falsify this hypothesis. If the items exceeds a certain number (4 or 5), accuracy and confidence decrease dramatically. Additionally, response times receive the huge increase of 250ms – 350ms per item in the display after the fourth, when compared to the 10ms-40ms per item in the 'subitizing' range (up to 4 items). This finding suggests that individuating up to 4 or 5 items is achieved by a mechanism different from the one involved in enumerating or estimating a larger number of objects.

In order to gain insight into the kind of mechanism that might be at play in the fast, reliable and accurate individuation of a small number of items, Trick and Pylyshyn (1994a,b) attempted to delineate the conditions under which subitizing is possible and the conditions under which it is not. The experimental paradigm they used had three conditions: (i) same size condition, (ii) different size condition, and (iii) concentric condition. In the same size condition, subjects were presented with a display containing rectangles delineated by four edges in a bound contour. All of them had the same size, which could be small, medium or large, and they were located in different locations in the display. In the different size condition, the setting was almost the same, the only difference being that rectangles could be of different sizes. In the concentric condition, all the rectangles had different sizes but they had the same centre. In each case, subjects had to simply say how many items were displayed.

Note that visual features - in this case shape, but also colour hue, saturation, luminance, position - need not be represented by *concepts*. Consider the human ability to distinguish different hues of the same colour, for example. We can distinguish red_{127} from red_{231}, but that might not surface at the conceptual level (e.g. as RED_{127} and RED_{231} concepts), due to a series of factors: (i) we may share the ability to distinguish between similar hues with other, evolutionary distant, creatures, (ii) these hue representations do not figure in thought productivity in the same way as RED and BLUE do, and (iii) the ability to distinguish very fine-grained hues may not be sufficient to 'lock' a cognitive system onto the right property, in the Fodorian sense of concept acquisition, among other factors.

In this study, Trick and Pylyshyn found that subitizing is possible in (i) the same size condition and in (ii) the different size condition, but not in (iii) the concentric

condition. Moreover, items that are too near one another or that are individuated by an instruction like 'Identify items on the same curve!' cannot be subitized (see Pylyshyn 2007, sect. 1.4.3). This seems puzzling, for the items presented in the concentric condition or in the curve condition were within the subitizing range. Trick and Pylyshyn have suggested that the visual system individuates objects by using coarse-grained information about their location in the display. This would engage a mechanism that determines which regions of the display are richer sources of information in an entirely bottom-up, data driven way, that is, before the stage at which we can describe the subject as having any kind of experience at all (see Olshausen and Koch 1995; Itti and Koch 2000 for arguments and a saliency based model of pre-attentive allocation of resources). Roughly, if the items are located in the same informationally salient region, then a mechanism of attention would have to focus on the concentrically positioned bound contours in order to determine which object is which; in other words, shifts of attention would have to spread serially. If these remarks are on the right track, subitizing engages mechanisms that distribute attention in parallel (see Pylyshyn 2007, 28, and Dehaene 1997).

Trick and Pylyshyn's suggestion seems to be further supported by neuroanatomical evidence from patients with Balint's syndrome. This condition is characterized by a general deficit in focal attention, generated by lesions in the post parietal cortex. Symptoms include the inability to see all the objects in the visual field simultaneously, difficulty in coordinating hand and eye movements, and the inability to shift attention towards another object. Dehaene (1997) reports that a patient with this condition failed to enumerate objects outside the subitizing range (more than 4 items) either by ignoring certain objects in the visual display or by counting the same objects multiple times. Nevertheless, the condition did not impair the ability to subitize.

In order to account for the data, Trick and Pylyshyn propose that subitizing is done by a mechanism of parallel object individuation that assigns visual indices to objects in the visual display. Pylyshyn calls such indices 'FINgers of INSTanciation', or 'FINSTS' for short, to remind us that they function as pointers. There are two fundamental characteristics of FINST assignment. First, they are assigned essentially in a causal, data-driven way, that is, information-flow is strictly bottom-up, coming directly from the objects perceived. Thus, saying that the visual system assigns indices could be a little bit misleading. Rather, the external objects themselves 'grab' certain indices given the way human cognitive architecture has evolved. According to Trick and Pylyshyn (1994a,b), there are two conditions that must be met for this 'grabbing' to occur: objects must have bound contours and they must be

located relatively away from each other. There is an important way in which such conditions differ from sortal concepts: they are stringent constraints on the information that serves as input to visual processing - if they are not met, visual indices are not assigned, whereas if an object does not fall under a concept it may fall under another concept. Pylyshyn claims that the representation conditions of the index assignment mechanism works as a *natural constraint*, in the sense of Marr (1982): if requirements imposed by the architecture of the cognitive system (e.g. module, network) are not met, the system cannot deliver its output. This *constraint* can be depicted by an *if - else* statement in computer languages:

(5)   *Paraphrase of the FINST assignment mechanism*

IF   *bound, move-as-a-whole, less or equal 4*
THEN *represent under the first index in the sequence: i,j,k,l,m*
ELSE   *abort*

Note that natural constraints are *architectural*, that is, the conditions in the IF statement belong to the hardware of the system of representation, so they would be *implicit* representations in the sense of Carey (2009). The condition in the ELSE statement illustrates an important difference between conceptual representations and the outputs of a natural constraint: if the latter misrepresents, no visual index emerges, whereas if an entity does not fall under a ANIMAL concept, it could fall under many other conceptual representations, such as PLANT, LIQUID, STUFF, and so on. I will now turn to experiments that suggest that this (alternative) mechanism of individuation can also capture the identity conditions of objects (in a very minimal sense).

3.3.3   The MOT data

In the subitizing experiments just reviewed, the task was to represent static objects, but how do humans deal with many moving objects in the display? Do they track these objects by their visual features, such as shape, colour, and size, by updates in their locations, or by something else? In order to answer these questions, Zenon Pylyshyn devised an experimental technique known as *multiple object tracking*. The idea is simple: subjects are presented with a display containing moving objects, some factors are manipulated, and, as a result, performance may improve or not. Sensitivity to the manipulated factors should shed light into the structure of the system of representation.

A typical MOT experiment runs as follows. Subjects are presented with a dis-

play containing two sets of objects: (i) target objects, which are supposed to be tracked during the trial, and (ii) non-target or distractor objects (t=1). Unsurprisingly, the number of target objects does not exceed four (respecting the subitizing range). After the display is introduced, targets usually blink on and off a few times so they can be distinguished from non-target objects (t=2). After that, targets move around the visually identical non-targets for 10 seconds (t=3). At the end of the trial, a given object is pointed at – the probe – and subjects are then asked whether or not the probe was in the target set (t=4). According to Pylyshyn (2007, 35), the task was extremely easy (around 90% accuracy) when up to 4 targets and a further number of non-targets were moving randomly, even passing in front of each other, at a reasonable speed (4-6 seconds to cross the computer screen on average, but that varied in different trials). By what means do subjects track these objects?
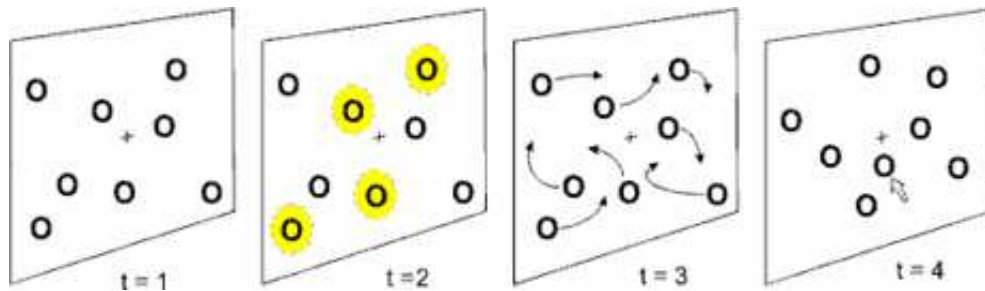
Figure 3.4: Diagram of the MOT paradigm in Pylyshyn (2001). Permission to reproduce this image has been granted by Elsevier.

As mentioned in the argument from re-identification, tracking or attending to an object as it moves in the display would seem to inevitably involve sortal concepts so that the object can be represented as the same object visualized at the beginning of the trial. Since the persistence conditions on material objects are captured by spatiotemporal continuity at successive times, one could cash out a sortalist model of object tracking, under the MATERIAL OBJECT concept, by claiming that subjects track multiple targets by storing their initial locations in memory and updating these continuously, as attention shifts to each successive position of the targets in turn.

According to Pylyshyn and Storm (1988), this proposal has two related difficulties. First, if objects move faster than the time that serial shifts of attention normally take, a huge decrease in performance would be expected. Second, if a target object is not at the stored location any longer, but a non-target is nearby, then updating the object's location could lead to greater probability of target/non-target confusion (see Pylyshyn 2007, 37). A tracking model which incorporates both factors - using the

average speed for human-like shifts of attention - yields a prediction of 30% tracking performance. However, Pylyshyn and Storm found that subjects performed at an accuracy level of 87%. Therefore, it appears that the parallel mechanism of object individuation could also be involved in representing objects' persistence conditions in the minimal sense of spatio-temporal continuity.

There is another way in which tracking might be done. By detecting the presence of certain shapes (e.g. *squareness*) and colours (e.g. *redness*) in the display, the visual system could track the object by tracking the features that co-occur in a given location (*squareness* $\wedge$ *redness* at l_2). Although the hypothesis does not support Sortalism, since features, such as *redness*, might not provide any criteria for individuation, it seems to support a form of descriptivism about object representation in the following sense: to experience an object amounts to experiencing the properties that uniquely describe it (see Treisman and Gelade, 1980; Campbell, 2002; Clark, 2004). Such a descriptivist strategy would predict that (a) if objects have distinctive features, (e.g. three objects consisting of one triangle, one circle, and one square) tracking them would be easier; and (b) if the objects' features remain stable during the trial, tracking would be easier as it does not require updating of features in memory. Dennis and Pylyshyn (2002) tested this and showed that having a uniquely discriminating shape does not improve tracking performance in MOT. Conversely, performance is no worse when objects' colours and shapes change randomly during the trial. Scholl et al. (1999) and Bahrami (2003) report that sometimes subjects were not even conscious of the object's changes. This further suggests that the objects were taken to be the same regardless of changes in their properties; that is, object tracking is not the same thing as tracking different bundles of perceptual features.

The results above strongly suggest that sortal concepts are not necessary for re-identifying the same object as it changes its position through time (the weaker notion of qualitative identity I have been assessing here). The MOT paradigm indicates that once the target objects blink, the early visual system assigns an index to the relevant objects in the display and the identity of the object is preserved under the identity of the index. Change in features would not change index assignment. Again, it must be noted that there are certain conditions not only for assigning visual indices, but also for preserving them. For example, Scholl et al. (2001) showed that objects' parts cannot be selected and tracked by visual indices; also, if objects liquefy and 'pour' from one place to another or if they move in a wormlike stretch and slink way, tracking is not possible (see VanMarle and Scholl, 2003). However, if target objects briefly disappear as if they were passing behind a 'trans-

parent' occlusion that renders them invisible, tracking is still possible (Scholl and Pylyshyn, 1999). Moreover, tracking is successful even if objects change the direction of their movement up to 60 degrees while occluded (Franconeri et al., 2006). Thus, in addition to the conditions of objecthood (i.e. bound contour, sufficient spatial segregation) unearthed by the subitizing studies, the MOT experiments suggest that visual indices can capture the weak notion of identity through space-time by a condition on object movement (i.e. spatiotemporal continuity). Interestingly, such object representations even survive momentary occlusions of the causal source.

In line with my previous remarks, the visual index mechanism relies heavily on spatio-temporal information: objecthood seems to be equated with instances of bound contours at a location and a time in the visual field. In order to test the role of location in object individuation, Blaser et al. (2000) asked a simple question: can humans track objects that are in the same location? In order to reach an answer, they devised a trial that consisted in tracking objects in *feature-space*, instead of physical space (within which they were static). The objects employed were two coloured, round, striped figures, called 'Gabor patches', one of which was superimposed on the other. As the Gabors had transparent backgrounds, the set of features pertaining to each object could be correctly identified. Now, subjects had to track movement in feature space, that is, track changes in the set of features of each object, which could be independently modified in the following way. Gabors could change colour, stripe width, or stripe orientation (moving clockwise or anti-clockwise). It would be possible for objects to occupy the same position in feature space with regards to one or two features (they could have the same colour and number of stripes, say), but at least one feature should remain distinctive.
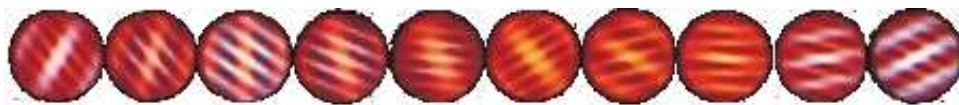
Figure 3.5: Diagram of the feature-space tracking paradigm Blaser et al. (2000). Permission to reproduce this image has been granted by Prof. E. Blaser.

Initially, subjects were designated a target Gabor as both of them 'moved' through feature space. After the trial, subjects picked out the right object with an accuracy rate of 90%, which strongly suggests that tracking targets did not engage an individuation mechanism based on spatial location. However, the question whether individuation is achieved by feature-based or by object-based mechanism theories of selection remains.

Let us suppose that the two gabor patches are individuated by their respective set of features, that is, objecthood equals a bundle of visual features (Treisman and Gelade, 1980; Campbell, 2002; Clark, 2004). Simplistically, this can be represented by a pair of colour - stripe orientation features, such as $_i$[*blue, anti-clockwise*] and $_j$[*red, clockwise*]. Here 'i' and 'j' supposedly represent the identity of the bundle, but it is important to stress that the difference between the two indices is explained by the difference in the feature sets. On the assumption that the subject distinguishes the gabors as soon as they start moving, it seems that the feature-based explanation of object individuation predicts that performance would be the same, whether the changing features belonged to the same object or different ones. To test this, Blaser et al. (2000) introduced pairs of small 'jumps' in the feature 'trajectories' of the objects in the display (i.e. sudden changes in orientation or colour, say). Such jumps were totally unpredicted, which has the advantage of working as data-driven stimulation. The pairs of sudden feature changes could either belong to the same object or to different objects. The task was simply to detect any 'jump'. In this experiment, subjects made faster judgements when the sudden changes belonged to the same object than when they belonged to different objects, (see Pylyshyn, 2007, 41). The selection of features seems, then, to be object-based and not feature-based.

3.3.4 More interim conclusions

I take my brief review of some experimental work on subitizing and the MOT paradigm to support the following claims.

First, it seems that the type of behaviour elicited by the subitizing tasks is governed by representations that do not quite fit the level of conceptual representations. The representations that guide behaviour in these tasks seem to be extremely encapsulated: the fact that an object has a uniquely discriminating shape, colour or label does not improve performance, suggesting that a conceptual token of, say, SQUARE, does not drive shifts of this type of attention. Plus, subitizing via conceptual representations, in response to the question 'How many objects lie in the same curve?' makes performance dramatically worse. Here, the concept CURVE seems to block a mechanism that was earlier engaged in individuation (Trick and Pylyshyn, 1994a,b). These results have been more or less duplicated in multiple object tracking (including feature-space) tasks. Additionally, subjects sometimes are not conscious of the changes the objects undergo[13] (Scholl et al., 1999; Bahrami, 2003), while concepts,

---

[13]There are many other systems for representing the identity of features of the world even though their respective stimuli changes. For example, humans are pretty good at preserving sameness of colour, size and shape even given radical changes in the distal stimulus impinging on the retina.

on the other hand, seem to underlie many of our conscious states. Finally, subjects are able to track co-located objects; as differences in location would be essential to distinguishing two instances of a concept or a set of features, it seems that representations that single individuals out descriptively are not operative in the feature space MOT task (Blaser et al., 2000).

The point of tracking (physically co-located) objects in feature space needs a bit of extra development. Blaser et al. (2000) suggest that spatial location is not that crucial for early object representation. As long as they have bound contours and movement is in a 'continuous' trajectory in feature-space, a bare representation of the object is established. If this is right, this study also suggests that individuating and preserving an object's identity in space-time (in this abstract sense) cannot be segregated, otherwise tracking in feature-space would be impossible. This also suggests that capturing spatio-temporal continuity happens at a more abstract level than just the occupying of adjacent positions in three dimensional space at successive times. If this is the case, the sortalist could claim that there is no real difference between tracking in feature space and tracking into three-dimensional space, and that both require some form of conceptual representation.

Let us now take stock and relate this body of evidence to the preceding discussion. In the first section of this chapter, we asked whether indirectly available concepts could integrate with the encoded meaning of expressions in order to express a proposition (i.e. the CMI question). The question was raised by certain accessibility puzzles displayed by a variety of linguistic phenomena (i.e. the asymmetry between the felicitous 'every man who has a wife should bring her to the party' and the infelicitous 'every married man should bring her to the party'), which motivated a series of constraints in linguistic theory, as, for example, Elbourne's conditions on NP ellipsis. Since such constraints are also designed to capture cases of 'deep anaphora' (e.g. 'Do most faculty members deposit it [*waving a paycheck*] in the credit union?'), they would have to appeal to a representational level that is not restricted to linguistic representations, but includes concepts that emerge from certain immediate cues in the environment. This type of constraint, however, results in a negative answer to the CMI question, and, counter-intuitively, identification-based descriptive pronouns cannot contribute to the proposition expressed. But if we want such pronominal uses to contribute to the proposition expressed, how can the accessibility pattern of donkey uses above can be explained? We are at an impasse.

Burge (2010) calls these identity preserving mechanisms 'natural constancies', which Pylyshyn's visual index system would be an example of. Thus, there are colour, shape and many more constancies in human representational systems in addition to the one involved in early object individuation.

In the last section, we looked at the necessity to ground, empirically, the notion of a directly available concept or an immediate cue for an interpretation. The philosophical idea that some form of representation is necessary to partition the world into discrete particulars, namely, sortals, has received empirical support from a variety of experimental paradigms. There is evidence suggesting that infants use concepts like MATERIAL OBJECT for a variety of cognitive tasks related to individuating and capturing persistence conditions on objects from a very early age. As more concepts become available, e.g. DOG, BALL, TOY, the agent can employ more specific information in cognitive tasks involving the objects that fall under them, such as individuation and the representation of an object's persistence conditions. Thus, these studies provide empirical support for the idea that certain conceptual representations are basic in the cognitive architecture of agents, for they are required for characterising the way humans experience objects. This basic status, in turn, can explain the direct availability of certain concepts for integration with linguistic information, as stated in the constraints just mentioned.

The studies reviewed in this section, however, go against this conclusion. There is evidence that the type of representation that governs object individuation and captures an object's spatio-temporal continuity is non-conceptual. The conditions on representing material objects (i.e. bound, move as a whole, etc.) can be captured in terms of hardwired *natural constraints* that spit out a bare object representation as output: a visual index. If this is the case, then the argument that establishes the direct availability of sortal concepts in virtue of their necessary status in partitioning the world into discrete individuals is threatened. As a consequence, we must seek an alternative explanation of how concepts (even non-sortal ones) emerge in cognition. As we shall see, this will pave the way for allowing indirectly available concepts to integrate with the encoded meaning of expressions.

Apparently, we have moved from a dilemma concerning communicators' intuitions to a dilemma involving divergent evidence. In the next section, I will present an interpretation of the empirical data that offers a way out of our second dilemma and paves the way to the solution of the first one as well. This will also provide an answer to the CMI question presented initially.

3.4   Levels of representation revisited

The preceding sections have raised interesting questions. Do these two sets of data provide evidence for a single system of object representation or two distinct ones? If they are the same, would output representations be conceptual or non-conceptual

in nature. If distinct, how would the different output representations be related. In order to assess possible answers, I shall first take a look at analogies between the two sets of experimental results and then proceed to their disanalogies.

3.4.1   Analogies between MATERIAL OBJECT and visual indices

Carey (2009), building on Carey and Xu (2001), points out three analogies between the two bodies of experimental data reviewed: (i) primacy of spatio temporal information in object individuation, (ii) a representational limit of four entities in working memory, and (iii) the ability to represent the distinction between existence through occlusion and cessation of existence. I will take each of these in turn.

Primacy of spatio temporal information means that human cognitive architecture privileges bound contours and spatio-temporal continuity in motion to individuate objects and preserve their identity through time and space. In the MOT experimental design, participants were quite good at tracking multiple objects as they moved around, regardless of whether or not they had uniquely identifying shapes, colours, or labels. Moreover, robust changes in the type of object motion (liquefying and pouring from place to place, disintegrating and reassembling (see VanMarle and Scholl 2003), severely impairs tracking skills. Similarly, Spelke, Carey, Baillargeon, Xu and colleagues tested children's expectations concerning the behaviour of objects in the world in terms of their spatio-temporal properties. In Spelke et al. (1995), children seem to use the gap between two screens as a way to determine how many objects are occluded by them. Replicating these studies using piles of sand instead of bound objects shows that children are unable to numerically individuate one object from the other (Huntley-Fenner et al., 2002). Also, Xu and Carey (1996) show that 8-month-olds cannot succeed at the object-permanence task if two objects are occluded behind a single screen. Carey (2009) rightly concludes that both bodies of experimental research support the claim that spatio-temporal information plays a crucial role in object representation.

A limit of up to four (or five) object representations has also been unearthed by both literatures. The first piece of evidence for this stems from the subitizing phenomenon: up to 4 objects can be easily enumerated, but performance drops dramatically if the set size is slightly increased. In the MOT paradigm, this limit becomes manifest quite clearly: up to four of five objects can be tracked simultaneously, but not more. Various studies in the infant literature converge on the same conclusion. Here I will confine myself to two. First, Rose et al. (2001) and Ross-sheehy et al. (2003) presented infants with an array consisting of a certain number of objects.

The array disappeared and then it emerged again either without any changes or with one of its objects (chosen at random) having its colour changed. The arrays could contain 1, 2, 3, 4 or 6 objects. If infants were able to create models for the objects in the first display, they would be more interested in the display with the change as opposed to the constant display. This hypothesis turned out to be true for one object, in the case of 4 and 6 month-old infants, and up to 4 objects in the case of 10-month-olds. Since the subitizing range is lower than 6, both groups of infants performed poorly in the 6 objects condition. This suggests a period of maturation of the subitizing range. Moreover, Feigenson et al. (2002) and Feigenson and Carey (2005) used a manual search paradigm to shed light on children's memory capacities. They placed a number of graham crackers into one bucket showing one at a time to 10 or 12 month-olds and then another number of crackers, in the same manner, into another bucket. Children would then crawl to a bucket and search for crackers. If infants were able to represent the number of crackers, they would choose the bucket that contains the most. Children performed as predicted when the choice was 1 vs 2, 1 vs 3, and 2 vs 3, but behaved randomly when one of the buckets contained 4 crackers (even in a 1 vs 4 condition). Thus there seems to be converging evidence for the idea that representing up to 4 objects has a privileged status in cognition.

Finally, *the ability to represent the distinction between the occlusion of an object and the cessation of an object's existence* is attested in both bodies of literature. In the MOT paradigm, Scholl and Pylyshyn (1999) showed that if objects pass behind a fixed zone in the display that renders them momentarily invisible and objects' movement patterns remain the same, tracking performance is unaltered. Also Franconeri et al. (2006) showed that if objects change their directions up to 60 degrees while 'occluded' in this way, they still can be tracked. If visual indices were re-assigned during these critical moments of occlusion, a drop in performance would be predicted. This is not the case. Similarly, in the infant literature, representing objects as persisting throughout occlusion is crucial. If the 1 + 1 task involving two objects hidden behind screens could be explained in terms of infants representing the destruction of a given object and the emergence of a new one, then there would be no surprise when only one of the objects appears as an outcome: the other simply got destroyed. This is not empirically borne out. Note, however, that this does not preclude this distinction from being captured *in different ways*: it is perfectly possible that, at the level of early representations of objects (visual indices), survival through occlusion is captured by sameness of index and cessation of existence by index re-assignment (reflecting reduction in performance). This would be achieved

by features of the hardware itself, not by the fact that a MATERIAL OBJECT concept is necessary for stating under which conditions a material object is occluded and under which it ceases to be.

Carey and Xu (2001) and Carey (2009, ch. 4), based on this common set of three 'signatures', argue for a single system of object representation. One that is conceptual in nature, given their arguments for the conceptuality of such representations, namely, the integrative role of such representations and their interface with linguistic systems (see section 3.2). Before assessing the soundness of the argument, it is worth taking a look at the disanalogies between the two bodies of evidence.

3.4.2 Disanalogies between MATERIAL OBJECT and visual indices

The similarity between the two systems of representation seems to break down in the ways they interface with either a given kind of stimuli or with the conceptual system. I will take up these two ways in which the point can be illustrated in turn. These basically amount to tensions between the interim conclusions of the two preceding sections.

First, it seems that visual indices can apply to *pre-3D* stimuli. In some of the MOT experiments (all presented using 2D displays), objects could pass through each other without any loss in performance, clearly violating Spelke's principle of cohesion (Baillargeon, 1987). Moreover, objects that do not stereotypically fall under MATERIAL OBJECT could, in principle, be visually indexed, such as shadows and holes (see Casati, 2004). More impressively, 'objects' can even be tracked in feature space. None of these forms of representation can be achieved under MATERIAL OBJECT. Stimuli represented under this concept, on the other hand, give rise to complex expectations concerning the possible interactions amongst different kinds of physical and animate entities in *real* space.

Second, based on the findings revealed by the MOT experimental paradigm, Pylyshyn argues that visual indices are unlikely to be *conceptual* because of their degree of encapsulation. Tracking multiple objects that have a uniquely identifying shape, colour, or size does not improve performance (this suggests that visual indices may be encapsulated even from (possibly) other non-conceptual representations). Crucially, giving a uniquely discriminating label for each object does not improve performance. These data contrast with the findings of Xu (1999, 2002, 2005); Surian et al. (2004); Dewar and Xu (2007). The success of 9-month-olds on the object permanence task involving two objects behind a single screen, usually achieved at 12 months of age, depends on cues that make information about object

kind salient to the child. Thus, providing a common noun for each object greatly improves performance. This does not seem to be attested in the MOT studies: *any object*, as long as they are bound and move continuously, can be indexed.

3.4.3   Accounting for the dilemma

The following dilemma emerges at this point. How can we explain the partial over-lap between the MOT based experimental findings (Pylyshyn and colleagues) and those discovered by the child expectancy-violation paradigm (Spelke, Baillargeon, Carey, Xu and colleagues), as well as their disanalogies? On the one hand, assuming that object representations are the output of a single, non-conceptual, encapsulated and data driven system allegedly faces the difficulty of accounting for the robust set of expectations and complex inferences that seem to be regulating children's be-haviour. On the other hand, assuming that object representations are concepts faces the challenge of accounting for the uniformity of the tracking performance across conditions that would allegedly make the task easier, such as an object's having a uniquely discriminating colour, shape, and, crucially, label.

I believe that the way out of this dilemma involves making room for the idea that visual indices and concepts play different, yet related, cognitive roles and pos-tulating a notion of *incremental object representation*.

The disanalogies reviewed above rest on the many interesting properties to the mechanism of visual index assignment. More specifically, the data reviewed in section 3.3 suggest the following claims: (i) visual indices play a role in the in-dividuation of objects based on pre-3D images as well as capturing their identity in the minimal sense of sameness construed as continuity through space-time, and (ii), given the need for mechanisms that determines that a set of features belong to (causally derive from) the same object (Treisman and Gelade, 1980), visual in-dices may provide a reference point to visual information processing units sensitive to other types of information, such as colour and shape. If information about the object's shape (say, [*square*]) and colour (say, [*yellow*]) are processed by different streams of activation in the brain, then adding information (say, [*square* $\land$ *yellow*]) necessitates a common reference point, provided by the visual index (say, $_i$[*square* $\land$ *yellow*]). Effectively, something like visual indices would play a fundamental role in the integration of information in the visual system. If this is right, concepts are the not the only representations that are content integrating (*contra* Carey's argument in section 3.2.3).

There is another angle from which this integrative role of visual indices can

be viewed. In the previous paragraph, I alluded to the addition of featural (non-conceptual) information in a rather static way. The process mirrors an inference, but at a possibly non-conceptual level: from [*square*] and [*yellow*] one can move to $_i$[*square* $\wedge$ *yellow*]. However, there is also a more dynamic perspective on object representations. In the MOT experiments, there were trials in which the object's properties changed constantly, yet were not reflected in performance. In order for changes of an object's properties to be perceived, the link between the object and a cognitive system must be preserved through time. This is captured by sameness of visual index. Thus, these nonconceptual representations not only allow for information to be added, but also to be updated through time by preserving an information channel between object and cognitive system (an important aspect of *cognitive dynamics*). This can regarded as yet another beneficial cognitive effect brought by such early-level object representations.

In conclusion, the review in section 3.3 argues for three inter-related claims, namely, (i) that there are representations that may fall below the conceptual level, (ii) that such a kind of representations plays a role in the individuation of objects as well as representing the conditions under which they remain the same in time-space, and (iii) that by playing the latter role, visual indices also maintain a link to the object, which allows information to be added or updated.

These claims seem to presuppose a notion of *incremental object representation*: our representation of an object evolves through time (via the preserved information-link achieved by something like a visual index). This temporal parameter introduces another important aspect: the richness of the object representation may be relative to certain cognitive tasks faced by the system. Some of the studies reviewed in previous sections raised interesting limiting cases for our ability to represent objects. Treisman and Gelade (1980) showed that *failures* in the 'binding' of visual features can occur, that is, when subjects are presented with a display containing a red triangle to the right and a yellow rectangle to the left for a very short time, they cross-linked shape and colour information, representing a red rectangle and a yellow triangle. This suggests that the system did not have sufficient cognitive resources to produce the richer information state in the right way due to cognitive load. Similarly, Scholl et al. (1999) and Bahrami (2003) showed that when subjects really concentrate on tracking multiple objects whose properties are changing constantly during the trial, the changes went unnoticed by the subjects. This suggests that some of the information derived from the object did not surface at the conceptual level, indicating that visual indices were enough to carry on with the task at hand.

The relativity to processing time, task achieved, and cognitive resources employed in doing so highlights one final aspect of the notion of incremental representation that I am putting forward: a prior representational state about an object is necessary for and informationally poorer than further representational states about the same object. Latter states engage more cognitive resources but allow for more information to be aggregated. Conversely, earlier states rely on shallower input analysers, but might provide a good enough representation of the object if current cognitive demands or environmental constraints necessitate a certain course of action - e.g. ducking an unrecognised object that moves towards one's head - before the system is in a position to process more fine-grained information about the object, such as which kind it belongs to.

If the foregoing remarks are on the right track, object representation seems to be an effect driven activity that takes place against the background of certain environmental constraints (both external and internal to the system). This characterisation of object representation displays clear parallels to the relevance theoretic *cognitive principle* outlined in the previous chapter. Sperber and Wilson (1995, 260-266) claim that cognition is geared 'towards the achievement of as many positive cognitive effects for as little processing effort as possible' (Cognitive Principle of Relevance). The cognitive regulations imposed by this principle, in effect, could enrich an object representation *modulo* the task at hand and the resources available. This provides an interesting tool for understanding how the gap between conceptual and non-conceptual representations can be bridged. *Conceptuality can be considered a cognitive effect*[14].

In a similar way that a visual feature bundle can be regarded as an enrichment of a bare object representation, early forms of non-conceptual representation and concepts can be related via an incremental notion of processing. The cognitive effects purchased by conceptual representations are obvious. Conceptual information can be stored as what psychologists call 'semantic memory' - the individual's database of knowledge and belief - and interact with other conceptually structured information inferentially (deductively or inductively) in the derivation of further information. Thus, as information is processed through the visual pathways, a series of mental states is established, with the goal of cognitive benefit. Object representations begin at the visual index level ($_i$) and proceed to $_i[round \land blue_{345}]$, the level of feature bundle (Campbell 2002) or object-file (Kahneman et al. 1992), and then

---

[14]Ness (2011) argues that concepts are necessary in order to provide the thematic unity of the phenomenal content of a given experience. Perhaps thematic unity could be regarded as another positive effect achieved at the conceptual level.

go to a conceptual level of information ($_{i(round \wedge blue_{345})}$[BALL]), at which propositional information associated with that concept (DOGS LIKE TO CATCH BALLS) is available for further cognitive tasks, such as inference and memory retrieval and storage.

This picture of object representation solves the dilemma at the beginning of this subsection in the following way. The analogies unearthed by the two bodies of experimental data are explained by the incremental nature of processing: given the dependency of a higher level state on a lower level one, it is natural to expect that features that constrain the latter (limit of 4 indexed objects) also constrain the former. However, the conceptual nature of the higher level states enables certain cognitive tasks, e.g. the possibility of interaction with semantic memory and inference - that are precluded of non-conceptual representations. This explains the disanalogies between the two bodies of experimental work. The encapsulated nature of representations involved in MOT, on the one hand, seems to be governed by early level representations, as the cognitive effects in question boil down simply to the preservation of the identity of objects moving in a display. The object-permanence studies, on the other hand, comprise more cognitively loaded tasks. Infants have to consider the transformations and interactions which objects might have undergone while occluded, which draws on more stable (less occasional) information.

The fact that words can improve children's performance suggest that they activate specific pieces of information (concepts) that are associated with knowledge relevant to the task at hand, as in the studies by Xu (1999, 2002, 2005) and Dewar and Xu (2007). The findings in these studies are also explained by the fact that incremental processing is geared towards achieving beneficial cognitive effects. Many of the expectations manifested in the two separate screens object permanence task of Spelke (1990) and Spelke et al. (1995) (among many others) seem to be underpinned by MATERIAL OBJECT, whereas the success in the object permanence task in the single screen studies of Xu and Dewar depends on the availability of two distinct conceptual representations (e.g. DUCK, BALL)[15] that allow for an inference of the following sort: if two kinds are present, then there must be two entities that instantiate them behind the screen. Children fail when they have not acquired the distinct concepts, because only one kind-level representation - i.e. MATERIAL OBJECT or TOY - is involved in the task, which would not naturally give rise to the expectation of two entities instantiating that kind being occluded behind the screen.

---

[15]The proposal made here is neutral with regards to whether these more specific concepts involve a re-conceptualisation of the world by the child and to whether certain object permanence tasks (Xu, 1999, 2002, 2005; Dewar and Xu, 2007) are explicable by beliefs associated with concepts acquired by the child (see Carey, 2009, 279-85).

This suggests that the notion of effect oriented incrementality also holds at the conceptual level (and throughout cognition, more generally): when children are able to use more specific conceptual representations, more effects can be generated.

Now, if we assume that there are non-conceptual representations (necessary for some early forms of object individuation and identity preservation) and that the incremental notion of processing mentioned above makes representations surface at the conceptual level at some point in the evolution of mental states, then it seems that *sortal are not necessary for individuating objects and capturing their spatio-temporal continuity*. Still, *concepts* play an important role in capturing an object's identity conditions in a broad sense: realising the totality of transformations an object can undergo whilst remaining the same object (Wiggins, 1967, 1980, 2001) is subject to the thinker's knowledge about the particular kind the object belongs to, which inevitability requires a concept that denotes that kind. Other roles include categorising and stating generalisations about entities that fall under the relevant kinds. Consider plants, for example: they are not self-propelled, so they do not fall under our innate concept for agents, but they also do not fall under MATERIAL OBJECT either. Thus, evaluating whether a seedling is the same as a large oak some time later requires not only a PLANT concept, but certain beliefs and knowledge about the transformations these objects can undergo. This indeed is a conceptually loaded cognitive task, but quite different from individuating and tracking simple objects moving in the visual field. In the next section, I will sketch how the notion of incremental processing can be extended from early mechanisms of object representation to (directly or indirectly available) conceptual representations that are suitable for linguistic integration.

3.5 Incremental object representation, the CMI question, and descriptive pronouns.

At this point, we must take a look at how the notion of incrementality of object representation gives us an answer to the question raised at the beginning of this chapter, namely:

(*CMI Question*): Can indirectly salient concepts be constituents of the proposition expressed by an utterance (i.e. count as a development of the logical form of an utterance)?

The initial negative answer to this question was motivated by a hypothesis that linked the notion of a sortal to that of concepts which would be directly available to the audience because they are necessary in the characterisation a particular visual experience as the experience it really is (i.e. as experiences of a discrete, bound objects that move as a whole, see the theses *Sortal Individuation* and *Sortal Identity* in this chapter's introduction). That is, the justification for sortal concepts constituting such a privileged level of representation relied on their necessary role for individuating objects as well as capturing their persistence conditions. The hypothesis offered a way to sensibly capture the notion of a *immediately available concept* used in linguistic constraints, such as Elbourne's condition on NP ellipsis or the Individuative-Representational constraint, which were motivated to explain some accessibility patterns in anaphora (and possibly descriptive uses), illustrated below:

(6)   a.   Every man who has **a wife** should bring **her** to the party.
       b.   *Every **married** man should bring **her** to the party. (intended reading: as (a) above)
       c.   Every groom thinks his bride is worth dying for.
       d.   Every groom thinks *she/?his [*pointing to a wedding dress*] is worth dying for.

Above the concept MARRIED makes the concept WIFE indirectly available and the gesture towards the wedding dress makes the concept BRIDE indirectly available as well. Therefore, such representations cannot be selected as the values of the respective pronouns. Moreover, justifying the immediate availability of certain concepts because they are required by a given perceptual experience enables us to deal with cases where certain dependencies are established perceptually, as in deep anaphora, below:

(7)   Do most faculty members deposit it [*waving a paycheck*] in the credit union (Jacobson, 2000)?

Singling out the object waved by the speaker would require a token of the PAY-CHECK concept, which would then be integrated with the linguistic meaning of 'it'. However, as a result all cases of identification-based descriptive uses of pronouns (as well as some cases of donkey anaphora) cannot contribute to the proposition expressed and an alternative way to explain this linguistic data would be motivated.

This tentative hypothesis has serious predictions and was assessed empirically. In section 3.2, I reviewed data in support of the idea that this basic level of repre-

sentation necessitates sortal concepts. In section 3.3, I reviewed data that argues against the very same idea. Thus, a new dilemma emerged in section 3.4: how to make sense of both sets of evidence. This section aimed to coherently explain both the analogies and disanalogies between the findings unearthed by the two bodies of experimental work reviewed previously. The answer involved a notion of incremental object representation, which falls out from the combination of general mechanisms that regulate the flow of information (e.g. the Cognitive Principle of Relevance) and an empirically based picture of how human visual processing mechanisms actually work, that is, a description of human visual abilities.

The adoption of this picture undermines the link between sortals and their role in individuating objects and capturing their persistence conditions in the spatiotemporal sense. As a consequence, the justification for sortals constituting a privileged level of representation in virtue of their role is lost. Still, incremental processing explains much of the data that was recruited in support of sortalism as well as data against it.

Representation begins by early (visual) processing mechanisms that individuate, capture the identity conditions of material objects, and aggregates (possibly nonconceptual) information about many of their properties, such as colour, shape, size, edge and so on. This accounts for the data against sortalism reviewed in section 3.3. We can think of these as visual counterparts of early auditory mechanisms involved in detecting many of the properties of speech, such as, tone and pitch, as well as identifying phonemes and determining word boundaries. Given the incremental nature of processing, these representations eventually surface at the conceptual level. Conceptuality could be regarded as a cognitive achievement which brings many beneficial effects to the thinker. This level of representation accounts for the data reviewed in section 3.2. However, the concepts that emerge do not do so because they are necessary to play an individuating role. That has already being achieved by early mechanisms. Still, we may preserve the idea that some *concepts* emerge *naturally* in the processing history of a visual episode, as, for example, when the seeing of a dog gives rise to a DOG concept, not because the concept is required for the individuation of the object, but because it is the representation that naturally emerges at the conceptual level and allows for cognitive effects that are semantic in nature (i.e. semantic memory retrieval and storage). A similar story holds for the processing of words: a DOG concept emerges in virtue of early level processing mechanisms that identify the word (e.g. 'dog') associated with it.

Note that this proposal makes room for an important unification. The emergence of non-sortal concepts is explained in the same way. Consider someone looking at

a large body of smoke. Early processing mechanisms (distinct from visual index-ing) would be able to represent certain properties of the entity in the environment (smoky stuff) in a way such that a concept (e.g. SMOKE) surfaces at some point due to the incremental nature of processing. This concept could then be available to feed the interpretation of various expressions in natural languages, even though it would not be a sortal. For example, it could be the referent of the 'it' or 'that' and it could be also used for the processing of certain identification-based descriptive uses of pronouns (e.g. one of two people engulfed in a smoke cloud at a demon-stration saying 'They are bringing reinforcements', where 'they' is interpreted as *the police*). In other words, even though the mechanisms responsible for the early processing of different kinds of aspects of the environment (smoke vs chairs, say) are underpinned by quite different output representations, they would eventually give rise to representations in a common format (e.g. SMOKE, CHAIRS concepts). Representations that emerge from standard material objects and those that do not figure in the same way as part of the context for interpreting overt stimuli.

Moreover, the picture sketched here answers the CMI question in the follow-ing way. The allegedly 'indirectly' available concepts *can* integrate with linguistic meaning in order to express a proposition, because the same notion of incremen-tality that would make concepts emerge after certain low-level processes fulfil their role would also make related concepts available. For example, seeing the White House involves processing information in a way such that the WHITE HOUSE or PUBLIC BUILDING concept eventually emerges in a person's mind, making White House-related information - such as the concept AMERICAN PRESIDENT and the propositions THE AMERICAN PRESIDENT WORKS AT THE WHITE HOUSE, THE WHITE HOUSE IS AN IMPORTANT BUILDING, and so on - salient or accessible for whichever cognitive task the system currently faces. In the case of interpreting (1)e, reprinted below, incrementality in object representation would render AMERICAN PRESIDENT accessible enough to satisfy the interpretative demands.

(8)   *Pointing to the White House in Washington DC*: Every time there is a war, he [~~the American president~~] has tough choices to make.

On the other hand, the gesture towards the wedding dress in the context of (6)d does not reach a threshold of activation that enables the integration between linguis-tic information and salient conceptual information (maybe due to the availability of linguistic forms that would establish the relevant interpretation at lower cognitive costs, or the availability of competing representations at a low cost). Thus, the incremental notion of processing allows a more empirically grounded solution to

not only the empirical dilemma concerning divergent data, but also the intuitive dilemma concerning whether a positive or negative answer to the CMI question should be given. Note that this type of explanation also carries over to some cases of anaphora where the antecedent is not stated explicitly but the dependency is successfully established, such as:

(9)   Every **iphone-owner** uses **it** for browsing.

Now, the role the incremental picture of processing (underpinned by the Cognitive Principle of Relevance) plays in the case above and in the White House scenario in (8) and the assumption that *Covariation* (e.g. between US presidents and wartimes, or iphone-owners and their mobiles) should be explained by some form of grammatical process has *important consequences* for our conception of *grammar*. That is, this chapter argues for an explanation that allows pragmatically available constituents (available through some form of incremental processing) to figure in the mechanism that explains the relevant co-varying readings. Conversely, the narrow string-based view of grammar sketched at the beginning of this chapter does not seem to get the data right (e.g. (8) and (9)), nor is compatible with the incremental picture presented here. In chapter 5, I will sketch a grammar (Dynamic Syntax, Kempson et al. 2001; Cann et al. 2005) that builds instrumentality in its very heart and does not suffer from such shortcomings. Its adoption can account for the *Co-variation* property of identification-based descriptive pronouns and cases of anaphora straightforwardly, thus offering an unified explanation for the variety of pronominal uses discussed in chapter 1.

Finally, the asymmetry between the felicitous 'Every man who has a wife should bring her to the party' and the infelicitous 'Every married man should bring her to the party' and the infelicity of certain identification-based descriptive pronouns should not be accounted for by the postulation of grammatical and/or representational constraints, but only by serious work in psychology aimed to establish the threshold of activation that certain representations need to reach in order for them to be recruited for integration with linguistic expressions in utterance comprehension. This would be in total consonance with the incremental view of processing put forward here. Certain cognitive effects, in this case, establishing certain interpretations, would be influenced by frequency and recency (as well as many other) variables. This would explain an interesting fact about the infelicitous cases mentioned above. As one repeats them to oneself, their acceptability rises. I believe this would be very much in the spirit of recent approaches to metonymy (see Falkum, 2010, ch. 6) and the same point could be made for bad cases of donkey anaphora

that motivated the formal link condition and the constraints stated at the beginning of the chapter.

In conclusion, the processing of descriptive pronouns can result from the integration of linguistic content and content delivered visually, in a way similar to integration underpinning deep anaphora. This paves the way to the full development of a *substitution-based* account of some descriptive cases, as suggested towards the end of chapter 2. Before doing so, we must assess some of the literature that attempt to explain the data presented in chapter 1 and that may appeal to some of the processes involved in visual representation discussed here.

Chapter 4

Descriptive pronouns: previous accounts

4.1   Introduction

This chapter marks the start of a second stage in the thesis: one in which we begin to assess different theories of identification-based descriptive uses. Thus, it is worth reminding ourselves of what has been achieved so far and how the conclusions reached in previous chapters shape possible accounts of the data.

In chapter 1, the properties of identification-based descriptive pronouns were thoroughly described. Comparing them to a variety of other pronominal uses and to one kind of non-literal interpretation (metonymy) brought both the core and extended properties of descriptive pronouns to light (i.e. *Identification Dependency, Descriptiveness, Connection*, and *Accessibility*). More specifically, such uses seemed to be most similar to donkey anaphora and metonymy.

In chapter 2, we have dug deeper into the latter similarity. Explaining descriptive uses of pronouns as non-literal interpretations seems impossible from a Gricean perspective and although this is more sensible in alternative theories, such as RT, there are important dissimilarities between descriptive pronouns and metonymies: since the former do not encode lexical concepts, postulating a 'substitution' between linguistically-driven conceptual representation and speaker-meant concept misses the point. Still, this chapter laid out some important cornerstones. Some descriptive uses seem to contribute to an utterance's explicature and some others to an utterance's implicatures. The best criterion for deciding which analysis is preferred seems to be provided by the scope test. While we have an account of implicature derivation (by global inference), contributions to the explicit level of meaning still

need to be explained by some mechanism. As we shall see in this chapter, this can rely more heavily on either grammatical structure or pragmatic principles.

In chapter 3, we have assessed whether identification-based descriptive interpretations result from processes that are radically different from donkey anaphora. In the literature on the latter type of dependency, many have proposed that concepts that are 'indirectly' available in the context cannot provide antecedents for pronouns, hence predicting that the two cases are essentially different phenomena. This hypothesis has been assessed thoroughly from both conceptual and empirical standpoints and it seems to be unwarranted. In the course of carrying out this task, we have reviewed many assumptions about the processes of individuating and identifying a material object in the environment which may be invoked by different types of proposals. In conclusion, treating identification-based descriptive pronouns as a special form of anaphora (dependent on an indirectly and extra-linguistically available antecedent, as it were) seems to be a viable option. This will be pursued in chapter 5. Note that this strategy is not exhaustive. As discussed towards the end of chapter 2, some identification-based descriptive uses are better captured as implicatures. The grammar to be described in chapter 5 together with the relevance theoretic pragmatic machinery introduced in chapter 2 can accommodate such pronominal uses at either level of meaning, which turns out to be a great advantage over competing accounts. Before laying out the proposal, however, we must look at alternative explanations.

In this chapter, I aim to review some of the accounts of the data proposed in the literature. I will begin with a reminder of the *desiderata* that a theory of descriptive uses of pronouns must meet (section 4.2). Broadly speaking, there are two families of theories that have attempted to account for these uses . In section 4.3, I will assess proposals that do so by postulating a more complicated grammatical (i.e. syntactic or semantic) machinery, that is, the so-called 'heavy-handed semantics' positions, in Neale's (2007) terminology. In section 4.4, I turn to theories that have tried to explain the phenomenon in pragmatic terms, that is, 'heavy-handed pragmatic' accounts, according to Neale (2007). The specific instances reviewed, unfortunately, do not exhaust every single position in the market. I have opted to focus on the more influential ones. The presentation of each within these two groups is followed by an assessment of its shortcomings. Section 4.5 concludes and paves the way to my own account of the data that assimilates the insights from other chapters, as mentioned above, and does not face the shortcomings of the proposals about to be reviewed.

4.2 The *desiderata*

As a quick review of the interpretative properties discussed in chapter 1, descriptive uses, such as:

(1)   a.   Someone pointing at a huge footprint in the sand says: 'He must be a giant' (from Schiffer 1981, 49) .

         b.   US Supreme Court Justice O'Connor says: 'We might have been liberals' (Nunberg, 1993, 14-15).

         c.   Bill Clinton saying: 'The founders invested me with the sole responsibility for appointing Supreme Court Justices' (Nunberg, 1993, 20).

         d.   John, expecting a call from his mother, answers a phone call from his colleague and says: 'Oh, I thought you were my mother' (Nunberg, 1990).

seem to have the following interpretative properties:

Descriptiveness: The information conveyed by the pronominal use is descriptive or general.

Identification Dependency: The pronominal uses depend on the identification of a particular entity as belonging to a certain kind.

Connection: An adequate account of the pronominal uses must explain the relation between *Descriptiveness* and *Identification Dependency* - that is, how the concepts used to represent objects as falling under a certain kind make the descriptive interpretation available in a principled way.

In chapter 1, these were referred to as the core interpretative properties of descriptive pronouns. Note that I am not saying that any use of these utterances *must* convey descriptive information at the explicit level, I am just saying they can or might. The sentences in examples (b) and (c) could easily be used to express singular propositions about Clinton and a group containing O'Connor, respectively (which could be used to implicate, or otherwise be used in the communication of, descriptive content). Then, in addition, uses like the one in (2) and the contrasts between (a) and (b) and that between (c) and (d) in (3) seem to motivate the extra *desiderata* below:

(2)   *Andy*: Well...**Every time there is a war**, **he** [*pointing to the White House*] has tough choices to make.

(3)   *Antecedent impossibilities of descriptive pronouns* (boldface indicate potential dependencies)

   a.   *Said by someone*: **A/The condemned prisoner** is traditionally allowed a last meal. **He** is also allowed to invite close friends for the execution.

   b.   *Said by a condemned prisoner*: **I** am traditionally allowed a last meal. ***He** [*said by same speaker*]/ **He** [*said by different speaker*] is also allowed to invite close friends for the execution.

   c.   *Said by someone:* If the Democrats had won the last few presidential elections, **the American Supreme Court Justices** might have been liberals. **They** would've guarded public interest better.

   d.   *Said by Supreme Court Justice O'Connor:* If the Democrats had won the last few presidential elections, **we** might have been liberals. *****They** [*said by same speaker*]/**They** [*said by different speaker*]/?**We** [*said by same speaker*] would guard public interest better.

Co-variation: The pronominal interpretations can co-vary with another expression in discourse.

Accessibility (descriptive): Certain mental representations can provide the right interpretations for (descriptive) pronouns, while others cannot.

   In short, we want a theory that is able to account for the core properties of the data, namely, *Descriptiveness, Identification Dependency* and *Connection* as well as the extended properties of the data, namely, *Co-variation* and *Accessibility*. Moreover, we would like a theory that explains the interpretative properties of other uses of pronouns, reviewed in chapter 1, in a unified manner. In the next two sections, we will take a look at how different families of theories try to achieve this goal. Anticipating things a bit, it seems that all of them face some problems, when it comes to explaining the extended properties of the data. We turn now to accounts that do so by postulating an intricate grammatical or semantic machinery.

4.3 Heavy-handed semantic accounts

In this section, I focus on theories that deal with the data at hand by postulating a complex formal semantic apparatus. I will concentrate on Nunberg's (1993) original insight.

4.3.1 Nunberg's three-component account

Nunberg's proposal

In order to account for the interpretative properties above, Nunberg (1993) assumes that personal and demonstrative pronouns encode three components: (i) *a deictic component*, which can be understood as descriptive information that works as 'a function from occurrences or utterances of an expression to elements of the context of utterance' (Nunberg, 1993, 8), for example, *the speaker* is the deictic component of 'I' and determines a particular entity in a context , (ii) *a classificatory component*, responsible for guiding the overall interpretation process and consisting of features such as gender and/or animacy, and (iii) *a relational component*, which constrains the relation between the object the deictic component determines, that is, *the index*, and the intended interpretation. In the case of first- and second person pronouns ('participant terms', according to Nunberg 1993), the relational component imposes that the 'index must be included in, or more generally, must instantiate the interpretation' (Nunberg, 1993, 9). Nunberg's notion of interpretation here is just the pronoun's truth-conditional contribution, which might be the deictic, or anaphoric, in addition to the identification-based descriptive interpretations discussed here.

This machinery explains the data as follows. Clinton's utterance in (1)c contains an occurrence of 'me'. Since first-person pronouns encode something like 'the speaker' as deictic component, it determines Bill Clinton as index. Now, the index is related to the final (and intended) interpretation via the classificatory component: the pronoun's animacy feature must be preserved by the final interpretation, which also must include or be instantiated by the index. In (1)c, the alleged intended interpretation is *the American President*, which is related to the index, Bill Clinton, in a way such that the information encoded by the classificatory and relational components are preserved (presidents are animate entities). The interpretation is also instantiated by the index (Clinton was a President). The three components equally account for the 'literal' deictic cases, with the exception that the intended interpretation and the index would coincide: instead of relating to a property like *being the American President*, the index (Clinton) would exhaust the intended interpretation

of the pronoun.

As suggested in chapter 2 section 2.2.2, Nunberg's argument against the conversational status of descriptive interpretations can be seen as one of the main motivations for the 'formal' account just mentioned. Thus, the treatment can be regarded as a function that takes a function from the *deictic component* and *context* to index as input and outputs the intended interpretation. But this formal approach does not answer a key question. What would determine whether the intended interpretation is deictic or descriptive? In earlier writings, Nunberg (1979, 160) borrows the psychologist's use of 'cue validity' to describe what determines the relevant reading. In some cases, the range of possible interpretations indicated by an index is exhausted by itself (e.g. a 'direct' cue, in the sense of chapter 3). These are the deictic uses. In some other cases, the index can be cue to an interpretation that is related to it (e.g. an 'indirect' cue, in the sense of chapter 3). These are the 'descriptive' uses. However, it is unclear whether cue validity should be understood in terms of the spread of activation amongst competing mental representations (Recanati 1993, to be reviewed shortly) or in terms of cues to specific intentions of the speaker (Sperber and Wilson 1995, reviewed in chapter 2). Nunberg (1993) mentions that speaker intentions would be necessary for fixing the relevant values of index and intended interpretation, which suggests that he has in mind something closer to the latter understanding. Thus, the take home message seems to be this: although Nunberg argued against implicature accounts, it seems that the appeal to the intentions of the speaker is essential for the explanation of deictic and descriptive readings.

The account generalises to other cases of pronouns as follows. The first person plural works pretty much along the same lines as the first person singular above, the only difference being that the classificatory component includes a plural feature that makes the intended interpretation to be about a group of entities. In the case of (1)b, O'Connor's use of 'we' determines her as the index. The intended interpretation is *the American Supreme Court Justices*, which is instantiated by the index, hence preserving the constraints imposed by the classificatory and relational components.

Differently from first person pronouns, the deictic component of second person pronouns corresponds to something like *the addressee*, which determines the specific person or group addressed as the index. The difference between singular and plural 'you' works along the same lines as the one between 'I' and 'we' above. In the case of (1)d, John's use of 'you' makes his friend the index. The intended interpretation is *the person calling*, which is instantiated by him (John's friend is making the call) and preserves the requirements imposed by the relational and classificatory components (the resulting interpretation is animate and is instantiated by the index).

Third person (singular and plural) pronouns - *non-participant* terms, according to Nunberg (1993) - fall a bit out of the picture sketched so far. Their deictic component does not encode enough information to direct the audience to any set of potential indices. Some uses of these expressions would require a demonstration to do the job (called 'deictics', Nunberg 1993, 23,36). Other uses can dispense with demonstrations, because an object is salient in the context in such a way that communicators take it as the pronoun's index (called 'contextuals', Nunberg 1993, 23,36). It is unclear here whether deictics and contextuals encode different deictic components, or whether they could be unified by encoding the same requirement of 'object salience' (demonstrations play the role of making a particular object salient, perhaps along the lines sketched in chapter 3) [1]. In the case of (1)a, the gesture determines the footprint as index, which is suitably related to the intended interpretation *the person who made the print*. Non-participant terms do not impose that the interpretation must include or be instantiated by the index. The classificatory component - male and animacy features - is preserved in the interpretation.

The *desiderata* in section 4.2 can be explained as follows. The deictic component encoded by the pronoun and the speaker intention determines the index: Clinton for (1)c. This is the stage at which *Identification Dependency* would be met, but it is unclear whether the account predicts that Clinton should be identified under a concept, like POLITICIAN, or as a bare individual, along the lines described in chapter 3 section 3.3. Since the latter would preclude the account from meeting *Identification Dependency*, I assume that the contextually salient object must somehow be identified under a concept and take this interpretative property to be properly accounted for at this stage in the interpretation. The relational component encoded by the pronoun and some form of pragmatic process, then, select an appropriate interpretation which is constrained by the classificatory component (e.g. the pronoun's animacy features) and, in the case of participant terms (first and second person pronouns), by the relational component's requirement that the index must be included in or instantiate the intended interpretation. The 'cues' in (1)c together with such linguistic constraints select *the American President* as the intended interpretation. This meets *Descriptiveness*. In this proposal, the relation between the identified object (index) and descriptive interpretation is a *linguistically-mandated-intention-sensitive* process. This accounts for *Connection* on partially inferential grounds (speaker intentions are necessary to arrive at the descriptive interpretation)

---

[1] Following Powell (1998), I take it that the distinction between two types of demonstrative pronouns would introduce an ambiguity for non-participant terms, which we hardly see evidence for. Therefore, I just assume that Nunberg assumes that both kinds of use require some form of object saliency.

and partially by encoded information: the requirements imposed by each component. The latter aspect makes descriptive pronouns much more similar to indexicality (to properties, in this case) than implicature derivation[2]. But what about the extended properties of the data?

Based on the proposal's ability to predict that identification-based pronouns can receive descriptive interpretations fact and the ability of descriptions to co-vary with other expressions in discourse, we could assume that the interpretation of the pronoun in 'Every time there is a war, **he** [*pointing to the White House*] has tough choices to make' contributes the description *the American President*, which co-varies according to different war times. However, this is a hand-waving explanation and one would really need a mechanism that can account for the co-varying readings. Elbourne (2008) provides an implementation of Nunberg's proposal, based on three cornerstones: (i) pronouns are definite descriptions, (ii) descriptions presuppose the uniqueness of their satisfiers, (iii) predicates describe minimal situations: world parts that contain only one individual, who instantiates only that property. Elbourne's implementation gets the Co-variation property right, thus I assume that this *desideratum* can also met[3] (Breheny 1999 has a proposal that is also able to account for co-variation, but unlike Elbourne's is deeply pragmatic in nature).

As for the *Accessibility* property, things are more complicated. Since Nunberg does not account for how dependencies across sentence boundaries are established, it is hard to see how he would explain the data below. The same difficulty can be raised for Elbourne's (2008) implementation.

(4) a. *Said by someone:* If the Democrats had won the last few presidential elections, **the American Supreme Court Justices** might have been liberals. **They** would guard public interest better.

b. *Said by Supreme Court Justice O'Connor:* If the Democrats had won the last few presidential elections, **we** might have been liberals. ***They** [*said by same speaker*]/**They** [*said by different speaker*]/?**We** [*said by same speaker*] would guard public interest better.

---

[2]Note, however, that the argument that Nunberg gives against implicature accounts of descriptive pronouns is a non-sequitur (see the argument against the conversational status of descriptive interpretations in chapter 2 section 2.2.2), but there are other reasons for not treating (some) descriptive pronouns as implicatures (as suggested by the arguments against Grice's picture of non-literal meaning, from circularity and from intrusion, discussed in the same section).

[3]I omit a full assessment of the mechanisms by which Elbourne (2008) accounts for *Co-variation* because the proposal is subject to the same criticism that may be raised to Nunberg (1993). Therefore, I have opted to take the explanation of this *desideratum* for granted and present a common criticism to both proposals.

Problems for the proposal

I will present three potential problems for Nunberg's proposal: (i) one concerning its explanatory insufficiency, (ii) another to the effect that the components seem to be too restrictive in some cases, and, finally, (iii) issues related to the *desiderata* that the account must explain. Arguments based on (i) and (iii) seem to be along the right tracks, but the status of arguments based on (ii) seem to be less clear.

First, Nunberg's proposal faces the problem of being an *insufficient* explanation. He rightfully takes the data discussed throughout this thesis as a motivation for postulating richer linguistic structure; in his case one that involves deictic, classificatory and relational components. However, these components do not explain, by themselves, how the descriptive readings are delivered. Nunberg (1993) appeals to 'speaker intentions' (or 'cue validity' in Nunberg 1979) to first fix the *index* and then to provide the intended interpretation that satisfies such linguistic components. However, we saw in chapter 2 that there are different ways in which a theory of human communication can incorporate sensitivity to speaker intentions. For example, Gricean pragmatics and Relevance Theory both appeal to speaker intentions in their explanations, but ultimately make quite different predictions. Moreover, as we shall see soon (section 4.4), the notion of 'cue validity' can capture extra-linguistic information in a speaker-intention free manner (Recanati, 1993). This adds more possibilities when it comes to the question of how extra-linguistic information is recruited to augment or modify linguistic meaning; but without the specific details of how pragmatic options elaborate on the three encoded components, the proposal cannot be properly assessed.

The second objection against Nunberg's proposal can be traced to Powell (1998) and runs along the following lines. In many cases, the account seems to be too restrictive in at least two ways: one related to processing order, the other related to the satisfaction of constraints imposed by the components themselves. First, consider the case of third-person non-participant terms. They do not encode a deictic component to fix the interpretation and rather require some form of 'object saliency' to fix their indices. For example, an utterance of 'They haven't signed the permission form yet', whilst a school teacher points to a little girl (Nunberg, 1993), has the index, namely, the little girl, fixed by the demonstration and then some form of mental process takes us to the interpretation: *the parents*. However, there are no arguments for this sequential view of processing. Since third person pronouns lack a deictic component, it could be the case that the index does not play a semantic role (i.e. a linguistically mandated stage) in finding the final interpretation. It could

be the case that third person pronouns just need some salient interpretation and in the case above pointing to a girl is a good enough way of making her parents the right interpretation of 'they' (as argued for in chapter 3). Nunberg simply does not rule out this possibility. I believe that this objection is well supported.

The second respect in which Nunberg's proposal might be too restrictive concerns the interplay among the different components in the theory. Many relatively similar arguments based on this point have been raised in the literature. I will try my best to address each claim individually, in order to avoid drawing misleading conclusions.

To begin with, it could be argued that the constraints imposed by a given component do not hold for certain interpretations (Grimberg, 1996; Powell, 1998; Galery, 2008). Consider the gender feature of the pronoun, which constitutes the *classificatory* component. In many cases, it can be overridden. For example, when I point to an old good male friend of mine and say 'She is going to be late', I convey, in addition to the proposition *that person is going to be late*, information associated with a certain (usually but not necessarily negative) female stereotype. On the assumption that expressing the relevant proposition requires certain gender features of the pronoun to be ignored by the audience, how can we be certain that such feature deletions do not occur in descriptive uses of pronouns? If nothing can guarantee this, then it could be argued that Nunberg's components would not play the role they are supposed to and hence could be dispensed with. In Galery (2008), I raised this objection, but now I think it misses the mark. The mismatch between features in the case of 'she' above plays quite a different role from the one involved in descriptive pronouns. In the former case, the feature-deletion occurs so information (e.g. regarding a negative female stereotype) collateral to the expression of a *singular proposition* is expressed. In the descriptive case, the feature mismatch (e.g. between the animacy of 'he' and the demonstrated footprint) is one of the triggers for the search for a descriptive interpretation. The fact that pronominal features are somehow preserved in these case can also be attested by the anaphoric data that motivates *Accessibility*; re-stated in section 4.2 above, i.e. descriptive pronouns usually convey descriptions, but these cannot serve as antecedents to other pronouns unless the person feature of the pronoun is of the right kind.

A similar complaint has been raised against the relational component. Grimberg (1996), quoted by Powell (1998), suggests that certain examples, like the one below, show that 'we' does not impose the restriction that the index must be 'included' in the interpretation:

162

(5)   ...a group of conservative middle aged businessmen [is] sitting around
             a boardroom table as one of them says: 'in a couple of years we'll
             probably all be women.' (Nunberg, 1993, 14, fn. 20)

Nunberg himself agrees that the speaker does not include himself in the interpreta-
tion, for that is a women-only group. Similar remarks apply for O'Connor's use of
'we' in (1)b, as she intended to induce a group-denoting interpretation that does not
include her. The example in the quote above could be treated along various lines.
The first of them consists in dismissing any constraints imposed by the relational
component of 'we', if any such thing exists (Grimberg, 1996). The second is to
actually say that 'we' above refers to the actual group of board members, but the
predicate 'probably all be women' receives a metonymic interpretation, (Nunberg
1993, 14, fn. 20 and Nunberg 1995), such as, *to behave in a feminine way*.

    However, as Powell (1998) suggests, none of these hard-liners need to be taken.
Nunberg postulates a disjunctive definition for the constraint of the relational com-
ponent of 'we': the index (speaker) must be included by or *instantiate* the inter-
pretation. In this case, the intended interpretation is *the boardroom members*, in
Donnellan's (1966) attributive sense. It is true that the speaker does instantiate the
property picked out by the description in the actual world, but the utterance is eval-
uated against a circumstance in which it is impossible for the speaker to instantiate
such property (as all the boardroom members are women). Still, nothing forces us
to assume that instantiation here must hold also at the circumstance of evaluation.
The fact that the index instantiates the property in the context of utterance (situation
or actual world) ensures that the right descriptive interpretation is 'loaded up' into
the proposition expressed. The criticism simply does not hold.

    Another complaint against Nunberg's components has been voiced by Recanati
(1993, 311) and this time the culprit is the index. In order to make his point, Reca-
nati urges us to consider the following example.

(6)   *John utters*: We are in red brick.

According to him, the truth-conditions of the utterance above is *the house we live
in is in red brick*. He then argues that determining the index (John) is not *sufficient*
to convey the intended interpretation. To him, what does the job is identifying the
group of people who live in that specific house, which he treats as a *trigger* (in the
terminology of Fauconnier 1985) and raises the saliency of the house itself (the *tar-
get*, according to Fauconnier). If the relevant group of people is not identified, there
is no descriptive interpretation. Recanati's remarks are indirectly associated with a

previous problem: the interpretation - *the house* - does not satisfy the requirements imposed by the relational component: it neither includes nor is instantiated by the index. However, these are separate claims: something that has not been recognised as such in the literature.

Nunberg (1993) argues his way out of this example in the same way he dodged previous criticism, namely, by assuming that the predicate in this case (*to be in red brick*) is mapped to another predicate (say, *to be **living** in red brick*, see discussion in chapter 1 section 1.7). Whether this process is predicate metonymy or coercion[4] is unimportant at this point, but note that Nunberg's way out is more appealing in this case than the previous one. The proposition that *some people, including the speaker, are living in red brick* seems to capture what is intuitively conveyed by the utterance. Moreover, as Recanati (2004) later points out, there is no principled explanation for assuming that such forms of (pragmatic) alterations to truth-conditional content target the subject as opposed to the predicate. Both are equally viable options (see chapter 1 section 1.7). In conclusion, the objection by Recanati (1993) seems to miss the mark. So, the second family of objections against Nunberg's account does not offer any knock down argument.

The third and final type of argument against Nunberg's account concerns its explanatory adequacy with regards to the *desiderata* mentioned in chapter 1 and reprinted here in section 4.2. The objection is quite simple. Nunberg's account does not predict the fact that certain uses of descriptive pronouns cannot serve as antecedents to other pronouns in discourse unless personal features of descriptive antecedent and anaphoric pronoun are of the right kind. To be fair, Nunberg did not consider this phenomenon in his original writings, but if he wants to account for this fact, substantial extensions to the theory must be made. The implementation offered by Elbourne (2008) seems to suffer from the same vulnerability. If pronouns are descriptions and utterances describe minimal situations, it is unclear why further pronominal uses cannot describe the situations (i.e. 'extend' the minimal situation, in his terminology) previously described by identification-based descriptive uses in some contexts but not can do so in others.

In conclusion, I have raised two problems for Nunberg's account (and Elbourne's 2008 implementation of it): first, in order to predict the readings that it allegedly predicts, it must spell out in detail how perceptually available content can integrate with linguistically encoded content. Second, it does not account for all the interpretative properties of the data, more specifically *Accessibility*.

---

[4]Consider, for example, 'John began the book,' where the semantics of the verb 'coerces' a richer reading due to the fact that it requires eventualities as complements: *John began reading the book*.

4.4 Heavy-Handed pragmatics accounts

This family of theories aims to account for the data by extra-linguistic means. In this section, I will deal with Borg's (2002; 2004) semantic minimalism account, Recanati's (1993) synecdoche account, and the relevance-theoretic accounts of Bezuidenhout (1997) and Powell (2003). Like the previous section, I will present these positions and their respective challenges in turn.

4.4.1 Semantic Minimalism

The proposal

Before addressing how Emma Borg deals with the particular details of the descriptive uses of pronouns I have in mind, I must outline her general view on linguistic understanding and the grammar/pragmatics interface.

Echoing Fodor (1975), Borg (2004) puts forward a position according to which language understanding is a formal, systematic and algorithmic process. The position assumes a level of linguistic comprehension, the semantic level, which is recoverable entirely from the lexical type encoded by the words used. This stage in linguistic comprehension can be described by a syntactically driven semantic machine: semantic representations are computed solely on the basis of the formal features of the linguistic expressions and formally tractable aspects of context (i.e. 'narrow' context, in the sense of Bach 1994).

This computational view of language understanding is underpinned by the requirement that a semantic theory should only account for the (pragmatics-free) outputs of the language faculty. The position is quite similar to the one held by relevance theorists (Sperber and Wilson, 1995; Carston, 2002) in the sense that semantic content is equated with encoded information. However, differently from relevance theorists, Borg further assumes that the outputs of the decoding phase in comprehension are always fully propositional (truth-conditionally evaluable) [5].

In addition to RT, it is useful to compare this position with that of Grice (chapter 2 section 2.2). To begin with, Borg seems to be making claims about representations that are psychologically realised. This picture is very different from one aimed at solving certain philosophical problems, like Grice's. Moreover, the semantic objects that are the output of the linguistic faculty are more minimal than the Gricean

---

[5]Bemis and Pylkkänen (2011) recently published some evidence against such a strong modularist claim: the composition of simple phrases containing just a noun (e.g. 'boat') and an adjective (e.g. 'red'), which results in a semantic object below the level of truth-conditions, requires the use of brain areas usually associated with general reasoning.

what-is-said, for they do not need to correspond to the level of information that the speaker intended to convey. In short, *M-Intentionality* does not constrain *semantic* representations for Borg. Still, there are two Grice inspired principles that inform her view: a minimalist principle according to which the minimal outputs of the decoding phase correspond to truth-evaluable objects (akin to *Truth-Evaluability of What-is-Said*) and a linguistic determination principle, which states that such objects are free from pragmatically derived constituents (akin to *Linguistic Constraint on What-is-Said*). From this last principle, an interesting consequence follows. Given this strong modularist position, the (E)pistemic condition on singular thinking, mentioned in chapter 1 section 1.2, does not characterise semantic (i.e. linguistically determined) contents. That is, the idea that the audience must *know which* object a word-use designates in order to grasp the singular thought expressed by the relevant sentence does not constrain semantic content. Therefore, either there is no *singular* content at the semantic level or *singular* semantic content is determined solely in virtue of the lexical types uttered. Borg assumes that the latter is the case and thus postulates a strong isomorphism between referential linguistic expressions and singular concepts introduced by them.

As stated, the position faces obvious difficulties when it comes to words whose lexical entries engage contextually available information in order to express singular thoughts about specific aspects of the context in which they were uttered, such as, 'I', 'you', 'here' and 'now'. How then could context dependency be captured in a purely syntactic manner? Following Kaplan (1977, 1989) and Bach's (1994) notion of 'narrow' context, Borg models contexts as containing parameters for each aspect of reality that is necessary for the existence of an information-carrying sentence use. For example, each sentence is used by someone to address someone else somewhere at a time. Thus, if I say 'I'm Brazilian', it is necessary that Thiago Galery is the speaker of that token, otherwise it would not be the token that it is (and similarly for the other essential features of context). Along Kaplanian lines, a (narrow) context c contains parameters for speaker, addressee, time, place and whatever other metaphysically necessary aspects for sentence evaluation are necessary. The resulting picture is this, a purely semantic treatment of context-dependent linguistic meaning amounts to a formal description of *sentences relative to context*: the objects described by our semantic theory. As a use of 'I', say, requires that someone produced it, the essential connection between producer and token allows the word to load the right referent into the semantic content expressed by the relevant sentence (see Borg, 2004, 158-63,166-7). The contribution of the first person pronoun roughly corresponds to a singular concept (i.e. α; a constant in LOT format) that

designates *the speaker of that token* (see Kaplan, 1978; Perry, 2001), regardless of whether or not the speaker is able to identify (in a extra-linguistic sense) the relevant individual in the context.

This picture faces a problem when one considers what would happen if the essential properties of the utterance event fail to hold. For example, if there is no speaker, there is no utterance, so no proposition is expressed. If there is no audience, linguistic material could still be produced by the speaker, but it would not be understood by anyone and no information would be conveyed. Now, consider a situation in which both speaker and audience are present and the speaker utters 'that is red' whilst pointing to an empty space. The absence of a demonstrated object does not bear on the existence of something (a proposition) that is conveyed to the audience based on the token uttered by the speaker. According to Kaplan (1989), third person demonstrative and personal pronouns cannot be traced to any formal model of the context. In short, a demonstrated object is not a necessary metaphysical feature of every communicative situation. How then could the content of such expressions be derived solely in virtue of the token of a certain lexical type?

Borg's treatment of third person personal and demonstrative pronouns borrows heavily from Higginbotham's notion of conditionalised truth-conditions (henceforth, CTCs), which can be summarized as follows:

> If the speaker of 'this is red' refers with the utterance of 'this' therein to x and to nothing else, then that utterance is true if and only if x is red. (Higginbotham, 1994, 92-3)

It seems that knowledge of the truth-conditions of utterances containing demonstratives involves the knowledge of an antecedent condition securing that in order for a demonstrative pronoun to have the content that it has, the speaker must have referred to a *specific* object. In loose terms, because utterances containing third person personal and demonstrative pronouns convey conditional truth-conditions in virtue of their linguistic type, they contribute singular concepts (e.g. α) to the semantic content expressed by the utterances they occur in. Similarly to other pronouns and indexicals, the singular concepts introduced by third person pronouns behave rigidly, that is, their content remains constant even if the utterance is evaluated against different possible worlds or situations. Still, their content is quite independent from the communicative intentions of the speaker or the audience's ability to identify the object in question (see Borg, 2004, 165).

The position can be made clearer by the following scenario. Suppose someone, looking at another person through the window utters 'She is running'. The hearer

sits far away from the window, and thus cannot integrate the linguistic information of the pronoun with the perceptual content that the speaker has access to. Still, Borg argues, the hearer is able to retrieve a singular content, which refers to the woman the speaker sees, solely based on the conditional rule invoked by Higginbotham above[6]. At this point let us grant Borg's claim that a certain class of linguistic expressions can introduce singular (i.e. rigid) concepts via their syntactic form alone and ask: what are the consequences of this position for the interpretation of descriptive pronouns?

According to the core data of descriptive uses, the communication of descriptive truth-conditions (*Descriptiveness*) depends on the identification of a specific entity in the environment as belonging to a kind (*Identification Dependency*) which provides the conceptual building blocks for the communicated proposition (*Connection*). Since extra-linguistic identification processes cannot provide constituents for semantic content, descriptive uses of pronouns cannot contribute descriptive information at this level of information.

Borg's picture is rather intuitive at least for some cases. Reconsider the scenario above, where the hearer does not have a perceptual link to the woman the speaker looks at from the window, but is still able to recover the concept that denotes the woman based on the linguistic form of the pronoun. Similarly, it could be argued that in the situation in (1)a, where the speaker points to a footprint whilst uttering 'He must be a giant', the concept introduced by the pronoun refers to the footprint-maker, even though the hearer does not know *who* that is. The similarity between the two cases can be highlighted by the separation Borg (2002) draws between the *ostensive act* (the physical gesture) and the *demonstration* itself. In the first case, the ostensive act and the demonstration coincide: the woman pointed at is the object the pronoun's referent. In the footprint case, these notions do not coincide, that is, the speaker ostensively points to the footprint, but in fact demonstrates the footprint-maker, which the encoded meaning of the pronoun is sensitive to (in virtue of the coditionalised T-schema). If one concentrates on the more interesting notion of demonstration, the descriptive readings of many pronominal uses are illusory at Borg's semantic level. Thus, the alleged core properties of the data are a mischaracterisation of what goes on in these cases (i.e. instances where the pronoun behaves rigidly).

However, Borg acknowledges that, in some other cases, descriptive interpreta-

---

[6]Note that one could object to Higginbotham's use of conditionalised T- schemas by demanding an elucidation of the meaning of 'refers to' in the quote above. Does it mean 'succeeds to refer to'? If so, can success in reference be separated from any appeal to intentions? This is a natural and powerful objection.

tions may arise from some form of pragmatic reasoning. The example she has in mind is this.

(7)   *Pointing to the American Presidential podium*: Every intern wants to have
        an affair with him.

Borg (2002, 500) argues that the speaker could continue her utterance by saying 'Even if he is deceitful, like Nixon, or boring, like Carter...', suggesting that 'him' above contributes *the American President*. Given the tenets of the proposal, the descriptive interpretation would not be part of the semantic content of the utterance, but rather fall under its implicatures. The idea is this: (7) expresses a proposition about the actual American President, Barack Obama, without any form of epistemic identification of the referent (the concept α introduced by the pronoun refers to him rigidly via the Higginbotham style T-schema) and then general reasoning processes use the semantic content in the generation of a descriptive interpretation at the level of speaker-meant implicatures. Borg (2002, 502-3) mentions that, in this inference, 'though we would view the speaker as having made some kind of mistake in his choice of sentence, we can still follow what he is trying to say'. Schematically, the core properties of the data would be explained along Gricean lines (see chapter 2 2.2), *Identification Dependency* would be captured by the singular proposition at the semantic level (or at a post semantic level at which visual information would be integrated with semantic content), *Descriptiveness* would be captured by the implicated descriptive proposition, and the inferential relation between the two propositions would explain how the building blocks of the descriptive proposition are generated from a more basic level of interpretation, thus accounting for *Connection*.

Before assessing this position, it is also important to consider two arguments *against* the idea that pronouns convey descriptive information at the level of semantic content (a position Borg 2002 calls 'Descriptivism'), which motivate the proposal sketched here.

The first argument, which I call the 'instability of content argument' (i.e. 'the variety of descriptions argument', in Galery 2008, 171), uses, as a premise, the claim that various descriptions can be associated with a given pronominal use. Borg (2002, 494) considers the following example:

(8)   *Someone points to a chair which a person sat on recently*: 'She is gone'.

Borg points out that there are various descriptions that the audience could associate with what the pronominal use communicates in this context, such as: *the previous*

occupant of that chair, *the person sitting there five minutes ago* or *the female owner of this piece of furniture*, for example. Since there are multiple semantically non-equivalent descriptions that are likely candidates of the pronoun's truth-conditional contribution, none of them could capture a stable level of meaning that is expressed by a token of the sentence above. Moreover, none of these descriptions needs to be intended by the speaker in order for the pronoun to express the content that it does. Here, Borg (2002) might be arguing for any of three positions, the first concerning content shareability; the second concerning semantic compositionality, the third concerning the status of such descriptive uses *vis-a-vis* a *semantic rule* that determines content.

On the first interpretation of the argument, if the speaker associates *the G* with the pronominal use and the hearer associates *the H* with the same word token, the truth-conditional content entertained by each may differ and it cannot be shared between communicators (see Borg, 2002, 494-5). There are several lines of response available to this. First, this may not be an undesirable feature of the system at all. Cases of miscommunication that emerge because a very basic level of information is not shared are common and a theory must capture the conditions under which they happen. Secondly, even though the descriptive readings are a significant departure from what pronouns literally encode, they are not 'free' from pragmatic factors. Thus, pragmatic principles may guide communicators in zooming into a common identification-based descriptive content. Borg's worries seem to be unwarranted in this case.

On the second interpretation of the argument, Borg could press on and abandon descriptivism on the grounds of compositionality. Since a token of 'he' can contribute different descriptions in different contexts, there is no possible explanation of how an utterance's truth-conditions are composed of the meaning of its parts. As a response, one might argue that compositionality does not need to hold at any pragmatics-free level such as Borg would like to. Rather, one might devise a compositional explanation which encompasses the relation between encoded content and the pragmatic processes that shape extra-linguistic information into truth-conditional contributions. In this sense, the very process of semantic composition requires adjustments on the lexical meaning of words uttered. Such adjustments may involve top-down pragmatic processes triggered by the bottom-up composition and/or highly salient information in the conversational setting (for recent arguments in this direction, see Recanati 2009; for the idea that a compositional explanation may hold at the level of encoded information, see Powell 2001; Carston 2002, 70-74). *Dynamic Syntax* (Kempson et al., 2001; Cann et al., 2005) is an example

of a fully compositional grammar model that incorporates pragmatically available contents into the computed representations. My own proposal for the descriptive pronominal data, sketched in the next chapter, draws on this framework. Thus, this instance of the argument also does not seem to be sound.

On the third interpretation of the argument, Borg (2002, 508-9) expresses some worries about perceptually-based descriptive information being treated as *semantic* information which begin by observing a distinction *between* cases of perceptually-based singular referring pronouns and perceptually-based descriptive pronouns. For example, in some cases, pointing to a person's shadow or the door she just slammed shut while uttering a pronoun may count as instances of singular reference to that person[7]. In some other cases, uttering the same pronominal expressions (accompanied by a similar ostensive act) could result in descriptive interpretations (e.g. *the person who just slammed the door shut*, assuming the audience does not know who did the door-slamming). The argument then introduced the following premise: *if the semantic content of such uses is descriptive, there must be a clear division between cases where a pronoun is interpreted referentially and cases where it is interpreted descriptively* (on the assumption that semantic content in order to be shareable, must be uniquely available for communicators).

One would have to assume then that referentiality and descriptiveness in interpretation emerge as a result of different linguistic rules, which in turn classify different types of situation which the pronoun can describe. This would carry direct consequences for learnability: if one has the means to produce and understand singular and descriptive uses of pronouns, then one would have mastered a method for distinguishing all the situations where a pronoun receives a referential interpretation versus those where it receives a descriptive one. However, this does not seem to be the case, there is a plurality of factors that may be involved in the availability of referential or descriptive interpretations, as for example, the lack of knowledge about the person who was just sitting on the chair pointed at by the speaker. Without a clear cut distinction between the circumstances under which a pronoun is interpreted descriptively and the conditions under which it is interpreted referentially, descriptive interpretations cannot emerge at the semantic level.

Although I quite agree with Borg's position that the descriptive interpretations cannot be traced back entirely to semantics, understood as encoded information, there are problems with the conditional statement used as a premise in this third interpretation of the argument (in italics above). First, referential and descriptive in-

---

[7]This in fact is one of the reasons why *Descriptiveness* is one of the essential properties of the core data.

terpretations of the pronoun are assumed to be explicable only in terms of different linguistic rules. Besides positing unwarranted ambiguity, this misses the possibility of a single underdetermined level of linguistic meaning (to be spelled out in chapter 5) giving rise to both types of interpretation. This possibility breaks the connection between antecedent and consequent: one could be able to understand and produce both types of pronominal use, without having a clear cut method for distinguishing all situations appropriately described by referential pronouns from those described by descriptive pronouns. After all, whether an identification-based pronominal use is referential or descriptive depends on a variety of factors (see chapter 1 section 1.2). Therefore, although Borg's conclusion, namely, that descriptivism is a semantic thesis (i.e. a thesis about what words encode) seems right, the argument does not show why pronouns cannot have descriptive interpretations at a level of content that is minimally truth-evaluable but pragmatically rich.

In summary, all the possible takes on *instability of content argument* can be resisted.

Borg's second argument, which I call 'the modal and scope argument', builds on the analysis of the pronoun's modal and scope behaviour. The idea is that if descriptivism is right, the descriptive readings conveyed by identification-based pronouns would give rise to scope ambiguities, when interacting with other operators (negation, quantifiers, propositional attitude verbs, etc...). As an illustration, consider the sentence below:

(9)   *Someone points to a parked car with a ticket on its windscreen and says*: 'He is **not** going to be sorry'.

On the assumption that the pronoun's truth-conditional contribution is a definite description, and given a Russellian (Russell, 1905) treatment of descriptions as complex quantifiers ('The F is G' is equivalent to $\exists x \ (F(x) \ \wedge \ \forall y \ (F(y) \ \rightarrow \ x = y) \ \wedge \ G(x)))$, the sentence above would have two scope readings: (i) 'there is a unique person who owns the car and that person is not going to be sorry', and (ii) 'there is not a person who owns the car and is going to be sorry'. However, no ambiguities arise when we interpret the negated sentence: only the first reading is available. A similar point can be based on the modal behaviour of the pronoun. Consider a non negated version of the sentence above embedded as the antecedent of a conditional statement, producing 'Even if parking tickets were cheap, he [*pointing to the car*] would be sorry'. Borg argues that the interpretation of 'he' is tied to the actual world or situation, that is, it refers to the actual owner of the demonstrated car. We do not have a potentially different car owner for each of the worlds/situations

against which the utterance is evaluated. In other words, the pronoun - even in 'descriptive' uses, behaves rigidly (see the test by Kripke 1972 in chapter 1 section 1.3). Borg (2002, 497) concludes that since scope interactions target the level of truth-conditional content and the predicted ambiguities are not borne out, the descriptive readings should not be captured by a semantic theory.

With regards to the modal behaviour of the pronoun, Borg seems to make an unwarranted move. As we have seen in chapter 1 section 1.3, Kripke's (1972) rigidity tests crucially relies on communicators' intuitions about the type of worlds/situations that are considered for the purposes of evaluating an utterance containing a modal operator. Intuitions about the truth-conditions cannot be used as *evidence* for ascribing an interpretative property - in this case, rigidity - to an expression type - in this case, pronouns - because they target an already pragmatically rich level of content (i.e. Gricean what-is-said, RT's explicature). The type of semantics/pragmatics interface assumed by Borg undermines the very evidence for assuming that referential expressions (pronouns, names, etc.) introduce singular concepts. Alternatively, Kripke's test could be rejected as a tool for determining whether certain expressions are rigid and pronouns could be classified as rigid designators. This move, however, throws out the baby with the bath water. There is no evidence for the alleged referentiality of pronouns as linguistic types and, given the plethora of descriptive interpretations, as, for example, bound, donkey and descriptive, stipulating an isomorphism between pronouns and singular concepts is unjustified. Note that if we let communicators' intuitions play a role in theory building, then we find cases, like O'Connor's utterance of 'We might have been liberals', where the pronoun selects different people in different worlds, i.e. a non-rigid designator.

In conclusion, both of Borg's arguments against descriptivism do not rule out the possibility of descriptive uses of pronouns contributing to a level of content that is pragmatic in nature but minimally truth-evaluable. I will now turn to some problems that the theory might face.

Problems for Borg's views on descriptive pronouns

Borg's theory aims at the right direction. Building psychologically realistic theories about linguistic and extra-linguistic content is a virtue. Recognising the crucial role of pragmatics in fixing the intuitive truth-conditions of utterances containing pronouns is another. However, there are two contentious aspects of the proposal just presented. The first concerns the status of the semantic representations as completely truth-evaluable objects that are entirely linguistic determined. The debate around this issue has been a hot topic for the past two decades. I will refrain from

entering into the details of it and focus only on aspects that concern the topic of this thesis. The second problem concerns how would a pragmatic explanation of the *desiderata* (section 4.2) would go *given the isomorphism between referential linguistic expressions and singular concepts*.

As detailed above, Borg has different proposals for descriptive pronouns. For some cases, such as when the speaker points to a footprint and says 'He must be a giant', the alleged descriptive reading is illusory. The idea is that although the speaker points at the footprint, the pronoun rigidly refers to the footprint-maker (the *demonstratum*) via Higginbotham's conditionalised T-schema. For some other cases, like the scenario where the speaker points to the American Presidential podium and says 'Every intern wants to have an affair with him', the pronoun contributes *the American President* to the utterance's implicatures. Let us assess each of these strategies in turn.

The first strategy faces a serious difficulty, namely, the conflict between the idea that semantic content is determined entirely by linguistic means and that the idea that the context-dependent contribution of third person pronouns and demonstratives (i.e. the demonstrated referent as opposed to the object ostensively pointed at) is determined by Higginbotham's conditionalised T-schema. The schema generates a certain truth-conditional contribution, if a condition holds, namely, that the speaker succeeds in referring to an object by the pronoun. Given that the relevant truth-conditional contribution (e.g. the demonstratum) is indirectly available (in the sense of chapter 3) in identification-dependent descriptive uses, it seems that whether or not the right truth-conditional contribution is loaded into the semantic content expressed by the sentence is not something that can be determined entirely by the lexical information of the words used (relative to formal aspects of the context). After all, as mentioned in chapter 3, there are descriptive uses that are infelicitous. As the antecedent of Higginbotham's schema makes clear, the success in establishing the right truth-conditional contribution seems to be *speaker* success. Alternatively, if we want to make the determination of linguistically conveyed content relative to certain conditions, these ought to be conditions on the existence of the utterance token itself (i.e. the formal aspects of context), or conditions imposed by the language faculty for the decoding of the linguistic material. Therefore, if this strategy is to be pursued, one of two things must give in: either we abandon the idea that semantic content is determined by linguistic form alone and keep the idea that one can refer to an object by ostensively pointing at another one, or we keep the linguistic determination principle, but sacrifice a sensible way to distinguish the ostensive pointing of an object from the referent actually demonstrated. Any of these

options would threaten Borg's strategy for dismissing the descriptive data.

The second strategy, namely, accounting for cases that do impart descriptive readings as instances of implicatures, also faces serious difficulties. According to chapter 2 section 2.3.3, that may be so for some cases, specially, uses of the first person pronoun, but ultimately an implicature-based account for all instances of descriptive uses seems to fail because the proposition that is necessary to violate a conversational maxim (e.g. *the footprint must be a gigantic person*) cannot be neither meant by the speaker nor determined solely by the meaning of the word alone, and hence the search for a descriptive proposition that preserves such maxims at the implicature level cannot be triggered. Alternative notions, such as making-as-if-to-say, do not carry enough speaker commitment to violate a maxim and thus the relevant implicatures cannot be generated. As a final possibility, Borg could try to marry her semantic minimalism with relevance theoretic pragmatics, so that the semantic content expressed by an utterance of 'He must be a giant' whilst pointing to the footprint, expresses *the salient male in this context must be a giant* at the semantic level and implicates *the footprint maker must be a giant* at the implicature level. This possibility, however, fails to account for the extended properties of the data. Consider again the co-variation and the asymmetries present in the examples below:

(10)   *Andy*: Well...**Every time there is a war**, **he** [*pointing to the White House*] has tough choices to make.

(11)   a.   *Said by someone:* If the Democrats had won the last few presidential elections, **the American Supreme Court Justices** might have been liberals. **They** would guard public interest better.

         b.   *Said by Supreme Court O'Connor:* If the Democrats had won the last few presidential elections, **we** might have been liberals. ***They** [*said by same speaker*]/**They** [*said by different speaker*]/?**We** [*said by same speaker*] would guard public interest better.

It seems that the co-variation between the description *the American President* and the adverbial 'every time there is a war' cannot be explained by implicature-based account. The descriptive reading falls within the scope of the adverb of quantification (over times), which motivates an explicature based treatment for some uses of the third person pronoun. In short, the isomorphism between referential expression types and singular concepts they introduce precludes any explanation of the *Co-variation* property, which are also present in bound and donkey uses (reviewed in

175

chapter 1). Moreover, the *Accessibility* pattern of antecedents in (11) above cannot be explained by variations in context other than shifts of perspective, suggesting that a simple pragmatic explanation may be insufficient. However, the division of labour sketched by Borg makes semantics a matter of *string* decoding, which precludes any explanation of this interpretative property that is not *entirely* pragmatic in nature. In sum, Borg's account is partial at best, for it does not explain the extended properties of the descriptive data and also cannot explain pronominal binding or anaphora, which raises serious difficulties for a unified treatment of pronominal expressions.

The foregoing observations suggest that *there is no evidence for supposing that the outputs of decoding are minimally truth-evaluable objects, nor for supposing that the alleged referential expressions, if there are any, introduce singular concepts via linguistic form alone* [8]. The decoupling of semantic content from the intuitive-truth-conditions expressed by an utterance undermines two of the underlying tenets of the proposal, namely, the claim that semantic representations are truth-evaluable objects, given that considerations about truth-evaluability and/or assertability impinge on the level at which intuitive content emerges, and also the claim that pronouns are rigid designators, given that the rigidity evidence for deictic uses (see chapter 1 section 1.3) depends on intuitions about the modal profile of the contents induced. In conclusion, each of the possible explanatory routes of the semantics/pragmatics interface proposed by Borg (2002, 2004) faces shortcomings. I shall now move to other heavy-handed pragmatic possibilities.

4.4.2 Truth-Conditional Pragmatics

Truth-Conditional Pragmatics comprises a family of theories that diverge from Grice's notion of what-is-said as an attempt to resolve tensions amongst the theses he held. In chapter 2, we saw that Grice held that (i) what-is-said corresponds to the linguistic meaning of the words uttered (*Linguistic Constraint on What-is-Said*), (ii) constitutes the minimal truth-evaluable object expressed by an utterance (*Minimal Truth-Evaluability of What-is-Said*), and (iii) captures the proposition the speaker intends to assert by a certain utterance (*M-Intentionality*). While some versions of semantic minimalism (e.g. Borg 2004, above) resolve the tension here by rejecting (iii) and slightly modifying the other tenets, TCP departs from this picture

---

[8]It could be suggested that a Perry (2001) style reflexive proposition or some form of existential closure over the linguistic form of the pronoun could capture a minimally truth-evaluable level of content (e.g. *the/a male is hungry* for an utterance of 'he is hungry'). Although this move is possible, it could characterise the meaning of pronouns as *descriptive*, hence threatening the isomorphism between referential expressions and singular concepts sought by Borg.

by reviewing the idea that the same level of information is solely linguistically determined and minimally truth-evaluable (claims (i) and (ii) above). It allows for pragmatic constituents to enter at least two levels of speaker meaning: the level of the intuitive proposition(s) expressed by an utterance (the explicature in Sperber and Wilson 1995, or the proposition generated by primary pragmatic processes, in Recanati 1993), and the level of the utterance's implicatures.

In the remainder of this subsection, I will concentrate on two TCP approaches that differ from each other with regards to the mechanisms that select pragmatic elements as constituents of the proposition expressed by an utterance. Such mechanisms can be association-based (Recanati, 1993, 2004) or fully inferential as in relevant theoretic accounts (Bezuidenhout, 1997; Powell, 1998, 2003).

Mixed association-based TCP (Recanati, 1993)

In the early 90s, Recanati (1993) attempted to solve a variety of issues in the philosophy of mind, language and pragmatics in a unified and empirically oriented framework. The main challenge was to provide a thought/language interface able to explain (i) Fregean insights about the cognitive value of certain natural language expressions or thought constituents, (ii) Kripke's insights about the scope and modal behaviour of a class of expressions, and (iii) the variety of meaning adjustments brought to light by people working in pragmatics (Grice, 1957, 1975) and speech act theory (Austin, 1962; Searle, 1969) (illustrated by the non-literal uses in chapter 2).

In order to account for (i), Recanati assumes a naturalistic version of Frege's notion of sense, that is, objects are not the kinds of things that are constituents of propositions; mental representations (of a certain sort) are. As we shall see shortly, these mental representations can be complex in nature and can be regarded as dossiers of information (Grice, 1969) or *mental files*. In order to account for (ii), he assumes that referential expressions, as linguistic types, encode a REF feature, which constrains the semantic value of the expression to be *singular* in nature, hence distinguishing this class of linguistic types from that of quantificational expressions (including definite descriptions). In order to account for (iii), he assumes that processes other than the bottom up decoding of linguistic information can participate in the building of the proposition expressed by an utterance as well as the derivation of the associated implicatures. However, the former, known as *primary pragmatic processes*, result from the spread of activation among certain mental representations by an association-based mechanism, while the latter, known as *secondary pragmatic processes*, involve general reasoning mechanisms that can

infer speaker-meant implicatures from the proposition expressed as a result from the contribution of primary processes. It is important to note that this duality of mechanisms does not exist in Relevance Theory.

Recanati (1993) accounts for the interpretative properties of descriptive pronouns by treating them as a special case of deictic uses of pronouns that undergo some sort of non-literal interpretation. Thus, some remarks about his views on deixis or indexicality are needed. The explanation of these cases begins by observing that deictic expressions or thoughts crucially involve a *non-descriptive* form of contextual dependency or token-reflexivity (Reichenbach, 1947). Consider the linguistic case. In a situation where John utters 'I'm hungry', 'I' refers to John in virtue of the fact that John is the utterer of that word token. In other words, the meaning of 'I' cannot be equated with the non-indexical description *the speaker*, because the connection between the specific speaker and the specific word spoken is lost. The description *the speaker of that token* successfully captures this connection, but note that a deictic element is re-introduced, namely, *that token*.

Recanati captures the non-descriptive nature of deictic or indexical context dependency at two levels. At the level of thought, the speaker or agent exploits certain *epistemic relations* that she has to aspects of the environment just by occupying it.

> ...by watching an object or by holding it in my hand, I gain (visual or tactile) information about that object; by standing in a certain place, I gain information about what is going on at that place; by being a certain person, with a certain body, I gain information about that person and that body (Recanati, 1993, 122)

.

Thus, what makes the thought *I'm hungry* be about John is the fact that John is the person entertaining it (similar remarks hold about the time and place the thought is about for example). This is subsumed by a thought mechanism - a psychological mode of presentation in the terminology of Recanati (1993), following Bach (1987) - capable of exploiting the special epistemic relation that one has to oneself, that is, information about oneself is gained from the 'inside' as it were (an egocentric relation). Another kind of epistemic relation is that between a thinker and a perceived object which possibly engages mechanisms described in chapter 3 (a demonstrative relation). These various types of epistemic relations have something in common: an instantiation of a particular type of relation in a context determines *a body of information* about the object the subject is related to[9]. The referent of John's *I'm hungry*

---

[9]Recanati mentions that the perceptual (demonstrative) relation may start at the non-conceptual

thought would be a *mental file* on John himself (in virtue of the 'inner' relation that John has to himself; files can also be less occasional in nature, like the *encyclopedic entry* on dogs that was established in virtue of many perceptual episodes regarding dogs). The specific token of the file (or information bundle) is individuated by the object it is about, but it is also associated with a body of information pertaining to the object, such as: *John has not eaten for 8 hours, John likes fish, John was born on the 8th of July*, etc.

At the level of language, Recanati holds that context-dependent expressions, such as pronouns, encode instructions that exploit these epistemic relations. This is done by two pieces of encoded information: one that tells which contextual parameter the expression depends on, in the case of 'I' this corresponds to 'the speaker' (i.e. *the linguist mode of presentation*; something like Kaplanian *character* or Nunberg's *deictic component*); another that tells the hearer to find out a mental representation about a *specific* entity. Recanati calls this last piece of encoded information the 'REF' (referential) feature. Suppose Jane is addressed by John's utterance of 'I'm hungry'. She is related to a variety of objects in the environment: herself (inevitably), John, the table he works on, and a book on that table, say. Allegedly, the epistemic relations between Jane and these objects would make mental files about them available in her mind. Furthermore, when John starts speaking, the piece of information ... *is speaking* gets added to Jane's file on him. The fact that this piece of information matches the information encoded by 'I' spreads activation from the decoded word to the mental file on John. Furthermore, the REF feature makes the file itself (its *label* or *address* in the terminology of Sperber and Wilson 1995), not information associated with it, the likely candidate for the semantic value of the word. This is an instance of the primary pragmatic process of *saturation*.

It is worth mentioning the similarities between the theoretical assumptions of Borg (2004) and Recanati (1993). Like Borg, Recanati aims to capture Kripke's (1972) intuitions about the modal and scope behaviour of expressions at the level of expression types, but while she assumes a strong lexical-type concept type isomorphism (as discussed in the previous subsection), he assumes a lexical feature that interacts with the spread of activation amongst semantic representations. This marks an important difference between the two positions. While the former is a syntactic based notion of singular thought, the latter bases the notion of singular thinking on extra-linguistic relations between communicators and environment that

---

level, possibly along the lines of Pylyshyn (2007), but surfaces at a level where certain combinatoric properties characteristic of conceptual representations (Evans, 1982) hold. See chapter 3. This model can also be extended to other types of epistemic relations.

can be exploited by special linguistic devices (the ones encoding REF). Recanati, echoing Bach (1987), assumes that the reference of deictic expressions in language and thought are determined *relationally*, as opposed to the satisfactional manner in which descriptions and other quantifiers make their truth-conditional contribution.

After this brief introduction to Recanati's basic framework, we can proceed to an assessment of his account of the interpretative properties of descriptive pronouns. Let us reconsider the famous footprint example in (1)a, reprinted here as:

(12)   Someone pointing at a huge footprint in the sand says: 'He must be a giant' (from Schiffer 1981, 49).

The proposal aims to explain descriptive pronouns by an additional primary pragmatic process. In the situation above, the perceptual experience of seeing the footprint creates a file on it that enables information about the object to be stored. The file is individuated entirely by the relation between the person who sees the object and the object itself. No conceptual representation or categorisation is involved (this relation can be regarded along the lines of bare object representation in chapter 3 section 3.3). Further cognitive processes add the information ... *is a footprint* into the file. At this point *Identification Dependency* is met. Now, the REF feature spreads activation to the file as a whole and selects it as the likely candidate for interpretation, but that cannot be the semantic value of the pronoun, because it would make the overall interpretation of the utterance nonsensical: a footprint cannot be a gigantic person. In addition to the whole file, other kinds of mental representation compete for being the semantic value of the pronoun, including ... *someone made this print*, which is written under the file on the footprint. At this point, the mechanisms of activation-spread select this bit of descriptive information for integration, the REF feature is suppressed, and the pragmatic process of *synecdochic* (part-whole) transfer kicks in. In summary, a description written in the file replaces the file itself for semantic purposes and the truth-conditions of the utterance above unpack as *the person whose foot made the print must be a giant*. At this point *Descriptiveness* is met. Finally, *Connection* is captured in terms of the association-based mechanisms that spreads activation from the file to information contained within it.

What about the other interpretative properties of the data? Since Recanati's account predicts that full-fledged descriptions can be the semantic values of pronouns despite the fact that they encode REF, *Co-variation* could be explained, because utterances containing pronouns can be truth-conditionally equivalent to utterances containing the corresponding definite descriptions, and the latter can covary with

other expressions in discourse. Although we would ultimately want a description of how co-varying interpretations are established, let us take for granted that they can be predicted by the proposal. The problem is that, under this assumption, the very same mechanism that captures *Co-variation* makes the contrast between (a) and (b), below, impossible to account for:

(13)   a.   *Said by someone:* If the Democrats had won the last few presidential elections, **the American Supreme Court Justices** might have been liberals. **They** would guard public interest better.

        b.   *Said by Supreme Court Justice O'Connor:* If the Democrats had won the last few presidential elections, **we** might have been liberals. ****They** [*said by same speaker*]/**They** [*said by different speaker*]/?**We** [*said by same speaker*] would guard public interest better.

It seems that the proposal does not account for *Accessibility*, that is, for why the descriptive interpretation of 'we' in (b) which proceeds through a singular mental file on O'Connor to the descriptive information '... is a member of the American Supreme Court' and selects this piece of information as the semantic value of the pronoun on the basis of the suppression of its REF feature cannot be the antecedent of a subsequent use of a third person pronoun.

Problems for mixed association-based TCP (Recanati 1993)

In this section, I shall point out four issues that arise for Recanati's proposal.

First, *the account is incomplete*. As my last comments in the previous section indicate, Recanati's proposal cannot explain, at least prima facie, the *Accessibility* property of descriptive pronouns, and hence cannot account for all the *desiderata*.

Second, *the account does not easily extends to other pronominal uses*. As reviewed in chapter 1, pronouns induce other forms of descriptive readings, more specifically bound and donkey interpretations. For example, how, in utterances like '**Every farmer** thinks **he** is a genius' and 'Every farmer who owns **a donkey** beats **it**', could the relevant dependencies (in boldface) be established give the presence of REF? It seems that they cannot. To be fair, Recanati (1993) warns the reader that his book focuses on singular reference; hence excluding bound and plural cases, but the conclusion seems to be that a unified explanation of the variety of pronominal uses seems precluded or at least left wanting.

At this point one could try to cash the co-varying readings above in terms of mental files and/or by the optional suppression of the REF feature, but I would

rather not. The reason is simple: REF is neither sufficient nor necessary to explain the data raised by deictic and descriptive uses of pronouns (see Galery, 2008, 180-3). This point can be developed along the following lines.

First, the mere existence of identification-based descriptive interpretations of pronouns suggests that REF is not *sufficient* to secure a singular interpretation of the pronoun. Such sufficiency was met in Borg (2004) by the stipulated isomorphism between referential expression types and singular concepts. Since Recanati allows pragmatic processes to modulate encoded information, referentially marked expression types can receive descriptive interpretations depending on the context they are uttered in.

Second, that REF is not *necessary* for establishing singular interpretations can be easily illustrated by Recanati's explanation of *referential uses of definite descriptions*. Let us suppose now that John and I are watching the trial of Smith's murderer. Both of us perceive the defendant from where we sit, and hence each of us has a perceptually based mental file on him. As the judge sentences the defendant as guilty, the description 'Smith's murderer' is added to our dossiers on the man, which already contained plenty of other descriptions, such as .... *is nervous, ...hired a bad lawyer*, etc. After the Court section, John says 'Smith's murderer is insane' to me. Recanati (1993, ch. 15) explains the fact that I could identify the specific man in the courtroom by the synecdochic relation between the description in the utterance and information that was just added to the mental file on the man (i.e. ... *murdered Smith*). The direction of activation flow is the converse from that of descriptive pronouns. In these pronominal uses, activation spreads from the singular file itself to a description in it. In the case of definite descriptions, activation flows from descriptive information in the file, which matches the descriptive information encoded by the words, to the file itself. Thus, the description 'Smith's murderer' does not figure in the proposition expressed by this referential use of the description and a singular thought about the demonstratively identified individual is communicated. This process does not involve appeal to REF features: the singular interpretation is determined via the matching of descriptive information and flow of activation to the file as a whole. So why suppose REF is encoded by pronouns in the first place, specially given that it makes a unified explanation of the variety of pronominal uses more difficult[10]?

---

[10]Much of Recanati's (1993) project aims to connect the linguistic meaning of pronouns to the context of utterance in a *special and direct way*. If the observations here and in Galery (2008, 181-2) are along the right lines, the same (mental) mechanism of synecdoche is available for definite descriptions and the linguistic meaning of pronouns would not connect to the context in a special way any more. In order to circumvent this consequence, Recanati (1993, 288-292) further characterises the linguistic meaning of pronouns: (i) as contributing to conditions on the context of utterance (as

Third, *the mixed association-based account cannot explain cases where pronominal reference crucially depends on the intentions of the speaker*. Carston (2007) raised the following difficulty for the solely association-driven nature of Recanati's *primary pragmatic processes*. She urges us to consider a scenario where she talks to one of her students, Sarah, who utters the sentence below:

(14)   Neil has broken his leg.

In this hypothetical scenario, Carston knows only two people named 'Neil': her young son and her colleague at the UCL linguistics department. Furthermore, suppose she is deeply worried about her son (represented by a NEIL1 concept) but not about her colleague (represented by a NEIL2 concept). Given that the process of delivering the proposition expressed by an utterance does not involve considerations about speaker intentions, the account predicts that the resulting interpretation would be something like NEIL1 HAS BROKEN HIS LEG, which is overall coherent and carries many psychological consequences (i.e. cognitive effects) for the addressee. Carston (2007, 28) concludes:

> Nevertheless, this is not the interpretation I give to the utterance - for the simple reason that I know that the speaker (one of my students) does not know anything about my family life, while she does know that I have a colleague, Neil, who teaches her syntax. So, even if initially my highly activated NEIL1 concept is the first one accessed, it is soon replaced by NEIL2. This interpretation is also coherent and sufficiently relevant.

Now, Recanati (1993, 1995, 2004) assumes that top-down shifts of activation, i.e. from the NEIL1 to the NEIL2 concept, are indeed possible, but they can only be explained in terms of cognitive schemas or scripts. These roughly correspond to abstract scenario types that are evoked by salient features of the situation in which the sentence was uttered. According to Carston (2007, 28), Recanati might assume

---

opposed to contributing to truth-conditions proper and (ii) as devices that exploit contextual information in such a way that it gives rise to thoughts that are *immune to error through misidentification*. In Galery (2008), I suggested that given these two properties, REF is completely dispensable because they are enough to mark referentiality at the lexical level: the level at which Recanati assumes Kripke's intuitions about the modal and scope behaviour of pronouns should be met. However, I do not think any of these properties hold. Regarding (i), consider a scenario where someone sees 'I am the greatest' written on the wall. There is no speaker present, but the hearer can retrieve a Perry-type reflexive proposition of the form 'the speaker of that token is the greatest'. Regarding (ii), we could imagine someone walking in front of a group of people which includes a person who utters 'I am the greatest'. The hearer knowns that someone uttered it, but could be mistaken about who did: he may assume Peter did it where in fact John did. It might be the case that the necessity of capturing a notion of referentiality at the type level is a theory-internal demand.

that a *university* schema provides top-down activation to the representations about particular lecturers, hence explaining the shift to the NEIL2 concept. However, Carston (2007, 28) still presses on:

> ...even supposing this account goes through for the example as described so far, we can change the speech situation so that it does not make salient university departments, lecturers or subjects like syntax. Suppose I run into Sarah at the local supermarket, and after some chat about the merits of organic vegetables, she says to me, 'I hear that Neil's broken his leg'. Given the high activation of my NEIL1 concept and the absence of any frame to effect an accessibility shift, the prediction of the automatic, non-reflective associative account seems to be that (contrary to intuitions) I will take her to have said that NEIL1 has broken his leg.

Her point is quite simple. One can evoke as many cognitive schemas and scripts as one likes, but for an account in terms of activation-spread, this move faces the following shortcomings: (i) it seems to capture the data (in this case our offline intuitions) in a rather *ad hoc* way, and (ii) in some cases the cognitive schema does not seem to deliver the necessary kind of top-down influence that the audience's expectations about the speaker's intention exercise on the selection of mental representations.

Now the point I would like to make is this. Given that the activation-based explanation faces problems for the simple linguistically mandated primary pragmatic process of *saturation* and that the optional pragmatic process of *synecdochic transfer depends* on the interpretative level determined by saturation, the explanation of descriptive pronouns inherits the difficulties raised by Carston. Moreover, in some of the descriptive cases at hand, say, the footprint case, the use of the pronoun seems to rely on information that is *occasionally* available. It is hard to see what a footprint schema, necessary for the relevant accessibility shift, looks like.

The fourth and final problem rests on *the cognitive architecture required by the complex relations between mental file and information associated with it*. We saw that the descriptive uses of pronouns (as well as referential uses of descriptions) are explained by a process of *synecdochic transfer* involving information written in the file 'standing in' for the file itself in some cases, and vice-versa for some others. But then how are the old metonymy cases (see chapter 1 section 1.7 and chapter 2), like the one below, explained?

(15)   a.   In a restaurant, the waiter says to the chef: 'The ham sandwich is getting restless' (Nunberg, 1978, 1979).

      b.   Proposition expressed by (a): *The ham sandwich orderer is getting restless.*

For Recanati (1993, 263), the delivery of the truth-conditions, in (15)b, of the utterance in (15)a involves the primary pragmatic process of *metonymic transfer*. I take it to be an optional pragmatic process that takes a mental representation about sandwiches as input and delivers a representation about their orderers as output, but how should it be understood given Recanati's mental file approach?

On the one hand, we could understand metonymic transfer along the lines of synecdochic transfer. Understanding (15)a calls up a HAM SANDWICH file or encyclopedic entry which has information associated with (written in) it. One piece of such kind information is *.... has been ordered by someone*. Thus, like the synecdoche case, activations flows from the file as a whole to a specific piece of information written in it.

If the proposal runs along these lines, it is unclear why the specific notion of *metonymic* (as opposed to synecdochic) was reserved to describe this process. Moreover, one of the motivations for the mental file approach seems to be challenged. One of the roles that could be attributed to files is to organise information in the mind. Suppose one raises the following question. If top-down processes are indeed involved in linguistic understanding, what would block my knowledge that *dinosaurs cannot play chess* from playing a role in the interpretation of (15)a (see the *frame problem* in Fodor 1983)? Two aspects of Recanati's proposal may provide an answer. It seems that the association-based mechanism would not make this proposition accessible enough for understanding the utterance. I take this point to be well grounded. Alternatively, one could claim that the proposition about dinosaurs does not play a role in the interpretation of the utterance about ham sandwiches because it does not figure in a ham-sandwich file. But if the synecdoche model can be extended from descriptive pronouns and referential descriptions to cases of metonymy, metaphor and possibly every other non-literal interpretation, we run the risk of ending up with one gigantic file, because in the right context any piece of information can be related to a single mental file, say, *ham sandwich*.

There are two conclusions that might be drawn from this. First, if the argument is sound, then the theoretical notion of mental file is not playing any role and we are no better off by adopting it rather than just having a pre-theoretical (or metaphorical) understanding that information about sandwiches (and people who order them) is somehow organised around my HAM SANDWICH concept (as far as I know, this point was first raised by Woodfield 1991 and has not been discussed

properly by mental file theorists). Second, if mental files are indeed psychologically real and do theoretically interesting work, we would need to describe mechanisms by which information in files is organised (e.g. stored, hierarchically related to other information, and retrieved), otherwise understanding an utterance would invoke an impressively large information database, whose access would have to be reduced in order for it to play an interesting role in utterance comprehension. This role could be played by the activation-spread mechanism adopted by Recanati (or the cognitive principle of relevance, under RT), *but appealing to it would make the information-constraining job allegedly done by files redundant: one would only need a mechanism of activation spread or something like the principle of relevance*[11]. In conclusion, appealing to files, without specifying the principles by which they are individuated, is helpless; however, specifying such principles renders the appeal to files theoretically inert.

On the other hand, metonymic transfer might be understood as an operation between files. For example, when I understand (15)a, my file on ham sandwiches serves as input to a process that delivers another file, namely, one about restaurant customers. This sheds light on Recanati's use of a specific name (e.g. 'metonymic' as opposed to 'synecdochic') for this type of transfer process. However, one could raise the question of how many files does the individual mind have? Does my knowledge that *cats chase mice* enter my cat file or my mice file? In short, what motivates the use of two distinct files to explain the understanding of metonymies, on the assumption that information that could potentially be stored in one file (say, on ham sandwiches) could also be stored in another file (say, on restaurant customers). This strategy may run the risk of assuming that information that constitutes the file (like clay that constitutes a vase) acts as an individuator of the file (like saying that a vase which was repaired by the addition of new clay is numerically different from the vase that it was before the repair). Thus, my file on ham sandwiches (containing *...is a nice food* in it) is different from my file on ham sandwich orderers (containing *...spend money needlessly*) because the information written under the files is different.

This assumption has severe consequences for the metaphysics of mental representations and runs counter to the idea that files are individuated by their referential relations to the entities they are about. If this conclusion is to be avoided, then it seems that the mental file theorist who explains cases of (metonymic or synec-

---

[11]One could say that files themselves are the result of activation patterns among different kinds of information that have been crystallised. Although this sheds light on how information can be organised as files, it concedes to the point I am trying to make: the activation spread mechanism is doing the heavy work, not files themselves.

186

dochic) transfer in terms of two distinct files must explain the principles (other than the constitution of information) by which data is organised. Without these principles, the theoretical (as opposed to the metaphorical) notion of a mental file does not give us better understanding of what goes on in understanding non-literal utterances and other phenomena. Again, this point has been made by Woodfield (1991) but unfortunately it has not been given proper attention in the debate about mental representations as files.

In conclusion, Recanati's account explains part of the data. In some cases, it is dubious how the mechanism of spreading of activation delivers the interpretations that seem to be intuitively available. Finally, it seems to postulate mental entities and processes that face some problems.

Variations on Recanati's TCP

In this section, I want to highlight a point that might have already occurred to some readers, but which I think is worth making more explicit. Recanati's account assumes the following: (i) *type-referentiality*, via the postulation of REF, (ii) a mixed associationist/inferential view on pragmatic processes and (iii) an approach to mental representations based on *mental files*. Now, it is important to notice that these are not part of a package deal. The rejection of any of these assumptions does not entail the rejection of any of the others. For example, Powell (1998), replaces Recanati's mixed picture in (ii) for a fully inferential mechanism based on the principles of relevance, while preserving some form of type referentiality, as in (i), through the idea that pronouns encode procedures (Wilson and Sperber, 1993) which trigger the search for a specific mental representation. Although variations on Recanati's proposal may bring insights on how the data can be explained from different (perhaps better) conceptual standpoints, they would eventually face the same difficulty: *Accessibility* cannot be properly explained. Thus, for reasons of space, I cannot do full justice to an assessment of positions that result from denying each of the assumptions in (i) to (iii). In what follows, I will examine a more radical alternative to Recanati's TCP.

The underdetermination view - Bezuidenhout (1997); Powell (2003)

The underdetermination view begins with the rejection of the idea that referentiality should be marked at the level of linguistic types. For Bezuidenhout (1997) and Powell (2003) referentiality is only a property of linguistic tokens, that is, only the thoughts these expressions induce can be singular or general. This move raises

two questions. First, if referentiality is not marked at the level of linguistic types, what does the linguistic meaning of pronouns look like? Is the meaning of 'he', for example, a description along the lines of 'the male'? Second, what are the mechanisms in virtue of which the same lexical meaning can be developed into either a singular or general mental state?

The first question warns of a potential danger that many working within the tradition initiated by Reichenbach (1947), including Recanati, were so keen to avoid: if the meaning of pronouns is captured entirely *descriptively*, the connection between the word token (e.g. 'I') and the specific aspect of the context it depends on (e.g. *the speaker*) might be severed. For this reason, token-reflexivists, like Perry (2001), make the linguistic meaning of the word reflect its token. Thus, 'I' means 'the speaker of u', where 'u' is a specific word token [12]. As an alternative to the idea that pronouns necessarily have an indexical element in their encoded meaning, Anne Bezuidenhout (1997) proposed that certain linguistic types encode *underdeterminate information* with respect to the type of proposition (singular vs general) they determine. Powell (2003) develops this insight in more detail and tries to apply it to a variety of phenomena, including referential descriptions and descriptive pronouns. For this reason, I shall concentrate more on Powell's account.

The gist of the proposal rests on the claim that pronouns 'are not semantically constrained, i.e. constrained by their linguistic meaning, either to express *de re* [singular] concepts or to express descriptive concepts: they can do either' (Powell 2003, 31, square brackets mine). In his terminology, expressions that can be used to pick out a unique individual, such as 'John', 'he', or 'the man', introduce what he calls 'individual concepts' in the mind of the audience. It must be noted that individual concepts should not here be understood as functions from worlds or situations to individuals, as commonly held in formal semantic frameworks (see Elbourne, 2005). For Powell, such conceptual representations can be used to express singular or general truth-conditions. The pronoun, 'provide[s] only pragmatic guidance to the interpretation' (Powell, 2003, 133).

It is unclear what exactly Powell's characterisation of the meaning of pronouns amounts to, but I take it to be a variable-like bundle of information that includes grammatical features, such as gender, case, number, animacy and, at least in the case of singular pronouns, the presupposition that the entity referred to is unique. This requirement is neutral about the way in which it is satisfied. If the individual concept describes the satisfier of a cluster of properties, the thought is general or

---

[12] A similar move has been made by Elbourne (2001, 2005, 2008). He postulates that pronouns have a structure containing indices, which can be assigned to specific individuals.

descriptive. Else, the individual concept denotes a specific individual on the basis of certain epistemic relations between the communicator and referent, that is, a singular thought is communicated. Furthermore, not all the encoded information needs to be satisfied in the process of interpretation, if its suppression leads to an overall more relevant interpretation. For example, Powell (2003, 133) mentions an example where someone refers to a building by the pronoun 'she'. In this case, some of the pronoun's features were 'deleted', so the hearer could grasp the proposition expressed.

With regards to the second question, namely, which mechanisms allow individual concepts to be developed into either singular or general thoughts, both Bezuidenhout (1997) and Powell (2003) adopt a full-fledged inferential view based on the relevance theoretic framework laid out in chapter 2. In a nutshell, RT assumes that cognition is geared towards the maximisation of relevance (Cognitive Principle of Relevance) and that ostensive stimuli (intentionally overt uses of words or gestures) are particularly important because they create the expectation that they are relevant enough (i.e. optimally relevant) for it to be worthwhile for the audience to process them (Communicative Principle of Relevance). Thus, by processing such stimuli, the audience forms certain hypotheses about what the speaker meant by her utterance or gesture. The most relevant hypothesis that counts as a development of the logical form of the utterance is attributed as speaker explicature(s), which is also used in the (mutual) shaping of the utterance's implicatures. Note that the audience's particular expectation of relevance may be shaped by specific information in the context (something richer than the presumption that the stimuli is worth the audience's processing effort). For example, a question usually enriches the expectations of relevance in the context: its reply must not only be worth being processed by the hearer, but also provide the means to an answer. Let us put these notions to work (borrowing from the relevance theoretic schemas used in chapter 2 section 2.3 and using '?' to mark the individual concept introduced by the pronoun). Consider the adaptation of our first example in (1), given below and my attempt to reconstruct a derivation based on Powell's ideas, as he unfortunately does not explicitly provide us with one:

(16)　　Someone pointing at a huge footprint in the sand says: 'He is gigantic!' (adapted from Schiffer 1981, 49)

　　　　a.　? IS GIGANTIC [partial decoding of the utterance's linguistic meaning (via the communicative principle of relevance and the presumption of relevance it conveys)].

b. THERE IS A BIG FOOTPRINT ON THE GROUND; SOMEONE MADE THAT PRINT; [assumptions made salient by the ostensive gesture towards the footprint (via the communicative principle of relevance and the presumption of relevance it conveys)] .

c. BEING GIGANTIC EXPLAINS WHY ONE WOULD LEAVE A BIG PRINT [implication (cognitive effect) derived from the encoded meaning of the words and the perceptual experience (supplementation of the expectation of optimal relevance by specific presumptions warranted in the context, given that no-one in the context is an immediate referent of the pronoun)].

d. THE PERSON WHOSE FOOT MADE THE PRINT IS GIGANTIC [unpacking of the individual concept into an explicature derived from (b) and (c)].

Now, Powell's theory is extremely flexible. In cases where the scope test suggests that the relevant descriptive interpretation should be captured at the level of the utterance's implicatures, one could add a couple of extra steps in the schema in (a) to (d) above, and the desired implicatures would come out. Moreover, this framework is flexible in the sense that its basic tenets apply generally to all pronominal forms (e.g. plurals; 2nd person, etc.). But how does it fare with the *desiderata* that needs to be explained?

The basic three core interpretative properties fall out straightforwardly. *Identification Dependency* is met by (16)b, if the object were not identified as a footprint, those mental representations would not be available. *Descriptiveness* is met by the descriptive explicature in (16)e. The relevance-driven inferential relation between the available assumptions and the explicature meets *Connection*. Moreover, since the individual concepts can be developed into either descriptive or singular content and given that the contend induced by descriptions can co-vary with other expressions in discourse, the framework might have the resources to explain *Co-variation*, exemplified below.

(17)   *Andy*: Well...**Every time there is a war**, **he** [*pointing to the White House*] has tough choices to make.

That is, the pronoun 'he' would unpack as the description *the American President*, which can depend on other quantifiers in discourse, possibly ensuing a covarying reading for the utterance. Like Recanati's proposal, Powell's does not lay out the mechanism by which these readings are derived (necessary for a full expla-

nation of *Co-variation*). Still, I take it that this property of the data can be success-fully accounted for by the proposal. However, the similarities with Recanati's do not stop short from the problems raised previously. The underdetermination view seems to overgenerate in cases where descriptive pronouns serve as antecedents of other anaphoric pronouns. The contrast between the utterances below is, once again, impossible to explain.

(18) a. *Said by someone:* If the Democrats had won the last few presidential elections, **the American Supreme Court Justices** might have been liberals. **They** would guard public interest better.

  b. *Said by Supreme Court Justice O'Connor:* If the Democrats had won the last few presidential elections, **we** might have been liberals. ***They** [*said by same speaker*]/**They** [*said by different speaker*]/?**We** [*said by same speaker*] would guard public interest better.

The underdetermination view can offer a solely pragmatic explanation for the *Accessibility* property, but this would not do justice to the role played by grammatical person in the asymmetry above. In short, the proposal characterises linguistic mean-ing as underdeterminate, but it does not say what linguistic meaning is. Describing it seems crucial for a proper account of the data.

Problems for the underdetermination view

In this subsection, I present three problems for the relevance theoretic underdeter-mination view.

First, like all the approaches discussed in this chapter, the proposal fails to ac-count for all the properties raised by the pronominal uses discussed in the thesis.

Secondly, the fact that the underdetermination view rejects any account of refer-entiality at the level of lexical types may be allegedly counter-intuitive (many would claim this would be the case). One could say that pronouns are used indexically most of the times, whereas descriptions are used to convey general propositions most of the time. Arguably, these intuitions should be captured. Powell (2003, 168-9) responds to this worry by claiming that 'we are cognitively designed primarily to track physical objects rather then properties'. I believe that given that pronouns do not encode concepts that enter truth-conditional content and rather make dependent contributions, they would be less costly to process than full-blown descriptions. Plus, pronouns are usually shorter. This explains the preference for pronouns over descriptions to convey deictic readings rather nicely, as descriptions would involve

unnecessary processing effort. Moreover, one could just reject the intuitions that pronouns are deictic by linguistic nature. Bound and donkey cases are genuine descriptive pronominal uses. In other words, once we factor these pronominal uses in, the intuition that these expressions favour singular interpretations loses its grounds. It seems that the second objection against Powell does not offer any knock down argument against the underdetermination view.

The third and final complaint rests on the nature of individual concepts and for this reason applies only to Powell's approach (not Bezuidenhout's). One could argue that saying that a mental representation is underdeterminate is misleading: words can encode underdeterminate information which can integrate with different types of mental representation, resulting in singular or general mental representations, but to say that such mental representations themselves are underdeterminate seems obscure. How can concepts be underdeterminate? I do not know how Powell might respond to this, but I also do not know whether the objection has any punch; for individual concepts themselves might be intermediaries to well delineated conceptual representations, rendering the objection terminological. Thus, I do not think that it challenges Powell's proposal significantly, unlike the first objection.

4.5 Conclusion

I have presented a variety of accounts of descriptive pronouns and discussed a variety of arguments against them. Since for each position there is at least one argument that seems to be sound, we have the motivation to look for an alternative explanation of the data. Moreover, many of them suffer from incompleteness in the sense that they do not account for the extended properties of descriptive pronouns (*Covariation* and *Accessibility*). In the next chapter, I will outline a proposal that does not suffer from this shortcoming.

Chapter 5

Descriptive pronouns: a reanalysis

5.1   Introduction

The aim of this chapter is to come up with an account of the core and the extended properties of identification-based descriptive pronouns (reviewed in chapter 1). Consider again the following examples:

(1)   a.   *Said by someone:* If the Democrats had won the last few presidential elections, **the US Supreme Court justices** might have been liberals. **They** would guard public interest better.

   b.   *Said by Supreme Court Justice O'Connor:* If the Democrats had won the last few presidential elections, **we** might have been liberals. ***They** [*said by same speaker*]/**They** [*said by different speaker*]/??**We** [*said by same speaker*] would guard public interest better.

In summary, in (1)b, 'we' contributes *the US Supreme Court Justices* to the truth-conditions of the utterance (*Descriptiveness*). In order for these truth-conditions to be communicated, the audience must identify the speaker (O'Connor) as belonging to the kind *Justice* (*Identification Dependency*). This shows that some form of categorisation by the audience provides the building blocks of the relevant communicated propositions (*Connection*). Moreover, the descriptive interpretation can co-vary with other constituents provided by other expressions in discourse. In (1)b, the Justices co-vary with possible worlds/situation introduced by the modal 'might' (*Co-variation*). Finally, certain descriptive interpretations can be antecedents for other dependent uses of pronouns, while others cannot (*Accessibility*). This is illustrated by the contrast between (1)a and (1)b.

In chapter 4, many proposals in the literature were reviewed. Although they were quite varied in terms of how they explained the data, they shared two drawbacks. First, the mechanism by which descriptive interpretations co-vary with other expressions in discourse was not properly explained (with the exception of Elbourne 2008). Second, and more importantly, the contrast between (1)a and (1)b (e.g. *Accessibility*) was not predicted. On the one hand, the more semantic-based accounts of the phenomenon (section 4.3) did not use a model of grammar that could recruit pragmatically available constituents that could be seen as 'antecedents' for pronominal interpretation. On the other hand, the more pragmatically oriented accounts (section 4.4) assumed that the phenomenon is entirely pragmatic and free from any grammatical constraints, failing to predict the asymmetry as well.

The proposal in this chapter overcomes this explanatory deficit. In order to do so, we need to outline a grammar that has the following properties: (i) it is *not* restricted to describing only relations within strings and effectively relates the representation of the proposition expressed to contextual information, (ii) it describes the encoded meaning of pronouns in such a way that they can receive singular or descriptive interpretations, (iii) it introduces a way to keep track of multiple perspectives in a dialogue, so that certain anaphoric impossibilities are explained, and (iv) it is a unified mechanism that explains not only the descriptive uses of pronouns, but also all their other types of interpretations (reviewed in chapter 1) straightforwardly.

These theoretical needs are met by coupling the relevance theoretic pragmatic framework (described in chapter 2) and the incremental view of object representation (described in chapter 3) to a grammar framework known as *Dynamic Syntax* (DS, for short, see Kempson et al. 2001; Cann et al. 2005, *inter alia*), which will be introduced in the next section. In section 5.3, I use the recently introduced grammar together with the other elements of the account to explain the data. Identification-based descriptive interpretations are handled in a heterogeneous way: some of them contribute to the utterance's implicatures, while others contribute to the utterance's explicatures. Determining to which level the use contributes to depends on contextual factors and also on the choice of the pronominal form (more specifically, the grammatical person and number encoded). In a nutshell, the proposal made here is a form of the *substitution-based account*, sketched towards the end of chapter 2. The rough idea is this. Pronouns contribute to variable-like constituents that can be replaced by indirectly available conceptual representations (along the lines described in chapter 3), in some cases, or by singular (object-dependent) representations that determine propositions that can contextually implicate or entail the relevant descrip-

tive interpretations, in some other cases. This makes the proposal quite similar to Powell's relevance-theoretic proposal in chapter 4 but describes, in DS-terms, what the meaning of pronouns is ina way such that Co-variation and Accessibility are properly explained. In section 5.4, I state concluding remarks.

5.2 Dynamic Syntax: a crash course

Dynamic Syntax (DS) (Kempson et al., 2001; Cann et al., 2005) is a grammar framework that aims to answer two challenges. On the one hand it seeks 'to state the interaction between order of words within the sentence' (see Cann et al., 2005, 2) and the structural constraints that it imposes. On the other, it is designed to incorporate contextual contributions to the interpretation, so various semantic and pragmatic phenomena are properly accounted for.

It is interesting to note the differences between DS and other grammar frameworks. The second theoretical aim mentioned above sets it against the orthodox view that only strings or sentences bear syntactic significance (i.e., the 'narrow' view of grammar briefly mentioned in chapters 1 and 3, see Chomsky 1957, 1965, 1981). As DS stresses, the relations between the constraints imposed by linguistic structure are constraints imposed by (pragmatically available) contextual information play a crucial role in linguistic theory. The data displayed by identification-based pronouns illustrates this nicely: we need a partially grammatical and a partially pragmatic account. Accordingly, DS aims to capture these relations in a unified way, so that the interplay between language and context output a single level of representation (structured propositional content in a Language of Thought, Fodor 1975). This sets DS against frameworks that achieve this goal by assuming two (or more levels of representation), such as postulating a Chomskyan-style grammar at the level of syntactic representation and, for example, (certain versions of) Discourse Representation Theory (DRT, Kamp 1981) at the semantic/pragmatic level. As we shall, these unification is achieved by assuming that lexical items encode a lot of information (including sets of procedures that allow the expansion of syntactic structure, i.e. the output representation is induced from lexical material in a bottom-up fashion) and that the means to represent encoded information can also be used to represent contextual information. Both the bottom-up aspect of the theory and the complexity of information it is able to represent approximate DS to Head-Driven Phrase Structure grammars (HPSG, Pollard and Sag 1994). The resulting grammar/pragmatics interface is unorthodox: the parse of the sentence is coupled with pragmatic aspects of utterance interpretation, yet there is a way to distinguish com-

petence from performance. As we shall see soon, the former can be regarded as the set of moves (computational operations) allowed by the system, whereas the latter can be regarded as a particular choice of operation (including pragmatic operations) at a given stage in interpretation.

The formal treatment of these theoretical tenets requires some basic ingredients. First, the fact that linguistic representations are hierarchically structured (see the discussion on binding in chapter 1 section 1.4) is captured in terms of binary LINKed trees, described by the *Logic of Finite Trees* (LOFT) of Blackburn and Meyer-Viol (1994): a modal vocabulary that describes the relations amongst the various nodes in a tree (points in logical space) via an accessibility relation (represented by symbols such as $<\uparrow>, <\downarrow>$). Second, information that holds at a node, such as truth-conditional contribution or syntactic features, for example, is captured by declarative units (DU) that 'annotate' or 'decorate' a given node. This comprises the *declarative* part of the framework. Third, linguistic underdetermination is modelled as mappings from sub-trees to further developed trees. Such mappings are regulated by three factors: (i) *requirements*, i.e. the fact that certain information must hold at a node, promoting tree-building (represented by a '?' preceding the required information), (ii) underspecified information (a node having unfixed position or underdeterminate content), and (iii) grammatical rules that allow for transitions between (partial) trees (such as rules that remove a requirement once it is met). This comprises the procedural part of the language.

The incorporation of the *Labelled Deductive Systems* of Gabbay (1996) into the framework allows for these various ingredients to be recruited, in a unified way, as descriptors of a (partial) tree, that is, a set of tree nodes. Each node is described by a triple of the following format: *[requirements, $<labels>$, formula value]*. The components of node descriptions as well as how the framework models linguistic understanding is the topic of the next subsection.

5.2.1 The language *DU*

Declarative Structure: A sketch of the Language DU

In DS, binary trees are rather simple. In addition to the root node, tree-nodes can be either *arguments*, always represented by the daughter node at the left hand side, or functors, always represented by the daughter at the right hand side. Thus, the scaffolding of the basic representations in the formal system is this:

(2)   Tree structure   root node

       terminal argument node   non-terminal functor node

                                   terminal argument node   terminal functor node
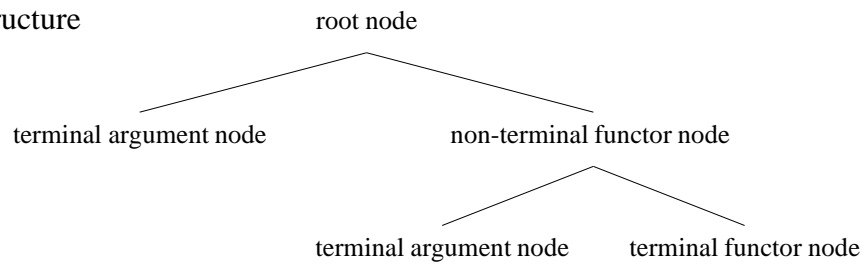
   As mentioned in passing, the language DU - the formal system that captures
the grammar - employs decorations or annotations to describe the nodes in a tree
as well as the relations amongst them. In the bare bone structure above, I have
informally shown this by annotating 'root node' or 'argument node' at certain places
in the tree. What we are going to do now is replace this intuitive vocabulary by a
formal description. Note that a (partial) tree is a set of tree node descriptions, for
example, { {node description 1},{node description 2}, {node description 3} ... },
each being described by a *[requirements, $< labels >$, formula value]* tuple. The
unique component of the formula value field is the predicate Fo. Additionally, the
predicates Ty (logical type), Tn (node position), $<\uparrow>$ or $<\downarrow>$ (relations amongst
nodes) inhabit the *labels* field. Let us begin with the first of these predicates.

The Predicate Fo

The first decoration to be introduced corresponds to the truth-conditional contribu-
tion of a given tree node. In DS, this kind of information is captured by the predicate
Fo (formula value). For example, the utterance in (a) can convey the structure in
(b) with certain formula values at each node. This structure in turn, is abstractly
(i.e. less graphically) described by (c).[1]

(3)   a.   John smokes

       b.
                  $Fo(Smoke'(John'))$
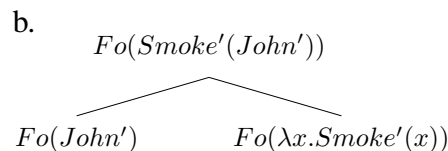
            $Fo(John')$ $Fo(\lambda x.Smoke'(x))$

       c. $\{\{..., Fo(Smoke'(John'))... \},\{...Fo(John')...\}, \{...Fo(\lambda x.Smoke'(x))...\}\}$

The Predicate Ty

---

[1]For the sake of simplicity, I will abstract from tense and treat names as logical constants.

The second decoration that we shall look at is the predicate *Ty*. It features in the field *label* and represents information concerning the logical type of a certain node. This imports the combinatorial properties of the lambda calculus into the DU language: for example, given the type information e and $e \rightarrow t$, we could deduce the type information t by modus ponens and the corresponding formula values via function application. The language DU assumes as basic types, *e* (entities), *t* (truth-values), *cn* (common nouns). A small variety of functional categories is defined in the specific lexical items of the language. Thus, the description in (3) can be enriched as follows:

(4)   a.

$$Ty(t), Fo(Smoke'(John'))$$

$$Ty(e) \qquad Ty(e \rightarrow t)$$
$$Fo(John') \qquad Fo(\lambda x.Smoke'(x))$$

         b. $$\{\{Ty(t), ..., Fo(Smoke'(John'))\},$$
$$\{Ty(e), ..., Fo(John')\}, \{Ty(e \rightarrow t), ..., Fo(\lambda x.Smoke'(x))\}\}$$

The Predicate Tn

The third type of tree node description that enters in the field *label* of the triple above is the predicate *Tn*. It represents the position of the tree node vis-a-vis other nodes in the tree. The Tn predicate employs a recursive tagging mechanism to describe node positions. By definition, the position of the root node is described by $Tn(0)$. The position of the argument daughter node(s) is recursively defined by appending an extra '0' to the number of the node of the mother. Thus, the argument daughter node from the root is described by $Tn(00)$. The position of the functor node is formed by appending a '1' to the end of the mother's Tn value. Thus, the functor daughter node from the root is described by $Tn(01)$. As we shall see shortly, the position of a tree node is used for a variety of purposes, such as keeping track of the point of the parse, supporting the descriptions of relations that hold between nodes, and describing the unfixed position of a node (which captures movement-like effects postulated by Chomskyan grammars). A richer description of the structure in (4) is:

(5) a.

$$Tn(0), Ty(t), Fo(Smoke'(John))'$$

$$Tn(00), Ty(e) \qquad Tn(01), Ty(e \rightarrow t)$$
$$Fo(John') \qquad Fo(\lambda x.Smoke'(x))$$

b. $$\{\{Tn(0), Ty(t), ..., Fo(Smoke'(John)')\},$$
$$\{Tn(00), Ty(e), ..., Fo(John')\}, \{Tn(01), Ty(e \rightarrow t, ..., Fo(\lambda x.Smoke'(x))\}\}$$

LOFT and relations amongst nodes

Given the necessity to describe relations among tree nodes, the language DU incorporates the vocabulary of the *Logic of Finite Trees* (LOFT, Blackburn and Meyer-Viol 1994). It amounts to a modal logic that treats nodes as points in logical space (in the same way that possible worlds are). Nodes then are related by two forms of accessibility relation: an existential one, symbolised by $<>$ and intuitively read as 'there is a node such that...', and a universal one, symbolised by [ ] and intuitively read as 'for all nodes...'. In order to be able to express relations about mothers, grandmothers, sisters, daughter, etc... LOFT employs operators symbolised as up and down arrows. Thus, from a given node, $<\uparrow>$ means 'at the mother node', and $[\downarrow]$ means 'for all the daughter nodes'. Furthermore, these modal operators can receive affixes describing relations that hold at argument nodes, functor nodes or underdetermined nodes (by the Kleene * operator). For example, $<\uparrow_1>$ **XYZ** means that 'the description XYZ holds at the functor node above the current one', $<\downarrow_0>$ **XYZ** means that 'the description XYZ holds at the argument daughter', and $<\downarrow_*>$ **XYZ** means that ' the description XYZ holds for some node (argument, functor, mother, grandmother, ...) below the current one'. Operators (also described as modal relations) can be iterated, so $<\uparrow_1><\downarrow_0>$ **XYZ** means 'the description XYZ' holds at the argument daughter of the functor node above the current one'. Here is a simplified example of the use of such vocabulary:

(6) $$Tn(0), <\downarrow_1><\downarrow_0> \textbf{XYZ}$$

$$Tn(00) \qquad\qquad Tn(01)$$

$$Tn(010) \,\textbf{XYZ} \qquad Tn(011)$$
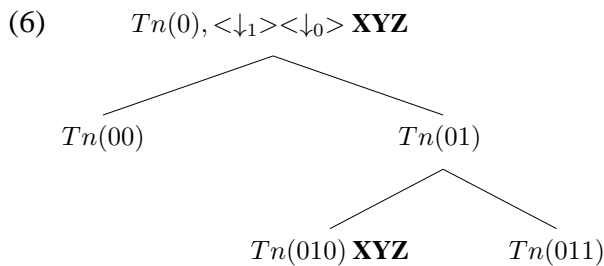
DU incorporates two more symbols from the LOFT formalism: the verum, $\top$, and

199

the falsum, $\perp$ (bottom restriction). The first interacts with LOFT's modal language by expressing that a given node description is 'true' for (some/every) tree node related to the current one (in a similar way a proposition is true at (some/all) accessible possible worlds),while the latter expresses the converse. For example, $[\downarrow]\perp$ expresses that no node description holds at any node below the current one. This indicates that the tree cannot be further developed below that point.

Procedural Structure: Partiality and Incrementality

Given the tools that describe the information that annotates a tree-node (the declarative structure), we are in a position now to state how the system accommodates the development of partial propositional forms into full blown truth-evaluable content. This amounts to the procedural structure of the system. It comprises a pointer (symbolised by '$\diamondsuit$') to indicate the current node of the parse (relevant for information processing purposes) and three main other elements: (i) requirements, (ii) actions (lexical or computational), and (iii) underdetermined content (e.g. metavariables) or structure.

Requirements

Roughly, requirements are node decorations that do not yet hold at that node. They model a temporary stage in utterance comprehension and, in fact, promote the development of the tree by triggering a series of operations that may achieve the required decoration (requirements are then written off by a computational rule and successful tree-building ensues). If requirements are not met, tree building fails and the utterance is declared as unacceptable, meaning that a representation for it cannot be produced given the available resources and/or stimulus.

The best example of a requirement is the starting point of the utterance comprehension. Simplifying for the sake of exposition, DS models utterance comprehension as the task of building a propositional type (which is taken to represent the speaker's intended meaning[2]). Thus, the first stage in tree development is an instance of the requirement below (called ' the Axiom', from which everything follows):

(7) $\{\{..., ?Ty(t), \diamondsuit ...\}\}$

Above we have an abstract description (in the same way as (3)c, (4)b and (5)b de-

---

[2]Though see Gregoromichelaki et al. (2011) for the idea that something more basic than the appeal to intentions may underpin some more basic forms of conversational exchange.

scribe their respective trees) of the type information that should hold at the root node. As we shall see shortly, the encoded meaning of words or rules of the grammar can use these requirements in procedures that extend the structure of the tree.

Actions (Computational and Lexical)

The procedures or instructions employed by DS are of two types: *computational actions* and *lexical* actions. Both can be considered as what is traditionally defined as the rules of a grammar. Like the model in Government and Binding (GB) theory (Chomsky, 1981), computational rules are entirely optional. They constitute the space of moves and choices that the performance system *may* draw upon [3]. The general format of such rules is conveniently displayed under the following pattern:

(8)   *General format of computational rules*

Tree Description 1

Tree Description 2

Rules allow certain transitions from an input description (Tree Description 1) to an output description (Tree Description 2). Transitions can be considered proofs (suppose tree description 1 contains $P \rightarrow Q, P$ and tree description 2 contains Q). By incorporating Gabbay's (1996) LDS, the language DU can use *any* decoration type (e.g. Ty, Fo, Tn, LOFT modalities) in the definition of grammatical rules. This adds expressive power and a more controlled way of deploying transitions between tree-structures.

The first rule that we shall look at is *Introduction*. It is used to break down the goal of building a type *t* node into the subgoals of building a type *e* node and a type $e \rightarrow t$ node. Formally, it is defined as follows (the definitions of the computational rules are reprinted in an Appendix at the end of the chapter, greyed out text indicates repeated information that is not the focus of the portion of the rule under discussion):

(9)   *Introduction*

$$\frac{\{...\{..., ?Ty(Y).., \Diamond\}...\}}{\{...\{..., ?Ty(Y), ? <\downarrow_0> Ty(X), ? <\downarrow_1> Ty(X \rightarrow Y)..., \Diamond\}, ...\}}$$

The symbols *Y* and *X* stand for any specifications of logical types. *Introduction*

[3] In DS the competence - performance distinction roughly mirrors the distinction between a set of rules that govern tree-building and specific choices of rule application in order to build a particular (partial) tree.

effectively creates new requirements at the current node (not on the nodes themselves) using the modal vocabulary of LOFT. It basically allows a move from a node that requires a propositional type (Y) to the requirements of an argument daughter of type (X) and a functor daughter of (functional) type $X \rightarrow Y$. These represent the yet to be built subject and predicate nodes, whose construction is done by another computational rule, namely, *Prediction* (shown below). It is important to observe here that type e nodes can accommodate content that represents what one would intuitively regard as objects and individuals, such as John, Mary, or this chair, but also, following Gregoromichelaki (2006), entities corresponding to situations/eventualities (parts of a world). The situation node (represented by $Ty(e_s)$) plays a major role in the evaluation of the utterance with regards to truth or falsity (see Papafragou 2000 on the actual world as a default for evaluation)[4]. After *Introduction*, the next rule is:

(10)   *Prediction* (Subject and Predicate)

$$\frac{\{...\{Tn(n),...? <\downarrow_0> X, ? <\downarrow_1> Y, \Diamond\}...\}}{\{\{Tn(n),...? <\downarrow_0> Ty(X), ? <\downarrow_1> Y\}, \{<\uparrow_0> Tn(n), ?X, \Diamond\}\{<\uparrow_1> Tn(n)?Y\}\}}$$

Prediction establishes two things. First, it creates the daughters (argument and functor) as required by *Introduction*. Second, it moves the pointer to the newly created argument node.

At this point, the definitions seem to be quite abstract, so it is best to illustrate with an example of initial stages in the parse. Consider a person hearing 'John smokes'. As a tentative approach to English (an SVO language), suppose that the subject and predicate nodes are built by a combination of *Introduction* and *Prediction*. On the assumption that recognising linguistic stimuli sets the hearer up with the task to build a truth-evaluable type, the axiom (see (7)) serves as input to *Introduction*. This move is represented by the move from (a), where the partial tree contains only the root node, to (b) below:

(11)   a. $Tn(0), ?Ty(t), \Diamond$
         b. $Tn(0), ?Ty(t), ? <\downarrow_0> Ty(e), ? <\downarrow_1> Ty(e \rightarrow t), \Diamond$

After *Introduction*, the application of *Prediction* creates the subject and predicate nodes and moves the pointer to the subject node. This makes the transition from (11) to:

---

[4]In this brief exposition, I will currently abstract away from situation arguments, which will be discussed in later sections.

(12) $Tn(0), ?Ty(t), ? <\downarrow_0> Ty(e), ? <\downarrow_1> Ty(e \rightarrow t)$

$Tn(00), ?Ty(e)\Diamond$ $Tn(01), ?Ty(e \rightarrow t)$
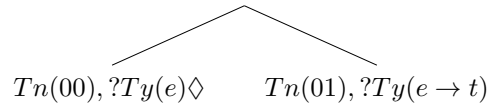
With the pointer positioned at the $Ty(e)$ node, the subject of the utterance can be parsed. Encoded information is represented by a similar form of procedure, namely, the lexical actions associated with a certain expression. In a rather simplified manner, let us suppose that the lexical entry of 'John' can be modelled along the following lines (again I'm simplifying things by assuming that a proper name contributes a constant denoting a specific person, as tentatively assumed in Cann et al. 2005, ch 2):

(13)   Lexical actions for 'John'

IF	$?Ty(e)$	Trigger
THEN	put $Ty(e), Fo(John')$	Actions
ELSE	abort	Else Statement

Lexical actions have the following structure. First they test whether a given condition obtains in the tree structure. This is stated in the trigger row above. If the conditions obtain, then a series of actions (represented in the Actions row by typescript) must be taken. If the conditions of the trigger do not obtain, then the system aborts through the else statement: the parse is incomplete and a propositional form cannot be assigned to the utterance. Given that in our example, the conditions are met, the tree is developed into[5]:

(14) $?Ty(t), ? <\downarrow_0> Ty(e), ? <\downarrow_1> Ty(e \rightarrow t)$

$?Ty(e), Ty(e), Fo(John'), \Diamond$ $?Ty(e \rightarrow t)$
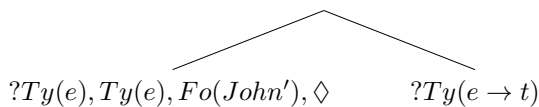
Note that the lexical actions for 'John' satisfy the type requirements of the node as well as decorating it with a formula value. We need now a rule that removes a requirement once it has been satisfied. This is performed by *Thinning*.

(15)   *Thinning*

---

[5]From now on, I will abstract away from DU predicates that are irrelevant for the computational rule under discussion.

$$\frac{\{...X, ...?X, ...\Diamond\}}{\{...X, ..., \Diamond\}}$$

Note that *X*, above, ranges over any kind of node decoration. Applying the rule makes the transition from (14) to:

(16) $?Ty(t), ? <\downarrow_0> Ty(e), ? <\downarrow_1> Ty(e \rightarrow t)$

~~$?Ty(e)$~~$, Ty(e), Fo(John'), \Diamond$ $?Ty(e \rightarrow t)$

What we need now is a rule that is able to register at the mother node the fact that a requirement has been met at a daughter node. This is the task of *Completion*. The rule takes as input tree descriptions that have the pointer at a daughter node containing a satisfied requirement and outputs tree descriptions with the pointer at the mother node and records the satisfaction of requirements there. The formal definition of the rule is:

(17)   *Completion*

$$\frac{\{...\{Tn(n)...\}, \{<\uparrow_i> Tn(n), ...Ty(X), ...\Diamond\}\}}{\{\{Tn(n), ... <\downarrow_i> Ty(X), ..., \Diamond\}, \{<\uparrow_i> Tn(n), ...Ty(X), ...\}\}}$$
where $i \in \{0, 1, *\}$

The application of *Completion* effectively makes the transition from (16) to:

(18) $?Ty(t), ? <\downarrow_0> Ty(e), <\downarrow_0> Ty(e), ? <\downarrow_1> Ty(e \rightarrow t), \Diamond$

$Ty(e), Fo(John')$ $?Ty(e \rightarrow t)$

A subsequent use of *Thinning* removes the $? <\downarrow_0> Ty(e)$ requirement at the root and we are now in the position to move the pointer to the predicate node, so information introduced by the verb can be processed. This is done by a rule that allows the pointer to move wherever any outstanding requirements exists. This is done by *Anticipation*. The rule basically takes as input tree descriptions that have the pointer at a node with requirements below it and outputs a tree description with the pointer at the nodes where the requirements hold. Formally, this amounts to:

(19)   *Anticipation*

$$\frac{\{\{Tn(n)..., \Diamond\}, \{<\uparrow> Tn(n), ...?X\}\}}{\{\{Tn(n)...\}, \{<\uparrow> Tn(n), ...?X....\Diamond\}\}}$$

The application of *Anticipation* makes the transition from (18) to the tree below, where the pointer is at the functor node:

(20) $?Ty(t), ? <\downarrow_1> Ty(e \to t), <\downarrow_0> Ty(e)$

$Ty(e), Fo(John')$ $?Ty(e \to t), \Diamond$

At this point, the verb can be parsed, resulting in the tree below. The following is a (simplified version of) the lexical actions associated with 'sleep'[6]:

(21)   Lexical Actions of 'sleep'

IF	$?Ty(e \to t)$
THEN	put $Ty(e \to t), Fo(Sleep'), [\downarrow]\bot$
ELSE	abort

(22) $<\downarrow_0> Ty(e), ? <\downarrow_1> Ty(e \to t), ?Ty(t)$

$Ty(e), Fo(John')$ $?Ty(e \to t), Ty(e \to t), Fo(Sleep'), [\downarrow]\bot, \Diamond$

An application of *Thinning* removes the requirement $?Ty(e \to t)$ at the functor node. Then a subsequent application of *Completion* records the predicate information in the root node, and *Thinning* (represented by strikethrough) removes the $? <\downarrow_1> Ty(e \to t)$ requirement at the root. The only requirement left is $?Ty(t)$. We are nearly done. Now we need a rule that performs modus ponens on the type-logical information and functional application over the formula values of the daughters and also records the result at the mother node. This is done by the *Elimination* rule, which ensues the following structure:

(23)   *Elimination*

$$\frac{\{<\downarrow_0> (Fo(a), Ty(X)), <\downarrow_1> (Fo(b), Ty(X \to Y))..., \Diamond\}}{\{...\{Ty(Y), Fo(b(a)), <\downarrow_0> (Fo(a), Ty(X)), <\downarrow_1> (Fo(Y), Ty(X \to Y))..., \Diamond\}\}}$$

---

[6]For purposes of presentation, I am abstracting away from tense (see Cann, 2011) and presenting a rather simplified entry; $[\downarrow]\bot$, the bottom restriction, indicates that the tree cannot be developed further; transitive and ditransitive verbs would involve the creation of more structure to accommodate object nodes.

(24) $?Ty(t), Ty(t), Fo(sleep'(John')), <\downarrow_0> Ty(e), <\downarrow_1> Ty(e \rightarrow t), \Diamond$

$$Ty(e), Fo(John') \qquad Ty(e \rightarrow t), Fo(Sleep'), [\downarrow]\bot,$$
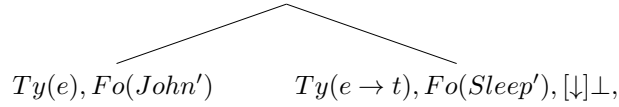
One last application of *Thinning* removes the final $?Ty(t)$ requirement introduced by the axiom. The parse is complete and a full propositional form has been retrieved on the basis of the linguistic meaning of the utterance. This ends the introduction to computational and lexical actions.

Content underspecification: Metavariables and pronouns

Metavariables, indicated by bold capital letters, such as **U**, **V**, **W**, represent a place-holder for an actual node decoration. While they are not proper values for the *Fo* predicate and hence need to receive some value from the context (via a process of *Substitution*, similar to Recanati's (1993; 2004) notion of saturation, see chapter 4 section 4.4.2), actual variables - u, v, w - are indeed proper values and hence can be bound by quantifiers (to be introduced in the next section) or indeed used in terms that substitute metavariables. Accordingly, metavariables are accompanied by requirements that will not be satisfied unless a proper value is found. DS assumes that pronouns (and other elements, like auxiliaries, see Cann 2011) introduce metavariables. The fact that they need the application of further (pragmatic) actions in order to make their truth-conditional contribution renders their meaning essentially procedural (as argued by Wilson and Sperber 1993, Powell 1998, Breheny 1999, among others). The lexical entry of the third person pronoun, for example, has the following form[7]:

(25)   Lexical actions for 'he'

|	
IF	$?Ty(e)$
THEN	put $Ty(e), Fo(\mathbf{U}_{male'}), ?\exists \, \mathbf{x}.Fo(\mathbf{x}), [\downarrow]\bot$
ELSE	abort

Note that the predicate in subscript represents the presupposition that the substituted formula must satisfy, that is, the gender feature of 'he' presupposes that the relevant formula value for *Substitution* selects male entities. As we shall see shortly,

---

[7]Abstracting away from case.

I will propose that person features encoded presuppositions about the perspective from which a formula value is determined. This can account for the *Accessibility* property of the data, but for the moment we will abstract away from this bit of encoded information. Now, suppose that instead of uttering 'John sleeps', the speaker utters 'He sleeps' whilst pointing to John. After *Introduction* and *Prediction* apply, the pointer finds itself in the argument node $?Ty(e)$ node, where the lexical actions of the pronoun can be triggered. They introduce type information $Ty(e)$, a metavariable **U** (i.e. a temporary formula value) and the requirement $?(\exists\, \mathbf{x}.Fo(\mathbf{x}))$ that a proper formula value must be assigned to this metavariable. In this context, identifying the person demonstrated may give rise to the formula value $John'$, which replaces the metavariable (via the application of *Substitution*) and satisfies the associated requirement (via *Thinning*). The remainder of the structure can be completed via the application of the same steps as before. Note that *Substitution* is not a grammatical operation, and hence should be considered a *pragmatic* action, rather than a lexical or computational one[8].

Structural underspecification: unfixed node relations

Given the vocabulary DS incorporates from LOFT, many of the structural constraints in the logical form of an utterance boil down to the position a node occupies. Thus, one of the ways in which linguistic underdetermination could be modelled is by rendering tree node addresses *underspecified*. This move is able to capture situations in which an expression occurs at a certain point in the linear order of the utterance, while its truth-conditional contribution is delayed until a compatible tree node is found (movement like effects). This allows DS to model dislocation phenomena that are usually captured in terms of movement by more traditional frameworks, as, e.g. topic-dislocated structures. This form of structural underdetermination is captured by imposing a requirement on the *Tn* decoration, that is, a tree node with the requirement $?\exists\, \mathbf{x}.\ Tn(\mathbf{x})$ is an *unfixed* node. In the tree diagrams, unfixed nodes are represented by a dotted line, such as in the tree schema below:

(26)

$$Tn(0),...$$

$$<\uparrow_*> Tn(0), ?\exists\, x.Tn(\mathbf{x}),.. \qquad Tn(00),... \qquad Tn(01),...$$

---
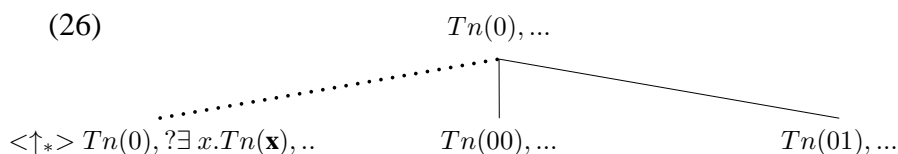
[8]Except in cases of grammaticalisation of the process or the resolution of anaphors, like 'himself', that depend on antecedents found at a particular position in the structure; see Purver et al. 2010.

Unfixed nodes are introduced by the *Adjunction* (star adjunction) computational rule, which is formally defined as:

(27)   *Adjunction*

$$\frac{\{\{Tn(a)..., ?Ty(t), \Diamond\}\}}{\{\{Tn(a)..., ?Ty(t),\}\{<\uparrow_*> Tn(a)..., \exists\,\mathbf{x}.Tn(\mathbf{x}), ?Ty(e), \Diamond\}\}}$$

The rule states that an unfixed node can be introduced, under the assumption its position resides somewhere under the root node. In addition, it requires that this node should eventually find a fixed position in the tree (which will replace the metavariable in $\exists\,\mathbf{x}.Tn(\mathbf{x})$. It also requires the formula value on this node to have the logical type e, thus being an alternative starting point to the rules of *Introduction and Prediction*.

The parse of an utterance of 'John, Mary likes' begins by the application of the rule to the axiom (the initial $?Ty(t)$ requirement), resulting in an unfixed node containing the decoration $<\uparrow_*> Tn(0), ?Ty(e), ?\exists\, x.Tn(\mathbf{x}), \Diamond$. At this stage, the lexical material of 'John' (see entry in (13)) can be parsed and the unfixed node is expanded into $<\uparrow_*> Tn(0), ?Ty(e), Ty(e), ?\exists\, x.Tn(\mathbf{x}), Fo, (John')\Diamond$. *Thinning* removes one of the requirements at this node. *Completion* moves the pointer to the root and records requirement-satisfaction there. With the pointer at the root, *Introduction* and *Prediction* can create the subject and predicate nodes (e.g. *Tn(00)* and *Tn(01)* above). The lexical actions of 'Mary' could then be parsed, and the structure would look like:

(28) $Tn(0), ?Ty(t), <\downarrow_*> Ty(e), ? <\downarrow_0> Ty(e), ? <\downarrow_1> Ty(e \rightarrow t)$

$Tn(00), Ty(e), Fo(Mary')\Diamond$ $Tn(01), ?Ty(e \rightarrow t)$

$\vdots$

$<\uparrow_*> Tn(0), Ty(e), ?\exists\, x.Tn(\mathbf{x}), Fo(John')$
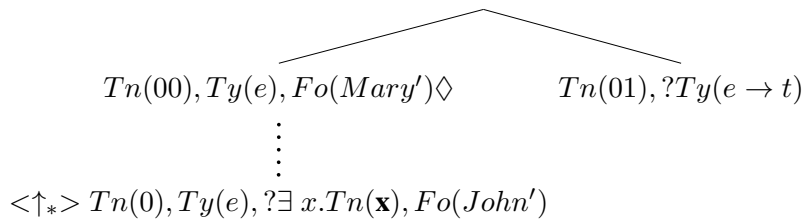
The dotted line under the subject node is a consequence of *The Normal Form Constraint* in Kempson et al. (2001) (see also Gregoromichelaki 2006, 85-6). This means that nodes bearing underspecified modal relations (as introduced by *Adjunction*) will appear as low as possible in the tree, thus making each argument position in the linear order of the parse available as a possible site for node-fixing. Now in the tree above, the unfixed node could be unified at the subject node, but

this would render the parsing of 'Mary' in 'John, Mary likes' impossible (we would be in a predicate node by the time 'Mary' is processed), hence ruling out this possible sequence of actions. The delay in the unification of the unfixed node is what allows 'Mary' to be parsed in subject position. Given the node is type completed, the pointer can be moved to the root (via *Completion*) and then to the predicate node (via *Anticipation*), where information from the predicate can be parsed.

With the pointer at the node containing a type $e \rightarrow t$ requirement, the system parses the verb 'likes'. Its lexical actions (see Cann et al., 2005, 48) include the creation of another argument node ('like' is transitive) and a functor node of type $(e \rightarrow (e \rightarrow t))$ with formula value $(\lambda x. \lambda y. like'(y)(x))$, association is to the right, if parentheses are omitted. The lexical actions also move the pointer to the internal argument node with a $?Ty(e)$ requirement and the dangling unfixed node with it.

(29) $Tn(0), ?Ty(t), ? <\downarrow_1> Ty(e \rightarrow t), <\downarrow_*> Ty(e), <\downarrow_0> Ty(e)$

$$Tn(00), Ty(e), Fo(Mary') \qquad\qquad Tn(01), ?Ty(e \rightarrow t)$$

$$Tn(010), ?Ty(e) \qquad\qquad Tn(011), Ty(e \rightarrow e \rightarrow t)$$
$$\diamondsuit \qquad\qquad\qquad Fo(\lambda x. \lambda y. Like'yx)$$
$$\vdots$$
$$<\uparrow_*>, Fo(John'), Ty(e), ?\exists x.Tn(\mathbf{x}), ?Ty(e)$$
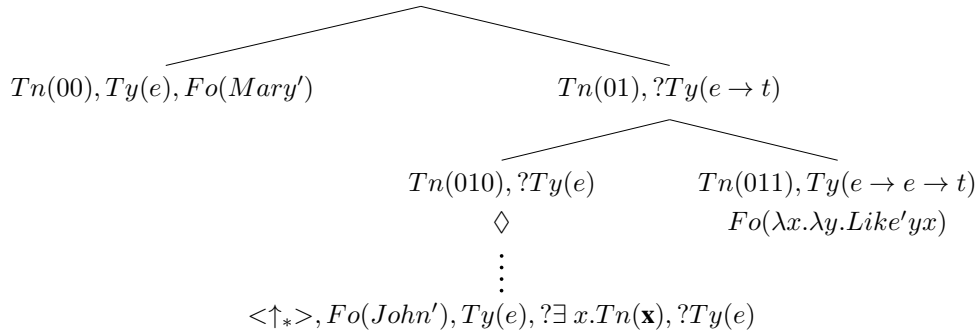
At this point, the requirements of the unfixed node match those of the node from which it hangs. The computational rule *Merge* can then be used to fix the node at this position. *Merge* basically unifies an unfixed node with a compatible position in a tree. More formally, the rule states the following where 'ND' stands for an arbitrary node description:

(30)   *Merge*

$$\frac{\{ND..., ND'...\}}{\{ND \sqcup ND'\}}$$
where $\diamondsuit \in ND'$ and $ND \cup ND'$ is compatible.

Above in (29), ND would be the position of the inner argument node $Tn(010)$ and ND' would stand for the unfixed node hanging below it. Given that these node descriptions are compatible, *Merge* can apply, and we have an internal argument node for the transitive 'like'. Successive applications of *Thinning, Completion* and

Elimination compile the remainder of the tree[9].

5.2.2 Quantification in DS

In the previous subsection, one of the ways in which content underdetermination can be modelled was done in terms of metavariable introduction. The assumption that pronouns introduce values that are yet to be specified is the first step towards an adequate account of identification-based descriptive pronouns. The second step in this direction consists in a non-standard view on quantification, described by the epsilon calculus (Hilbert and Bernays, 1939). In a nutshell, it imports two important features into the framework: (i) a single logical type for quantificational and singular terms, and (ii) a characterisation of quantifiers that allows them to reflect their surrounding environment (via the application of certain computational rules). Given the complexity of the topic, I will focus the presentation on aspects of the framework that are relevant to the interpretative properties of descriptive pronouns.

The epsilon calculus adds two new operators to predicate logic: the epsilon operator - ϵ (closely associated to $\exists$) - and the tau operator τ (closely associated with $\forall$). According to the DS implementation of the calculus, quantifiers have a tri-partite structure: (i) a variable x is combined with (ii) a nominal predicate, e.g. Man', forming the restrictor of the quantifier $(x, Man'(x))$, which is then added to (iii) the binder (i.e. an ϵ or τ operator binding the variable). The structure of a quantifier like 'Some man' would be: $(\epsilon, x, Man'(x))$.

Type-logically, quantifier terms in DS are of type $Ty(e)$. This may seem counterintuitive as utterances containing expressions, such as 'Every' or 'No' do not seem to be talking about anyone in particular. However, note that the grammars of almost all languages do not distinguish between quantificational and other NPs. To capture this uniformity, on the one hand, with a suitable semantics, on the other, DS employs epsilon terms. A term like $\epsilon, x, Man'(x)$ stands for an arbitrary entity that witnesses the property denoted by the restrictor (Man'). Eventually these initially incomplete terms will be augmented with further predications contributed

---

[9]As an alternative to *Introduction* and *Prediction*, Cann (2011) proposes that the parse of English utterances proceeds via an unfixed node which accommodates the argument introduced by the subject and then moves to the verb which creates the whole argument structure (including slots for situations/eventualities) from a $?Ty(t)$ requirement. I will remain neutral with regards to these proposals for the moment. Moreover, although the *Adjunction* rule comprises a general mechanism for explaining dislocated structures, it seems to be ill suited for capturing right dislocation, which requires immediate node fixing. This is done by another version of the rule (Cann et al., 2005), namely: *Local Star Adjunction*
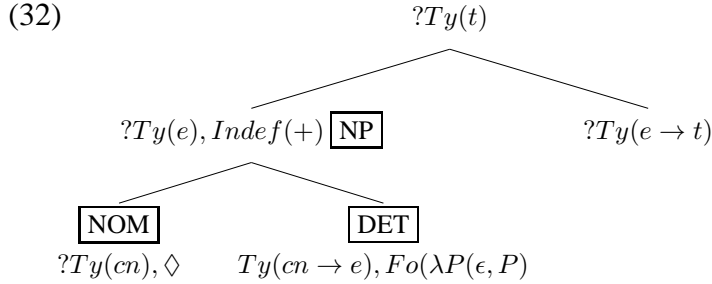
$$\{\{Tn(n)\}, ..\{Tn(m), <\uparrow_*> Tn(n), Ty(X), \Diamond\}, ...\}$$

$$\{\{Tn(n)\}, ..\{Tn(m), <\uparrow_*> Tn(n), Ty(X)\}, \{<\uparrow_*> Tn(m), ?Ty(X), ?\exists x .Tn(x), \Diamond\}...\}$$

by, e.g. verbs and incorporated in their restrictor (see below). Thus, through the introduction of appropriate operators relating restrictor and nuclear scope (predicate) at a final stage in the parse, the usual relations between sets A and B imposed by $\forall A\ B$ and $\exists A\ B$ can be captured: if an entity witnesses A, it also witnesses B ($\tau, x, A(x) \to B(x)$) and for some entity that witnesses A it also witnesses B ($\epsilon, x, A(x) \wedge B(x)$). The epsilon calculus thus imposes the equivalences $\exists x\ P(x) \leftrightarrow P(\epsilon, x, P(x))$ and $\forall x\ P(x) \leftrightarrow P(\tau, x, P(x))$, so we end up with a conservative but more expressive version of the predicate calculus[10].

This picture of quantification becomes clearer when we consider an example. Let us assume (following Gregoromichelaki 2006 and Chatzikyriakidis 2010) that the parse of (31) begins by *Introduction* and *Prediction* and that that the pointer is at the subject $?Ty(e)$ node (I am abstracting from the introduction of the situation node for the moment). Suppose now that the speaker utters:

(31)   A nurse saw every patient.

The lexical actions associated with 'a' decorate the $?Ty(e)$ daughter of the root with a $Indef(+)$ feature and create the binder $Ty(cn \to e)$ functor node bearing the formula value $\lambda P.(\epsilon, P)$ and an argument node with a $?Ty(cn)$ requirement (where the pointer lies)[11]. Basically, from the argument daughter, the determiner head induces the following structure:

(32)

$$?Ty(t)$$

$?Ty(e), Indef(+)\ \boxed{\text{NP}}$ $\qquad\qquad ?Ty(e \to t)$

$\boxed{\text{NOM}} \qquad\qquad \boxed{\text{DET}}$

$?Ty(cn), \Diamond \qquad Ty(cn \to e), Fo(\lambda P(\epsilon, P)$

---

[10]The semantics is defined via the introduction of a choice function Φ to the model $< D, I >$ of first order predicate logic, where D is the domain of individuals and I is a function that maps basic elements in the language onto elements in the domain. Φ selects arbitrary individuals in subsets of D: the witnesses (type $Ty(e)$) that stand for the relevant sets. Thus, $[\![\epsilon, x, Px]\!]^{\mathcal{M}, g} = \Phi([\![P]\!]^{\mathcal{M}, g})$, where Φ is the function given by $\mathcal{M} = < D, I, \Phi >$. It is assumed that restrictors that pick no individuals in the domain denote any arbitrary witness. For further discussion, see Gregoromichelaki (2006, 102).

[11]Here are the lexical actions for 'a'

IF	$?Ty(e)$
THEN	$\texttt{put}\ (Indef(+)); \texttt{make}(<\downarrow_1>); \texttt{go}(<\downarrow_1>); \texttt{put}(Fo(\lambda P(\epsilon, P)); Ty(cn \to e)$
	$\texttt{go}\ (<\uparrow_1>); \texttt{make}(<\downarrow_0>); \texttt{go}(<\downarrow_0>); \texttt{put}(?(Ty(cn)))$
ELSE	abort

The labels in the boxes aim to give a rough description of the noun phrase head, nominal, and determiner nodes. With the pointer at the NOM node, the common noun 'nurse', can be parsed. Its lexical actions are responsible for creating an argument $Ty(e)$ daughter with a fresh variable and a functor $Ty(e \to cn)$ daughter containing the formula $Fo(\lambda x.x, Nurse'x)$ [12]. The lexical actions unfold the current sub-tree into:

(33)

$$?Ty(t)$$

$?Ty(e), Indef(+), ?Sc(z)$ $\boxed{\text{NP}}$ $\qquad ?Ty(e \to t)$

$\boxed{\text{NOM}}$ $\qquad\qquad \boxed{\text{DET}}$
$?Ty(cn)$ $\qquad Ty(cn \to e), Fo(\lambda P.(\epsilon, P))$

$\boxed{\text{VAR}}$ $\qquad \boxed{\text{RESTR}}$
$Ty(e)$ $\qquad Ty(e \to cn)$
$Fo(z), \Diamond$ $\quad Fo(\lambda x.x, nurse'x)$
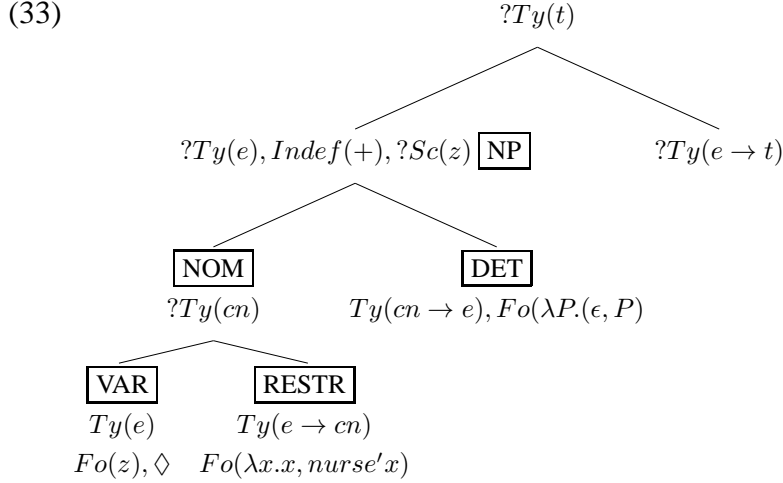
As before, the boxes represent labels for the variable and restrictor nodes. The pointer now rests at the $Ty(e)$ (VAR) node. Applications of *Elimination, Completion,* and *Thinning* compile the subtree until the NP node where the $?Sc(z)$ requirement lies[13]. At this node, the computational action for indefinites can kick in (from Gregoromichelaki 2006)[14].

---

[12] Here are the lexical actions for the common noun 'nurse' (`freshput` checks for the presence of a variable at the VAR node and enters that variable into the scope statement, mimicking the introduction of 'discourse referents'. For details, see Gregoromichelaki (2006, 110-1):

IF	$?Ty(cn)$
THEN	$\texttt{make}(<\downarrow_1>); \texttt{go}(<\downarrow_1>); \texttt{put}(Fo(\lambda x.x, nurse'x)); Ty(e \to cn)$
	$\texttt{go}(<\uparrow_1>); \texttt{make}(<\downarrow_0>); \texttt{go}(<\downarrow_0>); \texttt{freshput}(\mathbf{x}), Fo(\mathbf{x}); \texttt{put}(Ty(e))$
	$\texttt{go}(<\uparrow_0>); \texttt{go}(<\uparrow_0>); \texttt{put}(?Sc(\mathbf{x})); \texttt{go}(<\downarrow_0>); \texttt{go}(<\downarrow_0>)$
ELSE	abort

[13] The predicate *Sc* has the following definition $Sc(a) =_{df} (<\downarrow_0><\downarrow_0> Fo(a) \wedge Ty(e) \wedge <\uparrow * > (Ty(t) \wedge \exists y (Scope(a < y) \vee Scope(y < a)))$. Intuitively, this means that a variable must appear as a formula value of a VAR node and it also must be dominated by a type *t* node containing a scope statement relating that variable to some other variable (scope-taking element) in the tree.

[14] Scope actions for indefinites

IF	$Indef(+), ?Sc(\mathbf{x})$
THEN	$\texttt{gofirst}(?Ty(t)); \texttt{put} Scope(\mathbf{U} < \mathbf{x}),$
	$?\exists y(DOM(\mathbf{y}) \wedge Scope(y < x) \wedge \forall r(Scope(y < r) \to Scope(\mathbf{x} < \mathbf{r})))$
ELSE	abort

where *DOM* is defined as: $Tn \models_m DOM(a) \leftrightarrow Tn \models_m ?Ty(t) \wedge [(\downarrow\downarrow_* [Ty(e) \wedge <\downarrow_0><\downarrow_0> Fo(a)]) \vee (<\downarrow_0> Fo(a))]$.
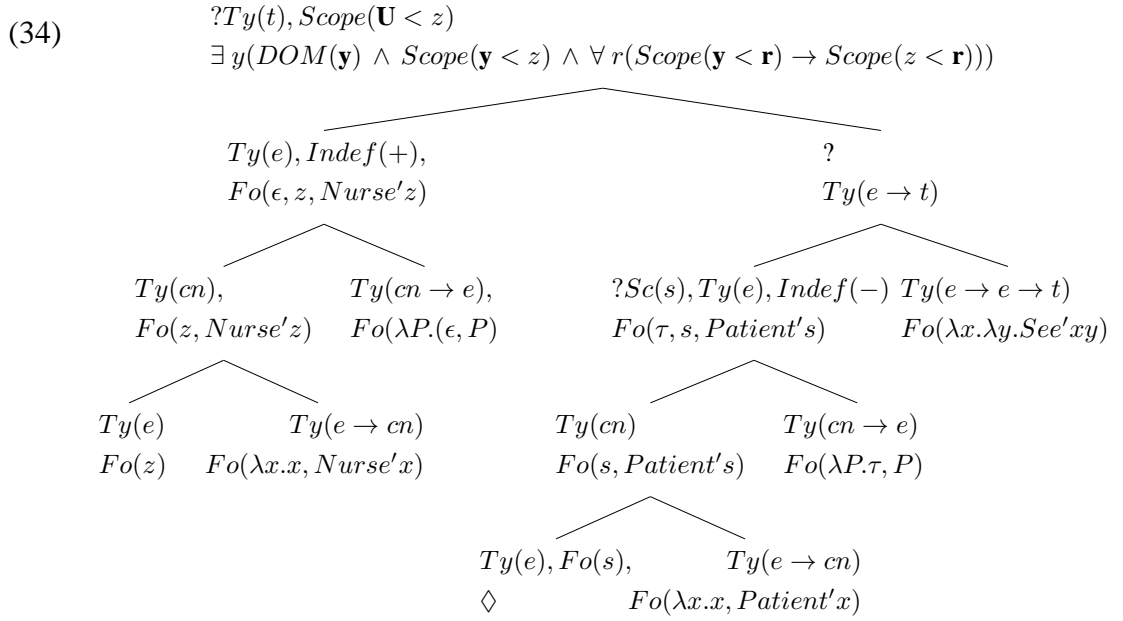
The actions check whether the pointer is at a node containing a scope requirement plus an indefinite lexical feature. If this is the case, they move the pointer to the first node containing a type *t* requirement and enter a scope statement containing the variable introduced at the (VAR) node. In this case, the decoration $Scope(\mathbf{U} < z)$ is inserted at the root, where $\mathbf{U}$ is a metavariable to be replaced by another scope taking element that occurs in the structure. The scope statement determines the order according to which the restrictor and nuclear scope will be unpacked (via computational rules to be introduced shortly). The link between a lexical feature (*Indef(+)*) and the determination of the scope statement allows for certain expressions, like indefinites, to induce different readings. In the specific case of indefinites, the associated lexical feature imposes conditions (via the *DOM* predicate) on the substitution of the metavariable $\mathbf{U}$ that allow for binders that occur later in the structure to precede the variables bound by the indefinite in the scope statement (e. g. scope reconstruction effects). Since this topic transcends the aims of the thesis, I will leave it aside. In summary, the actions discussed here introduce a scope statement $Scope(\mathbf{U} < z)$ at the root, containing a metavariable whose value is yet to be determined[15] and satisfying the scope (*?Sc(z)*) requirement at the NP node.

At this point, the pointer can move (after *Completion*) to the mother node and then (via *Anticipation*) to the functor $?Ty(e \rightarrow t)$ node. From this node, the lexical actions of the verb 'see' (abstracting away from tense) create an (internal) argument daughter (where the pointer rests) and a further functor daughter, decorated by $Ty(e \rightarrow e \rightarrow t), Fo(\lambda x.\lambda y.See'yx)$. With the pointer at the internal argument node, the material from the second NP, namely, 'every patient', can be parsed. The quantifier head 'every' has lexical actions quite similar to 'a' with the exception that it contributes $Indef(-)$ to the NP node and $\lambda P.\tau, P$ to the DET node[16]. Like 'a', the universal determiner also creates the $?Ty(cn)$ argument daughter (where the pointer resides). At this point, the common noun 'patient' is parsed. It creates a $Ty(e)$ (VAR) node, where a fresh variable is inserted and a $Ty(e \rightarrow cn)$ node, where the value $Fo(\lambda x.x, Patient'x)$ (see lexical entry for 'nurse'). The resulting tree is (I am omitting the boxed labels introduced previously):

---

[15]This form of content underdetermination enables scope to be underdetermined.

[16]*Lexical actions for 'every'*

IF	$?Ty(e)$
THEN	put $(Indef(-))$; make($<\downarrow_1>$); go($<\downarrow_1>$); put($Fo(\lambda P(\tau, P))$; $Ty(cn \rightarrow e)$
	go $(<\uparrow_1>)$; make($<\downarrow_0>$); go($<\downarrow_0>$); put($?(Ty(cn))$)
ELSE	abort

(34)
$$?Ty(t), Scope(\mathbf{U} < z)$$
$$\exists\, y(DOM(\mathbf{y}) \wedge Scope(\mathbf{y} < z) \wedge \forall\, r(Scope(\mathbf{y} < \mathbf{r}) \rightarrow Scope(z < \mathbf{r})))$$

Tree structure:

$$Ty(e), Indef(+),$$
$$Fo(\epsilon, z, Nurse'z)$$

$$?$$
$$Ty(e \rightarrow t)$$

$$Ty(cn), \quad Ty(cn \rightarrow e),$$
$$Fo(z, Nurse'z) \quad Fo(\lambda P.(\epsilon, P))$$

$$?Sc(s), Ty(e), Indef(-) \quad Ty(e \rightarrow e \rightarrow t)$$
$$Fo(\tau, s, Patient's) \quad Fo(\lambda x.\lambda y.See'xy)$$

$$Ty(e) \quad Ty(e \rightarrow cn)$$
$$Fo(z) \quad Fo(\lambda x.x, Nurse'x)$$

$$Ty(cn) \quad Ty(cn \rightarrow e)$$
$$Fo(s, Patient's) \quad Fo(\lambda P.\tau, P)$$

$$Ty(e), Fo(s), \quad Ty(e \rightarrow cn)$$
$$\Diamond \quad Fo(\lambda x.x, Patient'x)$$

Successive applications of *Completion, Elimination* and *Thinning* compile the subtree until the NP (type *e*) node of the universal, where the scope requirement $?Sc(s)$ lies. At this point, the scope action for non-indefinites applies[17].

These actions first check for the presence of the $Indef(-)$ feature and the scope requirement involving the variable *s*. If the descriptions hold, the pointer moves to the first $?Ty(t)$ node and inserts a scope statement about the variable *s* modulo the satisfaction of further constraints (e.g. described by the DOM^+ predicate). Intuitively, these further conditions aim to ensure that the scope of determiners with the $Indef(-)$ feature is determined linearly. In the formal definition, the variable of interest could depend only on variables of other $Indef(-)$ determiners or situation variables (not represented in the structure for the sake of simplicity). Let us assume, that this utterance is being evaluated with respect to a situation *v*. Thus, the actions would insert a $Sc(v < s)$ statement at the root, satisfying the requirement at the DP node and allowing the whole tree to be compiled (via applications of *Thinning, Completion* and *Elimination*).

---

[17]*Scope actions for non-indefinites*

```
IF      Indef(−), ?Sc(x)
THEN    gofirst(?Ty(t))
           IF                DOM⁺(y)
           THEN              IF        {∃ zDOM⁺z ∧ Scope(y < z)}
                             THEN      abort
                             ELSE      put(Scope(y < z))
           ELSE     abort
ELSE    abort
```

where DOM^+ is defined as $Tn \models_m DOM^*(a) \leftrightarrow Tn \models_m ?Ty(t) \wedge [(\downarrow\downarrow_* [Indef(-) \wedge <\downarrow_0><\downarrow_0> Fo(a)]) \vee (<\downarrow_0>]Fo(a))]$

(35)
$$Ty(t), Sc(\mathbf{U} < z; v < s), Fo(See'(\tau, s, Patient's)(\epsilon, z, Nurse'z)), \Diamond$$
$$\exists\, y(DOM(\mathbf{y}) \wedge Scope(\mathbf{y} < z) \wedge \forall\, r(Scope(\mathbf{y} < \mathbf{r}) \to Scope(z < \mathbf{r})))$$

$Ty(e), Indef(+),$
$Fo(\epsilon, z, Nurse'z)$

$Ty(e \to t)$
$Fo(\lambda y.See'(\tau, s, Patient's)(y))$

$Ty(cn),$ $Ty(cn \to e),$
$Fo(Nurse'z)$ $Fo(\lambda P.(\epsilon, P))$

$Ty(e), Indef(-)$ $Ty(e \to e \to t)$
$Fo(\tau, s, Patient's)$ $Fo(\lambda x.\lambda y.See'xy)$

$Ty(e)$ $Ty(e \to cn)$
$Fo(z)$ $Fo(\lambda x.x, Nurse'x)$

$Ty(cn)$ $Ty(cn \to e)$
$Fo(s, Patient's)$ $Fo(\lambda P.\tau, P)$

$Ty(e),$ $Ty(e \to cn)$
$Fo(s)$ $Fo(\lambda x.x, Patient'x)$

The only thing needed now is to fully determine the value for the metavariable **U**, so the formula value of the quantificational terms can be unpacked in a way such that the right relations between restrictor and nuclear scope are established. Note that the conditions associated with the (*Indef(+)*) feature allow for **U** to be instantiated by the situation variable *v*, resulting in the reading *there is a nurse who sees every patient*, or by the variable *s* bound by the universal, resulting in the reading *for each patient there is a nurse who sees him*. Let us suppose that in the context the second option has been taken (the indefinite would be outscoped by, or reconstructed below, the universal). Given the formal definition of such conditions (i.e. the *DOM* predicate), the resulting scope statement is $Sc(v < s < z)$ (note that the $<$ relation is transitive and irreflexive). The resulting tree is identical to the previous one with the exception that the decorations at the root have been reduced to: $Ty(t), Sc(v < s < z), Fo(See'(\tau, s, Patient's)(\epsilon, z, Nurse'z)), \Diamond$.

At this stage, the formula value can be re-structured using the quantifier evaluation rule below (differently from the definitions in Kempson et al. 2001 and Cann et al. 2005, I follow the definition in Gregoromichelaki 2006, 198):

(36)   *Q-Evaluation Rule*

$$\frac{\{...\{Ty(t), .., World(w[x_1]), Scope(x_1 < ...x_n), Fo(\phi[vx_n\psi_n/x_n]), ...\}\}}{\{\{Ty(t), .., World(w[x_1]), Scope(x_1 < ...x_n - 1), ..., Fo(f_{v_nx_n\psi[vx_n\psi_n/x_n]}(\phi)), ...\}\}}$$

where for x occurring free in ϕ and $w[x_1]$ = a world variable x_1 or $w[x_1] = v_1x_1\psi$, and the values $f_{v_nx_n\psi[vx_n\psi_n/x_n]}(\phi)$ for $v \in \epsilon, \tau, Q$ and $f_w[x1](\psi)$ are defined by:

a. $f_{\tau x\psi}(\phi) = \psi[a/x] \rightarrow \phi[a/x]$, where a $= \tau x(\psi \rightarrow \phi)$

b. $f_{\epsilon x\psi}(\phi) = \psi[b/x] \wedge \phi[b/x]$, where b $= \epsilon x(\psi \wedge \phi)$

c. $f_{Qx\psi}(\phi) = (\psi[c/x])(\phi[c/x])$, where c $= v_Q x((\psi)(\phi))$

d. $f_{w[x_1]}(\phi) = w[x_1] : \phi$

The rule above takes a node containing a strictly ordered scope statement (for example, $Scope(v < s < z)$) as input and outputs a new scope statement omitting the last variable and a formula value containing a connective appropriate to the binder of that variable which relates the predications contributed by the noun and the verb, i.e. the restrictor and the nuclear scope. In our example, the last variable z is bound by an ϵ operator. According to the definition (b) above, an ϵ term relates its restrictor (in this case 'nurse') to the nuclear scope (the predicate 'see' instantiating ψ in the rule above) via conjunction. The symbol 'b' in the definition (b) represents a name for the arbitrary witness that replaces the variable. Thus, the first application of the rule results in:

(37) $Scope(v < s < \not{z})$, $Fo : Nurse'(b) \wedge See'(b; \tau, s, Patient's)$, where b $= \epsilon, z, (Nurse'z \wedge See'z; \tau, s, Patient's)$.

Note that the equivalence between the name 'b' and the complex value is what ensures that terms in epsilon calculus reflect their surrounding environment (accounting for the incrementality in certain interpretations of quantified antecedents, see Evans 1977, chapter 1 section 1.5)[18]. Now, we must deal with variable s, which is bound by a τ term. As a first step, we must replace all occurrences of $\tau, s, Patient's$ in the formula in (37) by s (note that 'b' is a name containing variables, which will be represented in subscript). The result is:

(38) $Scope(v < s)$, $Fo : Nurse'(b_s) \wedge See'(b_s; s)$

Now, this complex formula will be an instance of ϕ in the scope evaluation rule above. The next step is replace s by the variable name used in sub-item (a) of the *Q-evaluation* rule, namely 'a' ($Fo : Nurse'(b_a) \wedge See'(b_a; a)$) and apply the instruction stated there, the result is:

(39) $Scope(v < \not{s})$, $Fo : Patient'(a) \rightarrow [Nurse'(b_a) \wedge See'(b_a; a)]$

---

[18]If we replace the name for the full formula description, we end up with the equivalent: $Nurse'(\epsilon, z, (Nurse'z \wedge See'z; \tau, s, Patient's)) \wedge See'(\epsilon, z, (Nurse'z \wedge See'z, \tau, s, Patient's); \tau, s, Patient's)$. For this reason, I will use the abbreviations.

where $b_a = \epsilon, z, (Nurse'z \wedge See'z, a)$, where $a = Patient'(s) \rightarrow [Nurse'(b_s) \wedge See'(b_s; s)]$, where $b_s = \epsilon, z, (Nurse'z \wedge See'z; \tau, s, Patient's)]$

The formula above means that if some individual is a patient then there is a nurse who sees him: the intuitive truth-conditions of the utterance. The final variable *v* is removed by the specification (d) in the *Q-Evaluation Rule*. It is designed to take situations or world variables and state that a given formula value holds at some situation/world. Thus, the formula is said to hold at a situation of evaluation v (e.g. the actual world).

This long detour introduced the following elements for an explanation of some identification-based descriptive uses. The account of quantification proposed here analyses quantifiers as type e expressions, whose formula values could replace the metavariables introduced by pronouns (as discussed in the previous subsection). This amounts to a preliminary account of the *Descriptiveness* property. Furthermore, the quantifier evaluation rules above unpack the formula values in a way such that certain dependencies are established. In the formula value above, the nurses co-vary with patients. This mechanism will be recruited for an explanation of the *Co-variation* property. We now turn to the final step towards an account of the data, namely, the introduction of a DS component that frames the contextual assumptions used in the substitution operation just mentioned and allows two tree structures to be connected.

5.2.3 LINKed trees

Since our grammar framework aims to capture linguistic phenomena that is not strictly confined to string-boundaries, we must introduce a mechanism that is able to capture cross-sentential dependencies. This is done in terms of LINKed representations. Roughly, a LINKed structure means that two tree-structured representations are connected in virtue of the fact that they have a common truth-conditional constituent. More formally, LINKed trees amount to an extension of the modal vocabulary of LOFT (Blackburn and Meyer-Viol, 1994). In addition to $<\uparrow><\downarrow>$ modalities, the language DU employs $< L >$ and $< L^{-1} >$ relations. The decoration $< L > Tn(m)$ means that tree node $Tn(m)$ can be found one step across the LINK relation, whereas $< L^{-1} > Tn(n)$ means that tree node $Tn(n)$ can be found one step across the *inverse* LINK relation (one step backwards through the LINK relation). This 'connection' is established via the sharing of type e terms, as stated in the rule below.

(40)   *Link Adjunction*

$$\frac{\{\{Tn(X)..., Ty(e), Fo(a)..., \Diamond\}...\}}{\{\{Tn(X)..., Ty(e), Fo(a)\}\{< L^{-1} > Tn(X), ?Ty(t), ? <\downarrow^*> Fo(a), \Diamond\}..\}}$$

The use of *Tn(X)* makes it possible for the rule to apply at any node of type *e* with a formula value *Fo(a)*. The application of the rule at a type e node effectively creates a new LINKed tree structure with the requirement that the formula value *Fo(a)* must occur somewhere in it. For purposes of illustration, let us suppose that the hearer is parsing an utterance of 'John, who smokes, runs'. After *Introduction* and *Prediction*, the lexical material of 'John' can be parsed at the subject $Ty(e)$ node. At this point, the LINK adjunction rule can apply, generating the following structure:

(41)

$$?Ty(t)$$

$$Tn(n), Ty(e)$$
$$Fo(John')$$
$$< L > Tn(m)$$

$$?$$
$$Ty$$
$$(e \rightarrow t)$$

$$< L^{-1} > Tn(n), ?Ty(t), ? <\downarrow^*> Fo(John), \Diamond$$
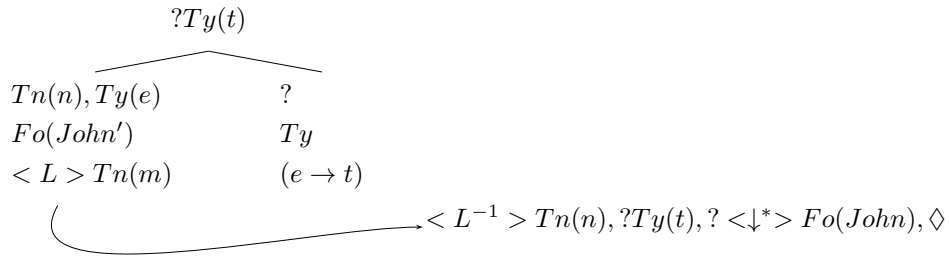
Given that the next word in the parse is 'who' and wh-relatives are usually involved in the interpretation of displaced (i.e. moved) constituents, we can apply the *∗ Adjunction* rule, which creates an unfixed node from a type *t* requiring node. The lexical actions of the relative 'who' [19] take the pointer at an unfixed node and recover the formula value at the type e node from which the LINK originates. This guarantees that 'John' and 'who' refer to the same person.

(42)

[19]Lexical actions for 'who$_{rel}$'

IF	$?Ty(e), ?\exists x Tn(\mathbf{x}), <\uparrow^*><L^{-1}> Fo(\mathbf{x})$
THEN	put $Fo(\mathbf{x}), Ty(e), [\downarrow]\bot$
ELSE	abort

$$?Ty(t)$$

$$Tn(n), Ty(e)$$
$$Fo(John')$$
$$< L > Tn(m) \qquad\qquad ?Ty(e \to t)$$

$$< L^{-1} > Tn(n), ?Ty(t)\text{, } ?\text{<}\!\!\downarrow^*\!\!\text{>}\text{ }Fo(John)$$

$$\vdots$$

$$<\uparrow^*><L^{-1}> Tn(n), \text{?}Ty(e), Ty(e), ?\exists x\, Tn(\mathbf{x}), Fo(John'), \Diamond$$
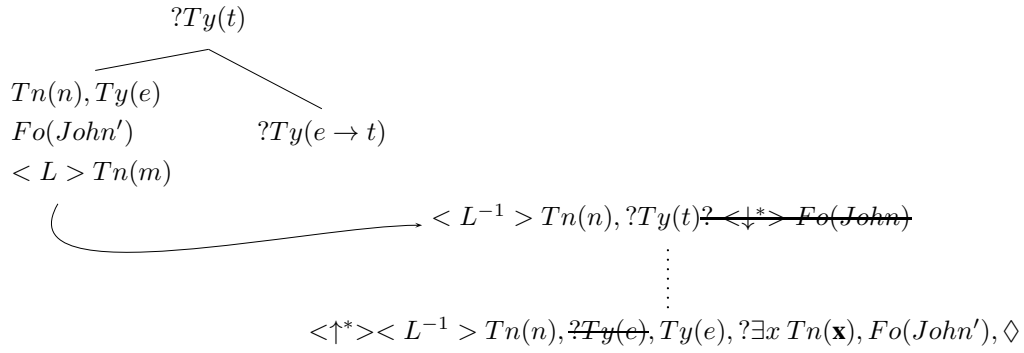
The requirements that have been met by the lexical actions of the relative are then removed via *Thinning* (shown by strikethrough above). The pointer then moves to the root node of the LINKed daughter and the rules of *Introduction* and *Prediction* build the subject and predicate nodes. *Merge* fixes the unfixed node as the subject, satisfying the $?\exists x\, Tn(\mathbf{x})$ requirement. The pointer is then moved to the predicate node of the relative clause and the verb 'smokes' is parsed. Applications of *Thinning, Completion* and *Elimination*, compile the reminder of the tree.

(43)

$$?Ty(t)$$

$$Tn(n), Ty(e) \qquad\qquad ?$$
$$Fo(John') \qquad\qquad\quad Ty$$
$$< L > Tn(m) \qquad\quad (e \to t)$$

$$< L^{-1} > Tn(n), ?Ty(t)$$

$$<\uparrow^*><L^{-1}> Tn(n), \qquad\qquad Ty(e \to t),$$
$$Ty(e), Fo(John') \qquad\qquad Fo(\lambda x.Smoke'x), \Diamond$$

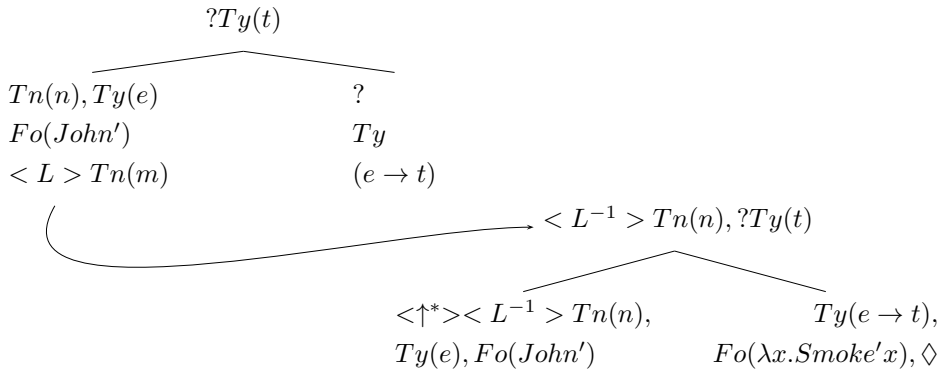Steps of *Thinning, Completion,* and *Elimination* complete the LINKed daughter and the pointer is moved through the inverse link relation to the subject of the LINKed mother tree. *Completion* and *Anticipation* move the pointer to the predicate node of that tree, where the verb of the main clause ('runs') can be parsed. After the type and formula values are inserted, repeated applications of *Thinning, Completion,* and *Elimination* compile the LINKed mother. As a result, we have two complete trees: one, from which the LINK originates, having formula value $Fo(Run'(John'))$; another (the LINK daughter) having formula value $Fo(Smoke'(John'))$.

At this point, one can apply an evaluation rule that combines the formula values of the two structures. In the case of non-restrictive relatives, the rule is:

(44)   *LINK Evaluation 1 (non-restrictives)*

$$\frac{\{\{Tn(n)...,Ty(t),Fo(a)...,\Diamond\}\{<L^{-1}>MOD\ Tn(a),Ty(t),Fo(b)\}...\}}{\{\{Tn(n)...,Ty(t),Fo(a)\wedge Fo(b)...,\Diamond\}\{<L^{-1}>MOD\ Tn(a),Ty(t),Fo(b)\}...\}}$$

MOD $\in\{<\uparrow_0>,<\uparrow_1>\}^*$

The result of the application of this rule takes the formula values of the root nodes of the two LINKed trees, namely, $Fo(Smoke'(John'))$ and also $(Run'(John'))$, and outputs the conjunction of those values at the root node of the LINKed mother tree, that is, $Fo((Smoke'(John'))\wedge(Run'(John')))$. Note that the optional nature of the rule can account for the intuition that in some cases, non-restrictive relatives communicate two independent (not conjoined) propositions (see Bach 1999, Neale 1999, Carston 2002, 131, fn. 26, and Gregoromichelaki 2006, 153 fn. 49)[20].

In the example above, the lexical actions of 'who' played an important role in guaranteeing that the wh-relative and 'John' refer to the same entity. However, the actions of relatives are not the only way through which formula values can be shared. For example, the discourse fragment 'John went to the shop. He bought fish' would be explained by the very same mechanism of LINKed structures. The LINK relation would originate from the node containing $Fo(John')$, providing an environment (i.e. a context) for the parse of the second string. Given that the LINK relation requires the sharing of terms and that pronouns introduce metavariables, the anaphoric pronoun's value would be $Fo(John')$ (via *Substitution*).

This suggests that the mechanisms of LINKing comprise a very general tool for modelling a plethora of phenomena, including contextual dependencies where antecedents are provided linguistically, i.e. donkey anaphora. This comes out more clearly when we consider tree structures which are slightly more complex than the ones displayed so far. As mentioned in the subsection on quantification, DS trees also have nodes corresponding to situations/worlds (left unrepresented for simplicity's sake) to capture the intuition that a formula value describes a situation or eventuality, which following Kratzer (1989) are taken as parts of a world (of evaluation). They minimally include time (subject to tense modification) and place parameters, but might also include other eventuality aspects. Gregoromichelaki (2006) pro-

---

[20]*Restrictive relatives* are explained by a very similar mechanism, according to which the the formula value of the relative clause is incorporated as a restrictor on the variable via the following rule. *LINK Evaluation 2 (restrictives)*

$$\frac{\{...\{Tn(Z),?Ty(cn)...\}...\{<\uparrow_0>Tn(Z),Ty(e),Fo(\mathbf{x})\}}{\{<\uparrow_1>Tn(Z),Ty(cn\to e),Fo(\lambda z.zPz))\},\{<L^{-1}><\uparrow_0>Tn(Z),Ty(t),Fo(Q)\Diamond\}...\}}$$

$$\frac{\{...\{Tn(Z),Ty(cn),Fo(\mathbf{x},P\mathbf{x}\wedge Q\mathbf{x}),\Diamond\}...\{<\uparrow_0>Tn(Z),Ty(e),Fo(\mathbf{x})\}}{\{<\uparrow_1>Tn(Z),Ty(cn\to e),Fo(\lambda z.zPz))\},\{<L^{-1}><\uparrow_0>Tn(Z),Ty(t),Fo(Q)\}...\}}$$

220

poses that, with respect to their logical type, situations are a subtype of type e (see Schlenker 2006, for a common semantic treatment of individuals and situations), indicated by a subscript s (i.e. $Ty(e_s)$) for presentation purposes. As such, situation arguments can be taken care of by any computational rule that accounts for the processing of noun phrases, thus providing a unified account of adverbial and argument modification. For our purposes here, situation ($Ty(e_s)$) nodes can be the inputs of the LINK adjunction rule, such that two or more formula values can be LINKed by the fact that they describe the same situation. Thus, the relevance theoretic assumption that context is a set of mutually manifest assumptions can be modelled as a context set $C =_{def} < T_1, T_{...}, T_n, >$, i.e. a series of structured propositions, against which a tree under construction T is interpreted (Cann et al., 2005, ch 9.2.1). The trees could be linked via a situation argument or by any usual type e node as indicated in the anaphora case mentioned above. Extended tree-structures containing situation nodes will be employed for the explanation of the data proposed in the next section (5.3.2).

For the purpose of accounting for the descriptive pronominal data, the contextually salient set of LINKed trees can be regarded as an environment which provides a term that replaces the metavariable introduced by the descriptively used pronoun. The only difference is that the term is descriptive. This will be developed for some instances of descriptive uses in the next section. Before doing so, I would like to mention yet another way in which context may be structured. DS also assumes that a record of the actions used to build the parse is stored in the system and may be recalled for the interpretation of certain utterances, as for instance, elliptical ones. In fact, a context containing a certain series of LINKed trees can be seen as isomorphic to a context containing a series of actions involved in the building of such trees (Kempson et al., 2011).

5.3 The proposal

In this section, I will detail how the relevance theoretic pragmatic framework (sketched in chapter 2), the incremental notion of processing grounded on it (proposed in chapter 3) and the particular grammar model (sketched above) come together in a unified explanation of the data.

The discussion in chapter 2 established that identification-based descriptive uses constitute a heterogeneous phenomenon. Accordingly, I will approach the data using a split account. The basic idea is this. Some uses are captured as instances of relevance driven implicatures, as sketched in chapter 2 section 2.3. Other descrip-

tive uses contribute to the utterance's explicature(s). For the latter case, the grammar formalism detailed here can do much explanatory work. First, it models context as a set of LINKed trees that provide an environment for the interpretation of the pronoun. As pronouns introduce metavariables, these can be replaced by conceptual representations from the context set. The grammar allows metavariables to be replaced by either quantified or singular formula values, since both are of type *e*. Finally, the tools for modelling context can be extended to capture certain perspectival aspects of interpretation (via the assumption that grammatical person presupposes certain perspectival instances), hence explaining why certain representations are not available for *Substitution*. The resulting picture can explain all the data in chapter 1 (summarised in section 5.1 of this chapter) without facing the shortcomings of a variety of theories, as reviewed in chapter 4, and using very general pragmatic and grammatical mechanisms: both independently required for explaining a plethora of linguistic phenomena.

I begin my proposal by accounting for the *core* properties of the third person singular pronouns and one of the *extended* properties, namely, *Co-variation*. I then move to an analysis of the other pronominal forms. Afterwards, I put forward an account of the *Accessibility* property of the data, which partially relies on the use of multiple pronouns with distinct person features.

5.3.1 Third person pronouns: the core properties

As reviewed in the introduction of this chapter, the basic properties that need explaining are *Descriptiveness* - the truth-conditions of utterances containing the relevant pronominal use are descriptive - *Identification Dependency* - the use of the pronoun requires the identification of an individual as falling under a concept - and *Connection* - the conceptual information used in the identification of the individual plays a role in the derivation of the right descriptive content. Let us consider a modified version of the famous footprint example (Schiffer, 1981).

(45)   A park ranger, is looking for trespassers in a forbidden part of the park. Her colleague, Andy, finds something on the ground. Pointing at it, he utters:

   a.   'He is gigantic'.

Following Relevance Theory (Sperber and Wilson 1995; Carston 2002; *inter alia*), I assume that the audience makes certain hypotheses about speaker-meaning based on the ostensive acts performed. For purposes of clarity, I will analyse the example above as two related ostensive acts: *demonstrating* something and *uttering*

something. Based on the incremental picture of processing sketched in chapter 3, I assume that the comprehension of the demonstration involves the following representational stages (justifications in square brackets).

(46)   *The pragmatics of object representation*

    a. $_i$ (bare object representation, i.e. visual index) [audience's singling out of an object in the environment based on cognitive constraints that, if met, output representations in a bottom-up data-driven way].

    b. $_i(Footprint')^{21}$ [conceptual representation that the visual index in (a) gives rise to, given the assumption that cognition is geared towards the maximisation of cognitive-effects/processing-costs ratio (Cognitive Principle of Relevance) and that *conceptuality* is an important cognitive effect (see chapter 3 section 3.4.3). *Identification Dependency is met*].

    c. $Make-a-Footprint'(\epsilon, z, Person'z), Fresh'(\epsilon, y, Footprint'y),$ $Close-By'(\epsilon, x, \underline{Footprint-Maker'(x)})$ [propositions made salient in virtue of the spread of activation from the concept in (b) and top-down influences arising from the audience's current expectation of relevance; licensed by the Communicative Principle of Relevance].

The propositions in (46)c are in the set of mutually manifest assumptions: the context against which the speaker's utterance is interpreted. Now, the decoding of Andy's utterance follows stages similar to the ones discussed in chapter 2 section 2.3, but it also incorporates a DS account of the incomplete logical form of an utterance that is developed into an explicature (see the definition of *Explicitness* in chapter 2 section 2.3.2). Here is a summary version of the comprehension process:

(47)   *Interpreting Andy's utterance*

    a. $Gigantic'(\mathbf{U}_{male'})$ [incomplete logical form that results from the decoding of Andy's utterance in (45)].

    b. $Gigantic'(\epsilon, x, \underline{Footprint-maker'(x)})$ [propositional form that results from the substitution of the metavariable in (a) by a formula value in one of the contextual assumptions in (46)c (in underline); warranted by the fact that *the footprint maker is gigantic* results in the optimally relevant propositional content in this context. *Descriptiveness is met*].

---

[21] The earlier way of writing concepts, e.g. FOOTPRINT, has been replaced by the equivalent notation employed by the language DU.

c. $Trespasser'(\epsilon, x, Footprint - maker'x)$, [implications (positive effects) that the representation in (b) gives rise to].

Note that the cognitive processes sketched in (46) and (47) probably happen in parallel and proceed as follows. The incremental view on processing, described in chapter 3 and underwritten by the cognitive principle of relevance, establishes a continuum from early levels of visual processing to a conceptual level which can be integrated with linguistic meaning so the utterance's explicature(s) can be derived, and from there to the level of the utterance's implicatures. Early object representations would surface at a conceptual level at some stage in visual processing, accounting for the classification of the relevant object as belonging to a kind (*Identification Dependency*). This concept, in turn, activates associated propositional information that is fed into the context set, as in (46)c. These assumptions (a set of LINKed trees) are used to develop the encoded meaning of Andy's utterance, i.e. (47)a, into an explicature, i.e. (47)b. This move is similar to the ad hoc concept proposal reviewed in chapter 2 section 2.3.3, in the sense that a set of assumptions is used to devise an occasion specific semantic value of a word, but the fact that the linguistic meaning of pronouns is underdetermined allows for a simpler solution: the assumptions provide a formula value that replaces the metavariable introduced by the pronoun, in a similar way to anaphora (see chapter 1 section 1.5)[22], that is, an instance of the substitution-based account sketched towards the end of that section. Since the replacement term is quantificational, the intuition that the truth-conditions are descriptive is properly accounted for (*Descriptiveness*)[23]. Finally, the inferential relation between the classificatory concept in (46)b and the term that replaces the metavariable (i.e. $(\epsilon, x, Footprint - maker'x)$) in (47)b accounts for the intuition that conceptual classification of the salient object provides the building blocks of the descriptive interpretation (*Connection*). The account successfully captures the *core* data.

At this point, it would be helpful to show how the grammar framework structures the context in such a way that the relevant substitution in this example can occur. Since the demonstration of the tree transitions is also necessary for the explanation

---

[22]Many of the descriptive uses discussed in the thesis have been paraphrased as *definites*. Traditionally, these expressions are taken to assert or presuppose uniqueness, which would not be captured by ϵ operator. Two options are available: to argue that definites do not induce uniqueness in virtue of their linguistic meaning (Breheny, 1999) or to model the data using the DS implementation of iota (ι) terms, which are ϵ operators that presuppose the uniqueness of their satisfiers (as 'he' presupposes its referent is male). For the purposes of this presentation, I follow the first option, but could implement the second, if necessary.

[23]Note that an important feature of the system is what allows this: postulating the same logical type for quantifiers and (allegedly) singular expressions.

of the extended properties of the data, I will leave this for the next section, where the account relies more heavily on grammatical mechanisms and subsumes steps that correspond to the explanation sketched above.

5.3.2 Third person pronouns: *Co-variation*

Although many of the accounts of descriptive uses discussed in chapter 4 *could* explain the co-varying readings between the contextually induced descriptive inter-pretations and other expressions in the utterance, their proponents have not shown how the relevant readings are properly derived[24]. The task upon us is to account for the *Co-variation* property using the grammar tools sketched above. I will base the analysis on the following example:

(48)   Every time there is a war, he [*pointing to the White House*] has tough choices to make.

The utterance in (48) could express a singular proposition about a specific indi-vidual, say, Bill Clinton, to the effect that *he* has tough choices to make for each wartime [25]. However, this is not the reading we are interested in. Additionally, the utterance above could communicate that for every time a war is being waged, the American President at that time has tough choices to make. The adverbial modifier 'every time there is a war' seems to play a crucial role in delivering this reading. Since this type of expression predicates something about a situation/eventuality, we must address now how these entities are incorporated into the DS formalism.

Since situations are modelled as individuals (type e_s expressions), computa-tional rules that operate on argument nodes can equally apply to situations. Follow-ing Gregoromichelaki (2006), let us assume that situation nodes are introduced by the application of *Introduction* and *Prediction* and that predicates encode an (extra) argument slot for the situation argument represented in the structure (although she discusses the possibility of the situation node being introduced freely)[26]. Therefore, from a $?Ty(t)$ node, the tree-structure is expanded into:

---

[24]Elbourne (2008) is an exception to this claim.

[25]This is a simplification, the utterance is subject to a lot of pragmatics that could in principle enrich the meaning of 'war' [involving NATO] or modulate the domain of quantification of 'every time' [the UN decides]. These will be abstracted away here.

[26]Alternatively, Cann (2011) assumes that the situation node is introduced by the lexical actions of verb morphology, which also insert a fresh variable at the situation argument node. My proposal is compatible with both ways of accommodating a situation argument.

(49) $?Ty(t), <\downarrow>_0 ?Ty(e_s), <\downarrow>_1 ?Ty(e \rightarrow t)$

$?Ty(e_s), \lozenge$ $?Ty(e_s \rightarrow t)$
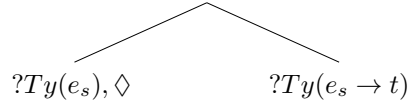
Following Gregoromichelaki (2006), I assume a *Situation Metavariable Introduction* rule, which introduces a situation metavariable whose value should be replaced by an appropriate formula value in the context. The rule has been independently motivated by condit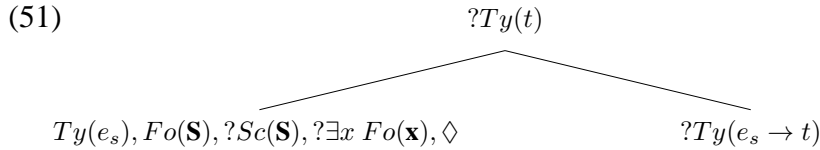ionals and successfully captures the idea that determining the situation described by an utterance constitutes an instance of saturation (see Recanati, 1999).

(50)   *Situation Metavariable Insertion Rule*

$$\frac{\{...Tn(n), ?Ty(t), \}, \{<\uparrow_0> Tn(n), ?Ty(e_s), \lozenge\}, \{<\uparrow_1> Tn(n), ?Ty(e_s \rightarrow t)\}\}...}{\{...Tn(n), ?Ty(t), \lozenge\}, \{<\uparrow_0> Tn(n), Ty(e_s), Fo(\mathbf{S}), ?Sc(\mathbf{S}), ?\exists x\ Fo(x)\}}$$

$$\{<\uparrow_1> Tn(n), ?Ty(e_s \rightarrow t)\}\}...$$

The rule licenses the transition to the following tree:

(51) $?Ty(t)$

$Ty(e_s), Fo(\mathbf{S}), ?Sc(\mathbf{S}), ?\exists x\ Fo(\mathbf{x}), \lozenge$ $?Ty(e_s \rightarrow t)$

This would be the point at which the adverbial 'Every time there is a war' can be parsed. Intuitively, the expression seems to focus on certain temporal aspects of the event described by the main verb. As such, it would be natural to model these expressions as modifiers of the situation argument. Following Kempson (2010) and Gregoromichelaki (2006), adjunction can be modelled as the addition of information to a tree via trees LINKed to it. Since the occurrence of adverbials (modifiers of situations/eventualities) is optional, I propose a *Situation Adjunction* rule that creates a LINKed structure attached to the situation node. This rule enforces identity of formula values between the LINK daughter and the value of the metavariable in the $Ty(e_s)$ node and is based, in part, on the structure induced by 'if' clauses as proposed in Gregoromichelaki (2006, 207-21), that is, what if-clauses do lexically, the *Situation Adjunction Rule* does generally and optionally[27].

(52)   *Situation Adjunction Rule*

---

[27]Ideally, adjuncts like 'every time' would introduce the LINK relation and impose the sharing of terms themselves, making the rule proposed here redundant. This alternative gives equivalent results but is more complicated in terms of presentation. Therefore, it will not be discussed here.

$$\{...Tn(n), ?Ty(t)\}, \{<\uparrow_0> Tn(n), Ty(e_s), Fo(\mathbf{S}),$$
$$?Sc(\mathbf{S}), ?\exists x \ Fo(x), \Diamond\}\{<\uparrow_1> Tn(n), ?Ty(e_s \to t)\}\}...$$

$$\{...Tn(n), ?Ty(t)\}, \{<\uparrow_0> Tn(n), Ty(e_s), Fo(\mathbf{S}), ?Sc(\mathbf{S}), ?\exists \ xFo((\mathbf{x}) \wedge \langle L\rangle Fo(\mathbf{x}))\}$$
$$\{<\uparrow_1> Tn(n), ?Ty(e_s \to t)\}\}, \{< L^{-1} ><\uparrow_0> Tn(n), ?Ty(e_s), ?\exists \ xFo(\mathbf{x}), \Diamond\}...$$

The rule allows for the transition to:

(53)

At this point, the content encoded by 'Every time there is a war' contributes to the LINKed tree. For the sake of simplicity, this complex expression will be treated as deriving the predicate 'Every war-time'. As seen in the previous section, the quantifier head 'every' makes the transition from a type *e* requiring node to a tree structure where this node has a type $cn_s \to e_s$ daughter with formula value $\lambda P.\tau, P$ (DET node) and type cn_s requiring node. The parsing of 'every' results in the following tree:

(54)

At this point the complex expression 'time there is a war' is parsed. As mentioned, since the anaphoric choice to be made does not turn on the specifics of this being a relative clause, for simplicity's sake I will treat the expression as deriving a complex predicate of type cn_s and formula value $Fo(s, WarTime'(s))$, i.e.

227

as projected by *every wartime*. According to this simplified exposition, the complex common noun creates two nodes: one containing a fresh variable ($Ty(e_s)$) and another containing the restrictor. ($Ty(e_s \to cn_s)$). The actions also requires the variable to participate in a scope statement. After the creation of the full determiner node, the subtree can be compiled by applications of *Elimination* and *Thinning*.

(55)

$$?Ty(t)$$

$$Ty(e_s), Fo(\mathbf{S}), ?Sc(\mathbf{S})$$
$$?\exists x\ (Fo(\mathbf{x}) \wedge\, <L>(\mathbf{x})) \qquad ?Ty(e_s \to t)$$

$$Ty(e_s), Fo(\tau, s, WarTime's), \Diamond$$

$$Ty(cn_s) \qquad\qquad Ty(cn_s \to e_s)$$
$$Fo(s, WarTime's) \qquad\qquad Fo(\lambda P.\tau, P)$$

$$Ty(e_s) \qquad Ty(e_s \to cn_s)$$
$$Fo(s) \qquad Fo(\lambda z.z, WarTime'z)$$
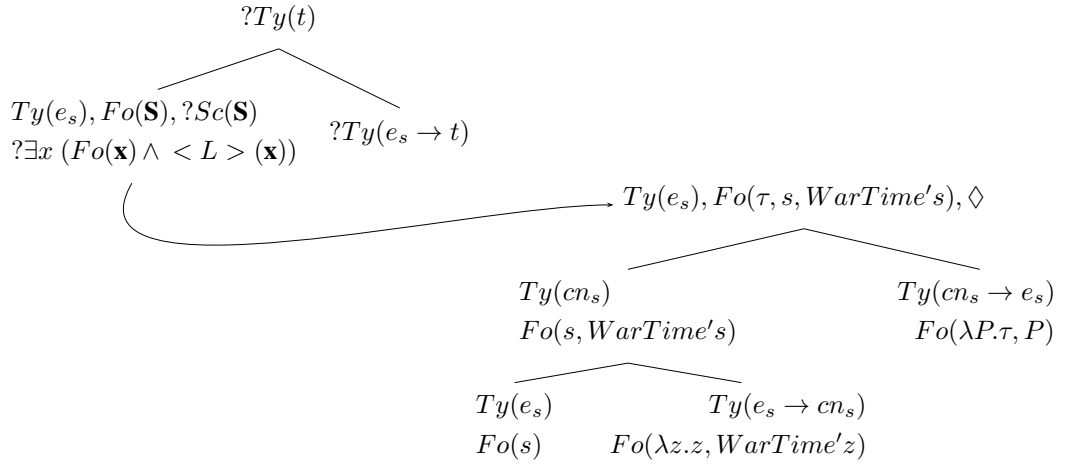
At this point, a member of the family of evaluation rules copies the formula value of the LINKed daughter up to the situation node, the LINK mother. This rule has been independently motivated by the effects induced by conditionals [28].

(56)   *LINK Evaluation Rule: Conditionals* (Gregoromichelaki, 2006, 217)

$$\{\ \{\ Tn(X), Ty(t), ...\}, \{<\uparrow_0> Tn(X), ?Ty(e_s), ?\exists x.Fo(\mathbf{x}), ...\},$$
$$\{< L ><\uparrow 0 > Tn(X), Ty(e_s)Fo(a_{[\mathbf{x}]}), \Diamond\ \}\ \}$$

$$\{\{Tn(X), Ty(t), Scope(\mathbf{U} < \mathbf{x})...\}, \{<\uparrow_0> Tn(X), ?Ty(e_s), ?\exists x.Fo(\mathbf{x}), ?Sc(\mathbf{x}),$$
$$Fo(a_{[\mathbf{x}]}), \Diamond\}, \{< L ><\uparrow_0> Tn(X), Ty(e_s)Fo(a_{[\mathbf{x}]})...\}\ \}$$

In addition to the copying, the rule has the effect of inserting an underdetermined scope statement concerning the variable of interest (*s* as an instantiation of **x**) at the root node. We thus have the following structure:

---

[28]Note that the metavariable **S** in the LINK daughter is the same one as in the LINK mother environment, so providing a formula value for the former means that the same value would be provided for the latter. This, in a way, dispenses with the evaluation rule above, but given that it is motivated for other linguistic phenomena, I have decided to state it here. Applying the rule gives the same results as applying *Substitution* to both occurrences of the metavariable **S**.

(57)

$$?Ty(t), Scope(\mathbf{U} < s)$$

$$Ty(e_s), \cancel{Fo(\mathbf{S})}, ?Sc(s), Ty(e_s), \qquad\qquad ?Ty(e_s \to t)$$
$$Fo(\tau, s, WarTime's), \cancel{?\exists x\, (Fo(\mathbf{x}) \wedge \langle L \rangle(\mathbf{x}))}, \Diamond$$

$$Ty(e_s), Fo(\tau, s, WarTime's)$$

$$Ty(cn_s) \qquad\qquad Ty(cn_s \to e_s)$$
$$Fo(s, WarTime's) \qquad Fo(\lambda P.\tau P)$$

$$Ty(e_s) \qquad\qquad Ty(e_s \to cn_s)$$
$$Fo(s) \qquad Fo(\lambda z.z, WarTime'z)$$
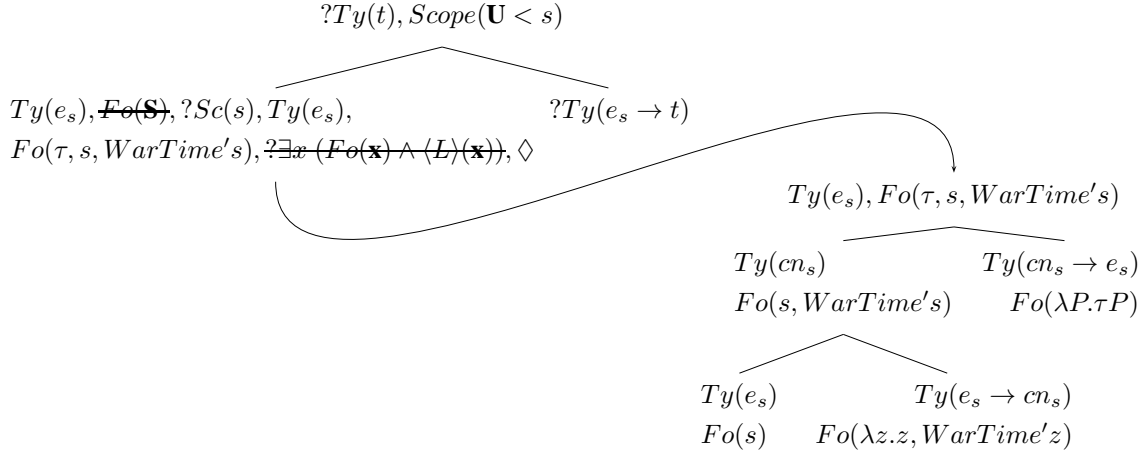
The pointer then moves to the $Ty(e_s \to t)$ node via *Anticipation*. With the pointer at the functor node, the rules of *Introduction* and *Prediction* create new argument and functor nodes. We thus have:

(58)

$$?Ty(t), Sc(\mathbf{U} < s)$$

$$?Sc(s), Ty(e_s), \qquad\qquad ?Ty(e_s \to t)$$
$$Fo(\tau, s, WarTime's)$$

$$?Ty(e), \Diamond \quad ?Ty(e \to e_s \to t)$$

$$Ty(e_s), Fo(\tau, s, WarTime's)$$

$$Ty(cn_s) \qquad\qquad Ty(cn_s \to e_s)$$
$$Fo(s, WarTime's) \qquad Fo(\lambda P.\tau P)$$

$$Ty(e_s) \qquad\qquad Ty(e_s \to cn_s)$$
$$Fo(s) \qquad Fo(\lambda z.z, WarTime'z)$$
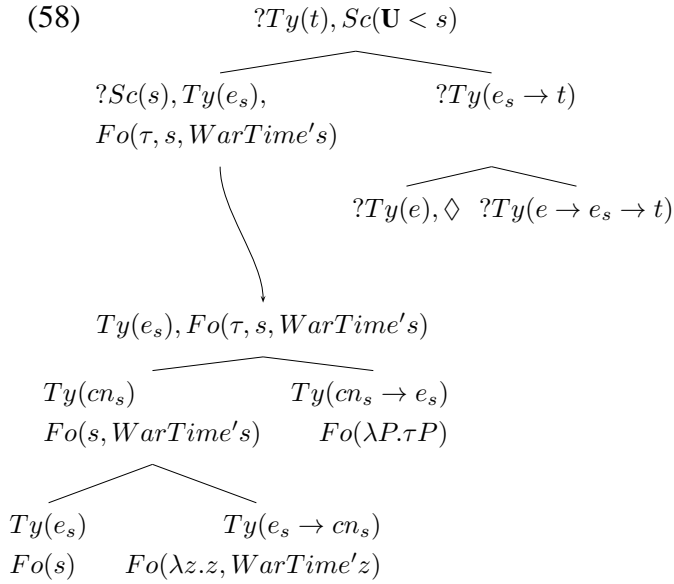
With the pointer at the subject type *e* node, the lexical actions associated with the pronoun 'he' can be parsed. The pronoun introduces a fresh metavariable **Z** and a requirement that it should be replaced by an appropriate formula value. We have the structure in:

(59)

$$?Ty(t), Sc(\mathbf{U} < s)$$

$$?Sc(s), Ty(e_s), \qquad\qquad ?Ty(e_s \to t)$$
$$Fo(\tau, s, WarTime's)$$

$$Ty(e), Fo(\mathbf{Z}), \qquad ?Ty(e \to e_s \to t)$$
$$?\exists x\ Fo(\mathbf{x}), \Diamond$$

$$Ty(e_s), Fo(\tau, s, WarTime's)$$

$$Ty(cn_s) \qquad\qquad Ty(cn_s \to e_s)$$
$$Fo(s, WarTime's) \qquad Fo(\lambda P.\tau P)$$

$$Ty(e_s) \qquad\qquad Ty(e_s \to cn_s)$$
$$Fo(s) \qquad Fo(\lambda z.z, WarTime'z)$$
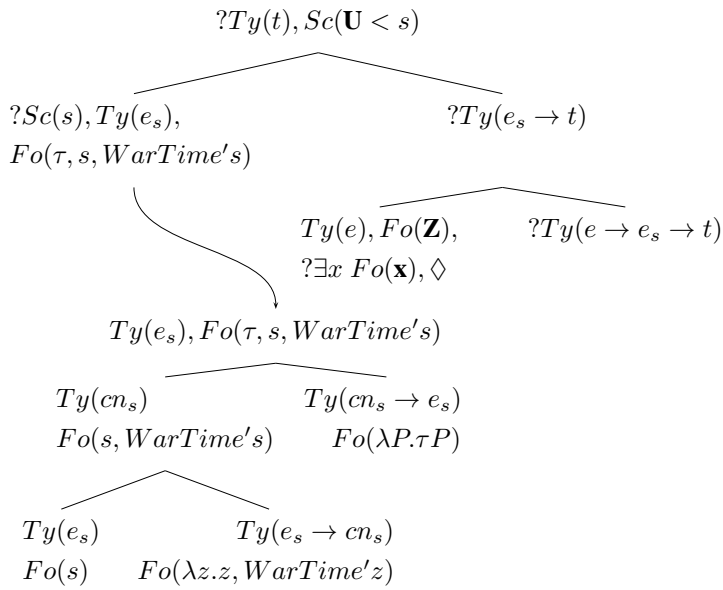
Given that the introduction of the metavariable satisfies the type requirement at that node, we can use *Completion* to move the pointer to the mother node and *Anticipation* to move the pointer down to the predicate node. In this way, we delay the *Substitution* operation until we parse the predicate node, so the hearer has more evidence in order to draw inferences about speaker-meaning. With the pointer at the predicate node, the verb and its complement, namely, 'has tough choices to make' can be parsed. For simplicity reasons, I will represent this verbal complex as a single predicate with formula value $Fo(\lambda t.\lambda v.MTC'tv)$ (abbreviating 'makes tough choices') which takes a subject and a situation variable. Thus, the structure would look like:

(60)
$$?Ty(t), Sc(\mathbf{U} < s)$$

$$?Sc(s), Ty(e_s), \qquad\qquad ?Ty(e_s \to t)$$
$$Fo(\tau, s, WarTime's)$$

$$Ty(e), Fo(\mathbf{Z}), \qquad Ty(e \to e_s \to t)$$
$$?\exists x\ Fo(\mathbf{x}) \qquad Fo(\lambda t.\lambda v.MTC'tv), \Diamond$$

$$Ty(e_s), Fo(\tau, s, WarTime's)$$

$$Ty(cn_s) \qquad\qquad Ty(cn_s \to e_s)$$
$$Fo(s, WarTime's) \qquad Fo(\lambda P.\tau P)$$

$$Ty(e_s) \qquad\qquad Ty(e_s \to cn_s)$$
$$Fo(s) \qquad Fo(\lambda z.z, WarTime'z)$$
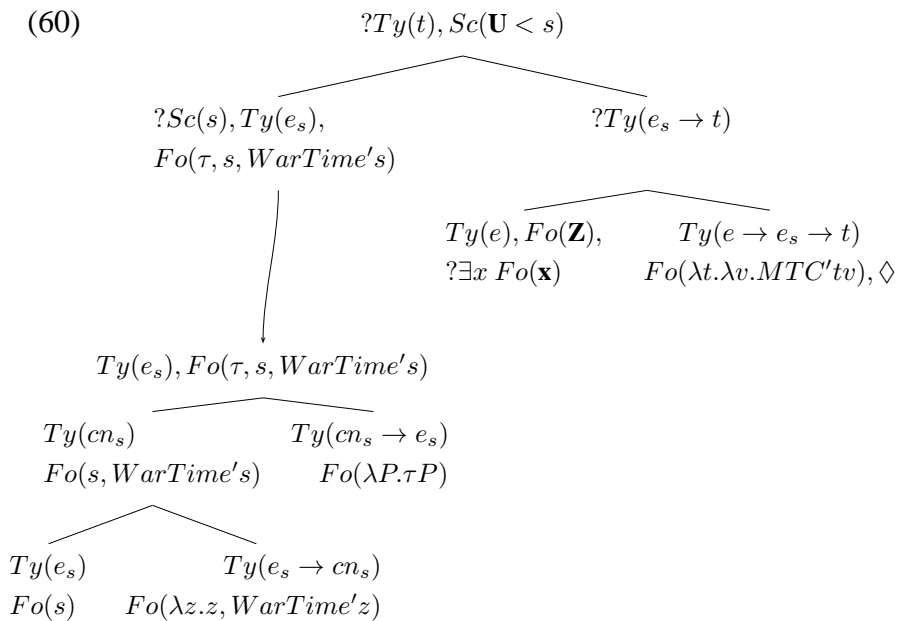
The application of *Completion* moves the pointer to the mother node and the outstanding requirement that a value for the metavariable must be found allows *Anticipation* to move the the pointer to the subject node, so *Substitution* can take place. In the next structure, I will present things slightly differently. First, I will omit the LINKed daughter that represents the adjunct 'every time there is a war' and present instead just the tau term $Fo(\tau, s, WarTime's)$ that derives from it (via the evaluation of that LINK transition). Second, I will include the assumptions that the demonstration towards the White House make salient. As argued previously, these are propositions LINKed to the tree under development, thus providing possible values for the pronoun. We have:

(61) $Ty(t), Fo(WhiteHouse'(d) \wedge WorksIn'(b,d))$

$\langle\downarrow^*\rangle Fo(b), ?Ty(t), Sc(\mathbf{U} < s)$

$?Sc(s), Ty(e_s),$ $?Ty(e_s \rightarrow t)$
$Fo(\tau, s, WarTime's)$

$Ty(e), Fo(\mathbf{Z}),$ $Ty(e \rightarrow e_s \rightarrow t)$
$?\exists x\, Fo(\mathbf{x}), \Diamond$ $Fo(\lambda t.\lambda v.MTC'tv)$
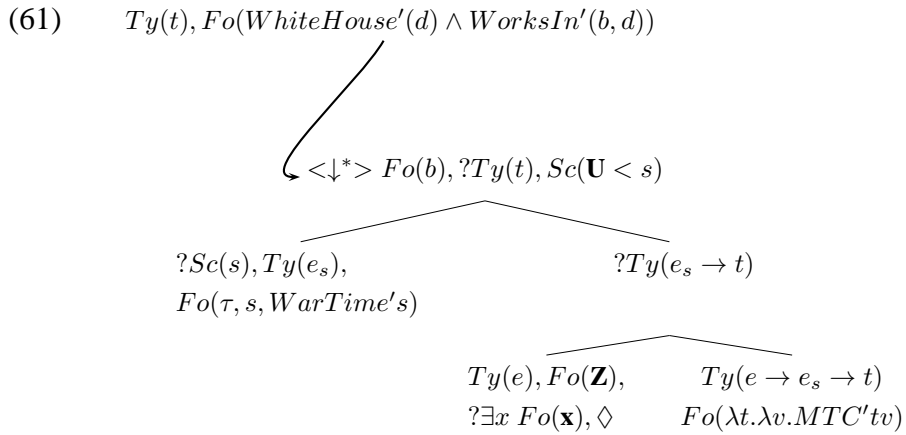
where;

d is a logical constant picking out the demonstrated building, and

$b = \epsilon, x, USPresident'(x)$

The LINKed mother tree (whose structure was abstracted away from) represents two conjoined propositions. The first conjunct describes that the bare individual denoted by the individual constant *d* is a White House. We can consider the assignment of constants to individuals in the world as a formal modelling of the psychological process by which visual indices (FINSTs) attach to the relevant objects (see chapter 3). In this case, the visual index is represented by the constant *d*, which refers to the White House (a given object in a model).

The second conjunct in the LINK mother above says that individual *b* works in *d*. Now, as we have seen in the previous exposition on quantification, the epsilon calculus allows names for arbitrary individuals to have inner structure corresponding to set relations. This means that *b* can regarded as an arbitrary witness that stands for *the US President*. Accordingly, the name *b* can have its content unpacked

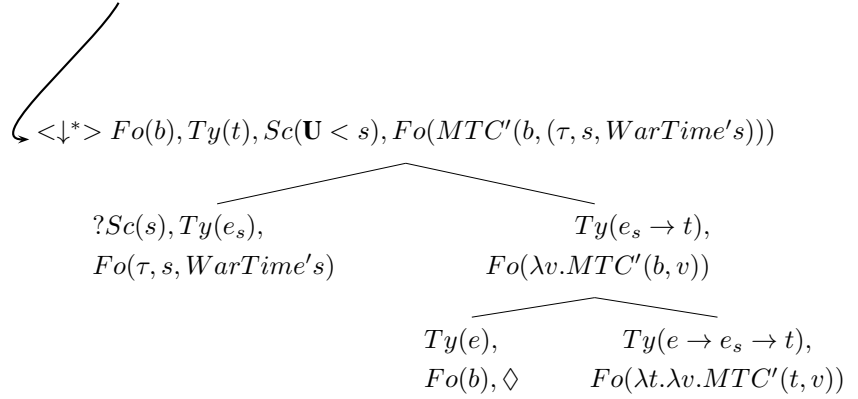as $Fo(\epsilon, x, USPresident'(x))$[29]. The name acts as a shorthand for this more complex, descriptive, representation. Now, as $Fo(b)$ is required to figure somewhere in the LINKed daughter, we have a strong candidate to replace the metavariable encoded **Z** introduced by the pronoun.

(62) $Ty(t), Fo(WhiteHouse'(d) \wedge WorksIn'(b, d))$

$<\downarrow^*> Fo(b), ?Ty(t), Sc(\mathbf{U} < s)$

$?Sc(s), Ty(e_s),$
$Fo(\tau, s, WarTime's)$

$?Ty(e_s \rightarrow t)$

$Ty(e),$ $Ty(e \rightarrow e_s \rightarrow t),$
$Fo(b), \diamond$ $Fo(\lambda t. \lambda v. MTC'(t, v))$
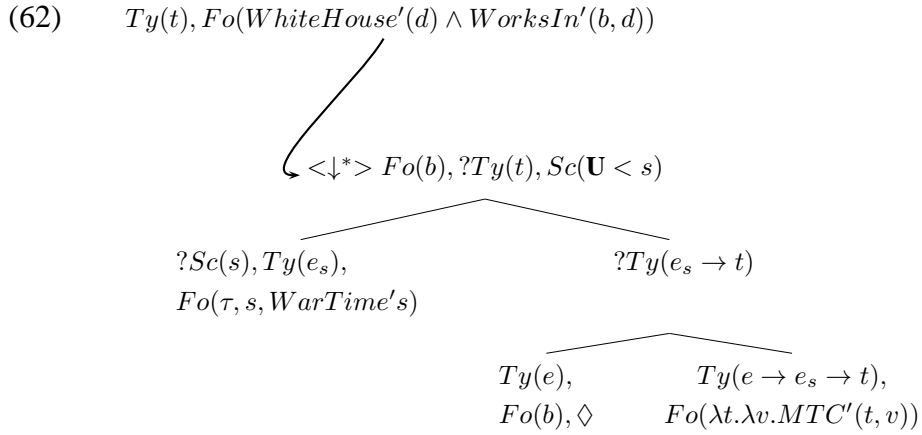
This example illustrates, how the incremental view of processing, sketched in chapter 3, and underwritten by the cognitive principle of relevance, presented in chapter 2, may make available conceptual representations that jointly feed a process of inference, thus producing assumptions that offer pragmatic contributions to the development of the logical form of the utterance (i.e. its explicature). After the *Substitution* operation takes place, *Thinning* and *Completion* move the pointer up to the mother node and then *Elimination* combines the values and types of subject and predicate. Another step of *Thinning* and *Completion* moves the pointer to the root node of the LINKed daughter and the situation term can now be combined with the remaining structure. A final *Thinning* removes the $?Ty(t)$ requirement. The result is:

---

[29] The terms d and b might be argued to be iota terms as the means of expressing their uniqueness, but I have opted for formulating these terms as an individual constant and an epsilon term. Nothing turns on this. There are further alternative ways of cashing out the name contents, in particular their time dependency. I will mention only two of them. First, as we shall see later, there are reasons for assuming that nominals carry a situation (meta)variable, so the arbitrary name would be unpacked as $Fo(\epsilon, x, USPresident'x, \mathbf{S})$. Although this might be a more appropriate way to represent things, as we shall see later, this alternative gives results similar to the position I have chosen to present. Thus, I rather opt for simplicity. Secondly, one could assume that the contents of the arbitrary name capture a complex term, resulting from evaluation rules; something like: $Fo(\epsilon, x, USPresident'x \wedge Works'x, d)$. The problem with this alternative is that quantifier evaluation rules presuppose scope related actions triggered by specific linguistic material (e.g. determiner heads). Since the representations in the LINKed mother result from inference and not decoding, this alternative does not seem to be well grounded. Nevertheless, I take this as an open topic, which the study of the epsilon calculus has much to contribute to.

(63) $Ty(t), Fo(WhiteHouse'(d) \wedge WorksIn'(b,d))$

$<\downarrow^*> Fo(b), Ty(t), Sc(\mathbf{U} < s), Fo(MTC'(b,(\tau, s, WarTime's)))$

$?Sc(s), Ty(e_s),$ $Ty(e_s \rightarrow t),$
$Fo(\tau, s, WarTime's)$ $Fo(\lambda v.MTC'(b,v))$

$Ty(e),$ $Ty(e \rightarrow e_s \rightarrow t),$
$Fo(b), \Diamond$ $Fo(\lambda t.\lambda v.MTC'(t,v))$

The only remaining issue that we must deal with is the scope statement $Sc(\mathbf{U} < s)$. It is important to observe that the name b corresponds to a complex formula value, namely, $\epsilon, x, USPresident'x$, which contains a bound variable. This term, however, emerged as a result of inference rather than by the parse of linguistic material (no *Indef(+)* nor *Indef(-)* are processed). For this reason, the variable x does not figure in the scope statement above.

The scope statement, as it stands, remains unresolved due to the presence of the metavariable $\mathbf{U}$. The role of the metavariable is to allow certain terms to outscope other quantified terms that linearly precede them. Since there are no other scope taking elements in the structure, we can assume that the metavariable coincides with the variable s, thus producing the statement $Sc(s)$ [30]. This statement feeds a *Scope Evaluation* rule that relates the restrictor of the term to its nuclear scope via the appropriate logical connective. The whole formula value is $Fo(MTC'$ $(b,(\tau, s, WarTime's)))$. Thus the restrictor of the tau term (universal binder) is $WarTime'$ and the nuclear scope is $MTC'(b,s)$ (where s marks the position occupied by the quantificational term, thus preserving the binding relation). Since the term is universal, the quantifier evaluation rule (see section 5.2.2 or appendix) replaces the whole term and the variables it binds by an arbitrary name a and relates restrictor (antecedent) and nuclear scope (consequent) via the introduction of a material implication sign. The output of the rule erases the variable from the scope statement at the root node ($Sc(\cancel{s})$) and decorates the root node with the final formula value of:

$Fo(WarTime'(a) \rightarrow MTC'(b,a))$

---

[30] This is a formulation which avoids any need of explicit representation of the time of the utterance as a fixed variable introduced by the Axiom, relative to which the tau term is dependent; see Cann 2011. Since I have retained the simpler form of the Axiom, following Kempson et al 2001 and Cann et al 2005, I adopt this scope-statement simplification instead.

As mentioned before, the arbitrary names 'a' and 'b' themselves impose set relations that reflect the surrounding environment. Thus, a is a shorthand for $\tau, s,$ $[WarTime'(s) \rightarrow MTC'(b, s))]$ and b is a shorthand for $\epsilon, x, USPresident'x$. Although substituting the simple names by their corresponding formula values would make the representation of the utterance content unnecessarily complex, it is important to realise what the set relations denoted by them do. The formula above says that *for any* arbitrary situation a that is a wartime, a is a situation in which an arbitrary individual b is having tough choices to make. Note that b itself denotes a set, namely, the non-empty set of American Presidents; thus imposing that whomever making tough choices at a must also be a US president in that situation. This captures the *Co-variation* between wartimes and Presidents at those times, which the utterance is able to convey. Since there is no other variable in the scope statement. The formula value corresponding to the utterance content cannot be unpacked further. The parse is complete.

We can conclude that the mechanism that generates the relevant formula value successfully accounts for *Co-variation*. Note that it does so in virtue of the ability of episilon terms to reflect their surrounding environment and of quantifier evaluation rules to insert the right connectives between representations denoting sets. This account generalises to cases of donkey anaphora, where the antecedent is provided explicitly in an overt utterance. Let us now move on to other pronominal forms which will provide the means necessary for an explanation of the *Accessibility* property.

5.3.3 First person singular pronouns

Nunberg (1993) argued that basically any kind of pronoun and indexical expression (eg. 'today' 'now') can receive a descriptive interpretation. According to him, these are derived via a linguistically mandated process (chapter 4 section 4.3.1) that incorporates the descriptive truth-conditional content into the proposition expressed. Here are some of the classical examples found in the literature and paraphrases of their respective truth-conditions.

(64)   a.   *Uttered by Clinton*: 'The founders invested me with the sole responsibility for appointing Supreme Court Justices' (Nunberg, 1993, 20).

          b.   *The founders invested the President of the USA with the sole responsibility for appointing Supreme Court Justices.*

          c.   *Uttered by a condemned prisoner:* 'I'm traditionally allowed to order

whatever I like for my last meal' (Nunberg, 1993, 20).

    d.   *A/the Condemned prisoner is traditionally allowed to order whatever he/she likes for his/her last meal.*

However, in chapter 2, the only reliable test for checking this assumption, the scope (embedding) test, gave mixed results as regards to whether the descriptive contents associated with certain uses of the first person singular pronoun fall within the utterance's explicature. This tension is exemplified below. On the one hand, the descriptive interpretation seems to be incorporated into the truth-conditions of the conditional, but, on the other hand, the first person pronoun does not generate a descriptive truth-conditional contribution when embedded under an adverb such as 'usually'. Consider:

(65)   a.   *Uttered by an Artist:* 'If the exhibition had more artists, I'd be in the basement' (Barrios 2011, see chapter 1 section 1.7).

           b.   *If the exhibition had more artists, the paintings of the speaker would be in the basement.*

(66)   a.   'He [*pointing to St. Peter's Basilica or Pope Benedict VXI*] is usually an Italian (intended reading: the pope is usually an Italian)'.

           b.   *Uttered by Benedict XVI*: *'I'm usually an Italian' (intended reading: the pope is usually an Italian).

Given the mixed status of the alleged descriptive interpretations of first person personal pronouns, I follow early insights raised in chapter 1 section 1.7 and chapter 2 section 2.3.3 and propose that they should be explained along lines that slightly differ from their third person cousins. My main claim is that the alleged descriptive conditions above do not fall within the utterance's explicature. I take the data in (66) to support this. Now, let us take a closer look at how the same claim applies to the other examples.

Regarding the examples in (64), I believe that the descriptive readings allegedly associated with the pronouns comes from the meaning of the verb. In (64)a, a singular interpretation containing a mental representation about Bill Clinton is possible, on a certain understanding of the verb 'to invest'. If we take it to be a transitive relation, it is natural to assume that if the founders invest a certain institution (the American State) with the power to delegate to another institution (the American Presidency) the duty to appoint the US Supreme Court Justices, then the founders would have invested Bill Clinton with that power. Note that the interpretation would be about Bill Clinton, but it also makes highly salient that for it to be true, it is nec-

essary for Clinton be an US President. This explains the 'descriptive feel' of the utterance. In (64)c, it has been suggested (Recanati, personal communication), that 'traditionally' is not a full blown adverb of quantification and rather means something like 'From some tradition it follows that...'. Thus, the descriptive feeling of the utterance would follow from the meaning of that expression and the audience's considerations of the type of situations that is properly described by the relevant tradition, that is, it follows from some tradition that, in a certain type of situation (one about prisoners), the speaker (Melvin, say) is allowed whatever he likes as a last meal. Recanati supports this claim by observing that translations of (64)c into another languages (e.g. French, Portuguese) simply do not convey the alleged truth-conditions of the English counterpart (i.e. they sound more like the Pope example in (66)b). The claim is further supported by the fact that, if we replace 'traditionally' by other adverbials, such as 'usually', the alleged descriptive reading is not so accessible.

The example in (65) is one of the best cases for supporting the idea that first person singular pronouns indeed convey descriptive interpretations, for it seems that descriptive interpretation falls in the scope of the conditional and hence should be treated as pertaining to the utterance's explicature. However, as argued in chapter 1 section 1.7 (see the analysis of 'I'm parked out back'), another process might be going on in this case. It could be that the predicate 'to be in the basement' is interpreted as *to have one's paintings exhibited in the basement*. Thus, the descriptive interpretation could come from some form of coercion or metonymic interpretation of the verb (see also the discussion on Recanati's REF feature in 4 section 4.4.2). The truth-conditions of (65) would be paraphrased as *If the exhibition had more artists, I'd be an artist whose work would be in the basement* (or else the predicate could be captured by an *ad hoc* sense modulation, i.e. TO BE IN THE BASEMENT*). The pronoun refers to the particular speaker of the utterance.

I conclude that first person singular pronouns do not seem to convey descriptive interpretations as part of an utterance's explicature (see chapter 2, section 2.3.3 for discussion on how a descriptive proposition can be communicated as an implicature).

5.3.4   Second person singular pronouns

Nunberg (1993) also claimed that 2nd person singular pronouns convey descriptive interpretations. Here are some examples:
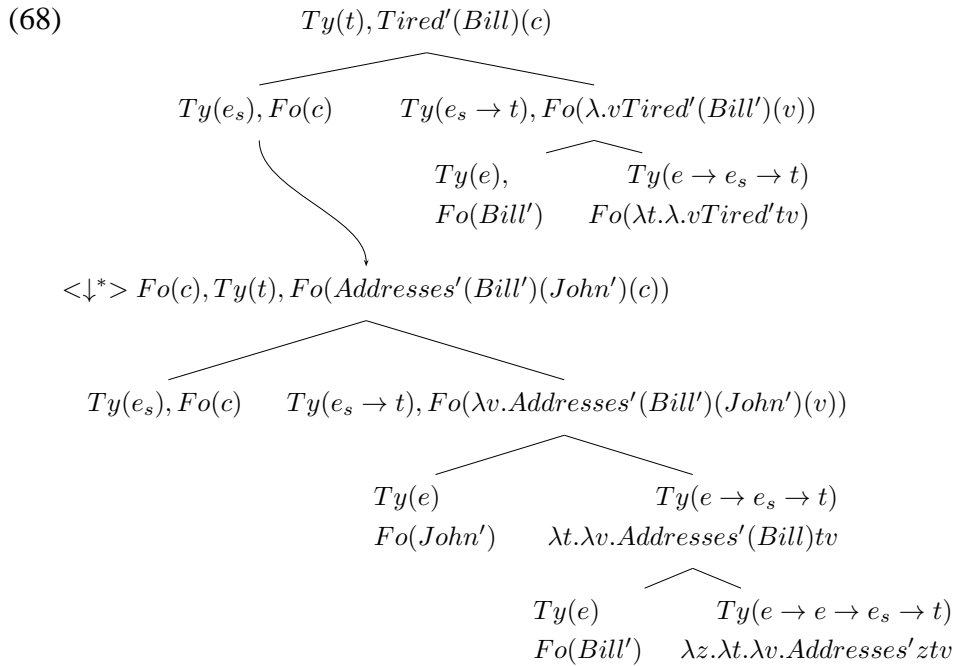
(67)   a.   *Chess teacher giving an introductory lesson to a student who has just*

played 4 N x P: 'According to all the textbooks, you often get in trou-
ble with that move' (Nunberg, 1990).

b. *John, expecting a call from his mother, answers a phone call from his
colleague and says:* 'Oh, I thought you were my mother' (Nunberg,
1993).

According to Nunberg, 'you' in (67)a is interpreted as *the chess player* and in (67)b
as *the person calling*. However, as in the first person singular case, the descriptive
feel of these utterances can be explained by other means. In the first case, it seems
that 'you' is just used impersonally and thus can be paraphrased by 'one'. This
strategy however does not work for the second case. Another explanation must be
sought.

John Perry (2001) famously proposed that utterances containing pronouns com-
municate *reflexive* truth-conditions in addition to other types of truth-conditional
content they determine in a given context. For example, John uttering 'you are hun-
gry' to Bill communicates *that Bill is hungry* (official or incremental content) and
also *that the addressee of u is hungry* (reflexive content). Perry claims that these
truth-conditional contents play different roles in the cognitive life of communica-
tors. The simplest way to illustrate this is to imagine contexts in which the audience
lacks knowledge about who was addressed by the utterance *u*. In this case, only
the reflexive proposition is retrieved. Cases like (67)b can be explained by a similar
approach. When the phone rings, the identity of the caller is unknown, but John
assumes it is his mother. Thus, the proposal made for third person singular pro-
nouns can be extended to capture the second person case as follows. Consider a
simple deictic use, as in the utterance 'You are tired' said by John to Bill. In DS,
deixis could be captured in terms of LINKed trees that share a situation argument.
Consider the structure:

(68)
$$Ty(t), Tired'(Bill)(c)$$

$Ty(e_s), Fo(c)$ $\quad Ty(e_s \to t), Fo(\lambda.vTired'(Bill')(v))$

$Ty(e),$ $\quad Ty(e \to e_s \to t)$
$Fo(Bill')$ $\quad Fo(\lambda t.\lambda.vTired'tv)$

$<\downarrow^*> Fo(c), Ty(t), Fo(Addresses'(Bill')(John')(c))$

$Ty(e_s), Fo(c)$ $\quad Ty(e_s \to t), Fo(\lambda v.Addresses'(Bill')(John')(v))$

$Ty(e)$ $\quad Ty(e \to e_s \to t)$
$Fo(John')$ $\quad \lambda t.\lambda v.Addresses'(Bill)tv$

$Ty(e)$ $\quad Ty(e \to e \to e_s \to t)$
$Fo(Bill')$ $\quad \lambda z.\lambda t.\lambda v.Addresses'ztv$

Above, we have two tree structures LINKed by a situation node. That is, the situation *c* (similar to Kaplanian context) in which Bill is tired is also a situation in which he is being addressed by John. This way of framing context allows us to write lexical entries for pronouns that exploit contextual information in the right way. Tentatively, consider the following lexical entry for 'you':

(69)   Lexical actions for 'you'

IF	$?Ty(e)$
THEN	put $Ty(e), Fo(\mathbf{U}_{Addressee'}), ?\exists \mathbf{x}.Fo(\mathbf{x}), [\downarrow]\bot$
ELSE	abort

Second person pronouns introduce metavariables and presuppose that their referents are being addressed by someone in the context. Treating grammatical person in this way is very similar to the DS modelling of gender, as in example (25), where we took the pronoun 'he' to introduce a metavariable and presuppose that the formula that replace it must denote male entities. In effect, the semantic significance of grammatical person, like gender, can be understood as constraints on proper values for pronominal expressions. This will play a decisive role in the explanation of *Accessibility*.

In the formal treatment of the context above, the LINKed daughter tree provides information about the satisfier of the presupposition. Since Bill is the addressee in the context, the metavariable introduced by 'you' is replaced by $Fo(Bill')$, as

shown in the top tree[31]. In cases where the identity of the addressee is unknown, there would be no formula value to replace the metavariable $U_{Addressee'}$ (this is a first approximation to the presuppositional content of the pronoun, my definite take on it will emerge shortly). Still, this temporary semantic value could be taken as referring to some entity that happens to be addressed in the context. Or else we can assume that the LINK daughter, in this case, corresponds to *Bill* (the speaker) *addresses someone*, where the existential replaces the metavariable introduced by the pronoun. We end up with a Perry style reflexive content. Thus, in the example 'I thought you were my mother', the speaker can be taken to communicate *that at time t, John thought that the person who he would be addressing at time t_1 (the caller) would be his mother* (see Recanati 1993, ch. 16, for a precursor to this idea, also note that much of the hard work of relating different temporal situations, which I am abstracting from here, would be done by the tense of the verb, see Cann 2011). This proposition captures the descriptive feel communicated the utterance. Moreover, the descriptive expression in italics can be generated by assuming the lexical entry needed for indexical cases. There is no need for extra semantic machinery or pragmatic operations of transfer. First person pronouns could be captured by a similar lexical entry: one which introduces $Fo(\mathbf{U}_{Speaker'})$. However, the presuppositions suggested here may be a bit too strong. This will be discussed in the next section and the one about the *Accessibility* property of descriptive uses, where an alternative way to state what pronouns presuppose will be sketched.

5.3.5 Plural Pronouns

Intuitively, plurals can be used to denote groups usually delineated by a certain conceptual representation. For example, a football player who utters 'We won!' conveys that a certain football team won, despite the hearer's lack of knowledge about the individual members of the team. In this case, the audience's understanding is mediated via some sort of conceptual representation (FOOTBALL-TEAM, or $Fo(Football-Team')$, see Sanford et al. (2008) and Filik et al. (2008) for empirical support). Therefore, some representation of sets, captured under certain concepts, is necessary *even in the simplest cases of literal uses of plural pronouns*. Based on this fact, one could argue that the descriptive uses of plural pronouns are an epiphe-

---

[31]Alternatively, one could invoke richer sets of data holding at each node, including information who performs and who is addressed by the speech act. In this spirit, Purver et al. (2010) incorporate elements of type theory with records (TTR) into the DS framework so certain puzzles about speech acts by many participants can be properly solved. This approach has relevance to the discussion here, but since much of the contextual structuring brought in by such records can be mimicked by LINKed structures, I will not discuss it in detail.

nomenon of plural morphology: plurals require representations that do not denote specific individuals, but rather sets of individuals or stuff. This would capture the descriptive feel of the utterances without the necessity for any special semantic machinery or pragmatic transfer operation. However, in this recent example, the use of 'we' refers to a team in the actual world that won the match and the descriptive use of 'we', below, denotes a group of American Justices whose members co-vary with non-actual worlds/situations.

(70)   *Justice O'Connor*: We might have been liberals.

Based on this difference, one could resist the claim that descriptive uses of plural pronouns establish interpretations that result from mere plurality. This claim, however, does not seem to hold. There are reasons to suppose that variables ranging over worlds/situations (which may not refer to the actual one) are not only restricted to verbs or predicates. Consider:

(71)   Every fugitive is (now) in jail (Enç, 1986).

Here, we must interpret 'fugitives' with regards to a past time/situation otherwise we would end up with a clash: an individual cannot be a fugitive and be locked up in jail at the same time. For this reason, some have proposed that every predicate expression (including nouns) comes with a situation variable (e.g. Elbourne 2005; Gregoromichelaki 2006 also mentions that situation arguments should be associated with all predicates, verbs and nouns alike: this is the reason for assuming the the representation about the US president in section 5.3.2 might contain a situation metavariable, see footnote 29). Thus, nominals seem to articulate the representations of non-actual situations or worlds that are relevant for capturing the intended interpretation. This idea together with the fact that plurals denote a group entity can explain the *descriptive* nature of many of the Nunberg cases. Appealing to a special semantic machinery to derive the intended readings is simply not necessary.

Let us look at how these ideas are implemented given the framework sketched above. It must be noted that plurality itself is a very complex topic, so I will limit myself to very basic observations. I assume that plural morphology induces the creation of a group individual of $(Ty(e))$ whose formula value is a metavariable ($\mathbf{X}$) ranging over sets. Group individuals have individuals (or stuff) as constituents and so admit certain mereological relations (see Link 1983 for groundbreaking work). For example, an underdeterminate group $\mathbf{Z}$ may have John and Mary as constituents ($\mathbf{Z} \wedge Mary' < \mathbf{Z} \wedge John' < \mathbf{Z}$).

In the O'Connor case, 'we' induces a metavariable $\mathbf{X}$ which is replaced by $American - Justice'(\mathbf{S})$, where $\mathbf{S}$ is a world/situation metavariable. The modal 'might have been' introduces a representation w about possible worlds accessible from the actual world w_0. Either world representation could replace the value of the metavariable ($\mathbf{S}$), and thus $American - Justice'(w)$ would denote a group of American Justices at some possible world, while $American - Justice'(w_0)$ would denote a group of American Justices at the actual world. Thus, we have the means to represent that the intuitive readings associated with the utterance in (70), which can be used by O'Connor to mean that he members of either group can be liberals.

A question arises at this point. If, like the 2nd person singular, 'we' introduces a group metavariable with a presupposition that the speaker should be part of the group ($\mathbf{X}_{Speaker'}$), then there are two possible interpretations for the pronoun: one that describes a group of Justices in the actual world that includes O'Connor and another that describes a group of Justices in some non-actual world that also includes O'Connor. Neither seem to capture a reading in which O'Connor is not part of the group. Such reading, however, is possible. Therefore, saying that the person feature of 'we' introduces a $_{Speaker'}$ presupposition is too strong. The initial assumptions about some of the presuppositional content of pronouns must be revised.

In order to capture the facts right, some perspectival elements may be needed to be incorporated in the proposal. Instead of the $_{Speaker'}$ presupposition, I propose that first person pronouns presuppose that the speaker (or thinker) has a certain perspective on a given individual in the environment (or aspect of reality, more generally). This is captured by specific mental predicates (P) that represent a certain perspective that a given discourse participant has in relation to a (discourse) entity. Perspectives can be considered as roles in the structure of a dialogue or narrative. The first person pronoun, for example, presupposes that its semantic value is determined from a first person perspective: the perspective of the agent of a certain mental or speech act. The perspectival predicate P could then be prefixed with subscripts that indicate the specific role presupposed. In the first person case, this is captured by P_1 and the lexical actions induced by 'we' look like:

(72)   Lexical actions for 'we'

```
IF      ?Ty(e)
THEN    put Ty(e), Fo(X_{P_1'}), ?∃ X.Fo(X), [↓]⊥
ELSE    abort
```

Presuppositions of perspective can then exploit the structure brought up by LINKed trees in order to be satisfied, thus determining the right values for the pronoun's metavariable. As in the tree-structure displayed in the 2nd person case, the tree representation of 'We might have been liberals' can be LINKed to another tree representation (via the sharing of a situation term: the context of a discourse or narrative, say) that captures the perspective of a given participant in the conversational setting.

Since O'Connor has 1st-person perspectival instance *vis-a-vis* the group of American Justices, the presupposition of the pronoun is satisfied and the formula value $Fo(American - Justice'(\mathbf{S}))$ can replace the metavariable introduced by 'we'. As a result the overall interpretation of O' Connor's utterance of 'We might have been liberals' does not need to include O'Connor in the relevant group in case the world of evaluation is non-actual. Note that this may also be required by uses that are more easily classified as literal discourse. The supporter of a football team who says 'We won!' is not included in the team, but induces the same interpretation that would be induced in case a team player uttered the sentence.

The proposal can be extended to other pronominal forms. Second person plurals might encode a similar presupposition: one that assumes that its value must be determined via a second person perspective (the recipient of the utterance/narrative), that is, they induce $Fo(\mathbf{U}_{P2'})$, and an interpretation identical to the one above could be induced by a speaker who addresses O'Connor with 'You$_{pl.}$ might have been liberals'. Finally, the perspectival element introduced here plays an important role in capturing some of the *Accessibility* patterns displayed by descriptive pronouns, which we will now turn to.

5.3.6 Extended properties: *Accessibility*

The only piece of data left to explain is the *Accessibility* property of descriptive uses. It involves anaphoric relations between pronouns with different person features, among other things. Roughly, explaining *Accessibility* amounts to accounting for the conditions under which certain pronominal interpretations can figure as antecedents for subsequent pronouns. More specifically, almost all the theories reviewed in chapter 4 failed to explain the pattern below, introduced in chapter 1 section 1.8.

(73)   a.   *Said by someone:* If the Democrats had won the last few presidential elections, **the US Supreme Court Justices** might have been liberals. **They** would guard public interest better.

     b.   *Said by Supreme Court Justice O'Connor:* If the Democrats had won the last few presidential elections, **we** might have been liberals. ***They** [*said by same speaker*]/**They** [*said by different speaker*]/??**We** [*said by same speaker*] would guard public interest better.

The contrast is this. O'Connor cannot use the third person pronoun 'They' to refer to the descriptive interpretation of 'we', namely, *the US Supreme Court Justices*, but another speaker can. This is particularly puzzling because third person pronouns usually can pick up salient entities in discourse, which is precisely what much of the data in this thesis suggest, and also because such anaphoric link is allowed under a different speaker. How can the proposal made here account for these facts?

The idea is to extend the perspectival elements introduced in the previous subsection to cover these cases as well. We can assume that the third person feature presupposes that the interpretation of the pronoun must not be established from a 1st or 2nd person perspective. The 3rd person feature introduces a metavariable with a P'_3 presupposition; equivalent to $\neg P'_1 \wedge \neg P'_2$ (not from a first nor second person perspectives). In the previous subsection, we described O'Connor's use of 'we' as determining the formula value $(American - Justice'(w))$ via the satisfaction of the presupposition of a first person perspectival instance. Given that anaphoric uses are modelled via LINKed trees (see Gregoromichelaki 2006, ch. 8, section 2.c.2), the dependency of the subsequent use of 'they' by O'Connor on the $(American - Justice'(w))$ value would be established via the sharing of terms. However, as third person pronouns presuppose that substitutes for the metavariables they induce cannot be established via the first person nor second person perspective, the anaphoric relation is blocked. Note that the anaphoric relation is allowed by another speaker's use of the same expression, because this (other) speaker did not determine the formula value $(American - Justice'(w))$ through the first person (i.e. her) perspective.

At this point, the following objection could be made. To some native speakers, a second use of 'we', by O'Connor, that is anaphoric on the descriptive interpretation of the first use 'we', namely, $(American - Justice'(w))$, is as degraded as the anaphoric use of 'they' discussed above (Robyn Carston, p.c.), but surely the anaphora in this case should be allowed, since the pronoun presupposes a first person perspective and the interpretation has been established through this perspectival instance. The point is well taken. However, general pragmatic mechanisms also play a part in the explanation of why certain representations are more accessible than others. Such an anaphoric use of 'we' contrasts with an utterance, by

O'Connor, of 'if the Democrats had won the last few presidential elections, **we** might have been liberals and *guarded public interest better*' which is easier to process, given that the coordinated VPs are contained in the consequent of the conditional, facilitating an interpretation where the people guarding public interest better are a group of Justices (possibly not including O'Connor) in a non-actual situation. Given the availability of a sentence that demands less processing effort to achieve the same effects, an anaphoric use of 'we' which depends on a descriptive interpretation of an antecedent use of the same expression is a bad word choice. The proposal sketched here successfully captures the *Accessibility* property of descriptive uses, meeting a *desideratum* which was not touched on by previous accounts.

There are other desirable features of the proposal made here. For starters, it employs two independently motivated mechanisms of utterance interpretation. The first of them is the structure created by LINKed tree representations, independently motivated to capture anaphora and other linguistic phenomena, such as the processing of relatives and cross-over effects. Thus, we are simply re-using an already necessary theoretical vocabulary.

The other general mechanism of interpretation consists in the perspectival predicates that capture specific perspectival stances. These mechanisms have been required by other linguistic and mental phenomena. For example, Higginbotham (2002) claims that linguistic competence with (indexical) pronouns enables language users to track different perspectives. Tracking of perspectives here would explain, among other things, the informativeness of uses of different pronominal forms that refer to the same individual. For example, Perry (1979) considers a speaker who sees a person whose trousers are on fire and describes the situation by an utterance of 'His trousers are on fire', but unbeknownst to the speaker, he is looking at himself in the mirror. When the speaker is in a position to realise the thought expressed by 'My pants are on fire', an important perspectival shift has occurred: even though pronominal reference remains the same, the speaker acquired a relevant piece of information. In the framework above, this difference would be captured in terms of different tree structures that are LINKed to the tree-structured representations of the utterances (which presuppose different perspectival stances according to the pronouns used). Similarly, some linguists (Hagège, 1974; Clements, 1975) have been intrigued by the ability of some pronominal expressions (logophors) to shift their value according to the perspective of a narrative. This has recently re-surfaced in the literature, as counter examples to Kaplan's (1989a,b) theory of indexicals (against the rigidity facts that Kripke 1972 allegedly raised, see chapter 1 section 1.3). For example, in some languages, like Amharic, a report equivalent to the

English sentence 'John believes I am a hero' could induce the interpretation *John believes John is a hero*, that is, 'I' in Amharic would determine its reference with regards to a shifted situation, one in which the agent is the subject of the report (i.e. the believer) and not the speaker of the utterance (Schlenker 2003; Anand and Nevins 2004; for the relation between the grammatical person and the representation of perspectives, see Sauerland 2008). Hopefully, the proposal sketched here could be extended to cover such shifts as well, but this transcends the aims of this thesis.

Finally, the proposal has a second explanatory advantage. It strikes a fine balance between semantics and pragmatics. On the one hand, it leaves pragmatics with the task of determining relevant contextual information. For the data discussed here, there are two types of contextually available information that plays a crucial role: information concerning the terms that provide the descriptive terms that replace the metavariables introduced by pronouns and information about the perspectival stance of the conversational participants. On the other, it leaves grammar with the task to recruit such information in order to deliver the attested co-varying readings (via quantifier evaluation rules) and also to describe which representations are accessible given a word use and a particular perspectival stance (via specific contextual assumptions represented as LINKed structures).

5.4 Conclusion

In this chapter, I have discussed how the combination of relevance theoretic pragmatics (chapter 2), the incremental notion of processing (chapter 3) and the language DU (the DS grammar framework) explains all properties of descriptive uses of pronouns. The proposal made here does not suffer from the shortcomings faced by previous proposals (chapter 4) and explains the data using general and independently motivated mechanisms of interpretation. Let me now present some concluding remarks.

5.5 Appendix: Computational rules

Introduction

$$\frac{\{...\{...?Ty(Y)..,\Diamond\}...\}}{\{...\{...?Ty(Y),? <\downarrow_0> Ty(X),? <\downarrow_1> Ty(X \to Y)...,\Diamond\},...\}}$$

Prediction

$$\frac{\{...\{Tn(n),...? <\downarrow_0> X, ? <\downarrow_1> Y, \Diamond\}...\}}{\{\{Tn(n),...? <\downarrow_0> Ty(X), ? <\downarrow_1> Y\}, \{<\uparrow_0> Tn(n), ?X, \Diamond\}\{<\uparrow_1> Tn(n)?Y\}\}}$$

Thinning

$$\frac{\{...X,...?X,...\Diamond\}}{\{...X,...,\Diamond\}}$$

Completion

$$\frac{\{...\{Tn(n)...\}, \{<\uparrow_i> Tn(n),...Ty(X),...\Diamond\}\}}{\{\{Tn(n),... <\downarrow_i> Ty(X),...,\Diamond\}, \{<\uparrow_i> Tn(n),....Ty(X),...\}\}}$$
where $i \in \{0, 1, *\}$

Anticipation

$$\frac{\{\{Tn(n)...,\Diamond\}, \{<\uparrow> Tn(n),...?X\}\}}{\{\{Tn(n)...\}, \{<\uparrow> Tn(n),...?X....\Diamond\}\}}$$

Elimination

$$\frac{\{<\downarrow_0> (Fo(a), Ty(X)), <\downarrow_1> (Fo(b), Ty(X \rightarrow Y))...,\Diamond\}}{\{...\{Ty(Y), Fo(b(a)), <\downarrow_0> (Fo(a), Ty(X)), <\downarrow_1> (Fo(Y), Ty(X \rightarrow Y))...,\Diamond\}\}}$$

Star Adjunction

$$\frac{\{\{Tn(a)...,?Ty(t),\Diamond\}\}}{\{\{Tn(a)...,?Ty(t),\}\{<\uparrow_*> Tn(a)...,\exists x.Tn(\mathbf{x}), ?Ty(e),\Diamond\}\}}$$

Merge

$$\frac{\{ND..., ND'...,\}}{\{ND \sqcup ND'\}}$$
where $\Diamond \in$ ND' and $ND \cup ND'$ is compatible.

Local Star Adjunction

$$\frac{\{\{Tn(n)\},..\{Tn(m), <\uparrow_*> Tn(n), Ty(X), \Diamond\},...\}}{\{\{Tn(n)\},..\{Tn(m), <\uparrow_*> Tn(n), Ty(X)\}, \{<\uparrow_*> Tn(m), ?Ty(X), ?\exists x .Tn(x), \Diamond\}...\}}$$

Link Adjunction

$$\frac{\{\{Tn(X)..., Ty(e), Fo(a)...,\Diamond\}...\}}{\{\{Tn(X)..., Ty(e), Fo(a)\}\{< L^{-1} > Tn(X), ?Ty(t), ? <\downarrow^*> Fo(a), \Diamond\}..\}}$$

Q-Evaluation Rule

$$\frac{\{...\{Ty(t),..,World(w[x_1]),Scope(x_1 < ...x_n),Fo(\phi[vx_n\psi_n/x_n]),...\}\}}{\{\{Ty(t),..,World(w[x_1]),Scope(x_1 < ...x_n - 1),...,Fo(f_{v_nx_n\psi[vx_n\psi_n/x_n]}(\phi)),...\}\}}$$

where for x occurring free in ϕ and $w[x_1]$ = a world variable x_1 or $w[x_1] = v_1x_1\psi$, and the values $f_{v_nx_n\psi[vx_n\psi_n/x_n]}(\phi)$ for $v \in \epsilon, \tau, Q$ and $f_w[x1](\psi)$ are defined by:

a. $f_{\tau x\psi}(\phi) = \psi[a/x] \to \phi[a/x]$, where a $= \tau x(\psi \to \phi)$

b. $f_{\epsilon x\psi}(\phi) = \psi[b/x] \wedge \phi[b/x]$, where b $= \epsilon x(\psi \wedge \phi)$

c. $f_{Qx\psi}(\phi) = (\psi[c/x])(\phi[c/x])$, where c $= v_Qx((\psi)(\phi))$

d. $f_{w[x_1]}(\phi) = w[x_1] : \phi$

Situation Metavariable Insertion Rule

$$\frac{\{...Tn(n),?Ty(t),\},\{<\uparrow_0> Tn(n),?Ty(e_s),\Diamond\},\{<\uparrow_1> Tn(n),?Ty(e_s \to t)\}\}...}{\begin{array}{c}\{...Tn(n),?Ty(t),\Diamond\},\{<\uparrow_0> Tn(n),Ty(e_s),Fo(\mathbf{S}),?Sc(\mathbf{S}),?\exists x\, Fo(x)\}\\ \{<\uparrow_1> Tn(n),?Ty(e_s \to t)\}\}...\end{array}}$$

Situation Adjunction Rule

$$\frac{\begin{array}{c}\{...Tn(n),?Ty(t)\},\{<\uparrow_0> Tn(n),Ty(e_s),Fo(\mathbf{S}),?Sc(\mathbf{S}),?\exists x\, Fo(x),\Diamond\}\\ \{<\uparrow_1> Tn(n),?Ty(e_s \to t)\}\}...\end{array}}{\begin{array}{c}...Tn(n),?Ty(t)\},\{<\uparrow_0> Tn(n),Ty(e_s),Fo(\mathbf{S}),?Sc(\mathbf{S}),?\exists\, xFo((\mathbf{x}) \wedge\, <L> Fo(\mathbf{x}))\}\\ \{<\uparrow_1> Tn(n),?Ty(e_s \to t)\}\},\{<L^{-1}><\uparrow_0> Tn(n),?Ty(e_s),?\exists\, xFo(\mathbf{x}),\Diamond\}...\end{array}}$$

LINK Evaluation Rule: Conditionals

$$\frac{\begin{array}{c}\{\,\{\,Tn(X),Ty(t),...\},\{<\uparrow_0> Tn(X),?Ty(e_s),?\exists x.Fo(\mathbf{x}),...\},\\ \{<L><\uparrow 0> Tn(X),Ty(e_s)Fo(a_{[\mathbf{x}]}),\Diamond\,\}\,\}\end{array}}{\begin{array}{c}\{\{Tn(X),Ty(t),Scope(\mathbf{U} < \mathbf{x})...\},\{<\uparrow_0> Tn(X),?Ty(e_s),?\exists x.Fo(\mathbf{x}),?Sc(\mathbf{x}),\\ Fo(a_{[\mathbf{x}]}),\Diamond\},\{<L><\uparrow_0> Tn(X),Ty(e_s)Fo(a_{[\mathbf{x}]})...\}\,\}\end{array}}$$

Chapter 6

Conclusion

It is now time to state the concluding remarks of this thesis. First, I will summarise the contributions it has made to the debate on identification-based descriptive uses of pronouns. Then, I will anticipate a possible objection to the proposal sketched in the previous chapter.

My first task, set out in chapter 1, was to extract the key properties of identification-based descriptive uses of pronouns. On the one hand, this was done by looking at previous literature. Since Nunberg (1993), it seems that such pronominal uses depend on the identification of an individual in the context (i.e. a property known as *Identification Dependency*) to communicate descriptive information (i.e. *Descriptiveness*) in a way such that the conceptual representation involved in identification provide the building blocks (in a rather indirect way) of the descriptive interpretation (i.e. *Connection*). On the other hand, this task was also carried out by comparing such uses to other pronominal and non-literal interpretations. In this comparison, I have argued that descriptive uses seem to bear closer similarities to cross-sentential (donkey) anaphora and metonymy.

Like cross-sentential anaphora, descriptive uses seem to co-vary with other expressions in discourse (i.e. *Co-variation*), even though such 'antecedents' are not given linguistically. Like metonymy, such pronouns seem to be involved in a situation where the concept most salient in discourse (e.g. FOOTPRINT) provides a gateway to the concept relevant for interpretation (FOOTPRINT-MAKER). However, pronouns encode person features that seem to presuppose certain discourse perspectives, whereas metonymical uses, in virtue of exploiting words (e.g. 'dog') associated with concepts (e.g. DOG), are not bound by such perspectival takes. This was used to show that descriptive uses of pronouns of a given grammatical person cannot provide antecedents for subsequent pronominal uses of different person features (i.e. *Accessibility*). The last two were identified as the extended properties of the

data, not only because they depend on the core properties, but also because they have not been discussed extensively in the literature (with the exception of a few authors). Moreover, the variety of different types of pronominal interpretations examined in this chapter raised difficulties for a unitary account of such expressions.

The task of chapter 2 was to introduce principles that govern the audience's use of contextual information to figure out speaker-meaning (i.e. Pragmatics) and examine how such principles would explain the core features of the data. Our starting point was Gricean pragmatics. In a nutshell, a reconstruction of a Gricean explanation of descriptive pronouns treats them as non-literal word uses, and thus would classify them as instances of particularised conversational implicatures. In order to derive this level of content, the audience must be able to recover what-is-said by the speaker and engage rational principles governing communication in order to derive further levels of speaker meaning. Many difficulties for a Gricean-inspired account were raised. The most problematic ones were related to (i) the structural dependency between implicatures and what-is-said, given the difficulty of describing how this latter level of content looks like in the case of descriptive pronouns, and (ii) the fact that it excluded pragmatically available constituents to figure as part of the proposition expressed by the utterance, hence making some of the properties of the data (e.g. *Co-variation*) very difficult to account for.

Relevance Theory was presented as an alternative pragmatic approach. It does not inherit the shortcomings of the Gricean view (i.e. i-ii, above), allowing for pragmatically derived constituents to figure in the proposition expressed by the utterance. However, many of the options it allows, such as an implicature-based account or the *ad hoc concepts* construction account seem to be ill suited for explaining the data in a unified manner. It seems that an account that treats some descriptive cases as implicatures and some other cases as explicatures fits the data best. Thus, the thesis has contributed positively to the debate by breaking the taboo that a linguistic phenomenon must be treated *either* as contributing to implicit or to explicit content. Within explicature-based possibilities, more specifically, it seems that descriptive uses result from a mere substitution of variable-like component encoded by the pronoun by the appropriate descriptive representation. Within implicature based accounts, the variable-like entity would be replaced by a singular interpretation and implicate the descriptive reading. The problem is how to describe the linguistic meaning of these expressions in a way such that the substitution operation is both licensed and explanatory; a task chapter 5 was dedicated to.

In chapter 3, we looked at arguments in the literature against the idea that extra-linguistic (perceptual) stimuli can provide 'antecedents' for pronouns in sub-

sequent discourse. These arguments motivate constraints, such as, the *Anaphoric Island Constraint* (Postal 1969, also known as the formal link condition) or the *Individuative-Representational Constraint*. More specifically, we examined how mechanisms of visual processing make conceptual representations available for a variety of cognitive tasks, including utterance comprehension, from both conceptual and empirical standpoints. Against such constraints, I have argued that concepts that are indirectly available in a given visual experience can be integrated with the encoded meaning of pronouns in virtue of the nature in which they emerge from earlier (possibly non-conceptual) levels of information. Detailing such mechanisms and how they can be seen in the light of general principles of information processing (such as the Cognitive Principle of Relevance) opens up the possibility of seeing dependencies on visually available representations in a way similar to dependencies on linguistic antecedents: another important contribution to the debate. If the availability of such antecedents is underpinned by an incremental view of processing and by activation coming from the hypothesis about speaker-meaning, it seems that developing the logical form of an utterance is an activity that evolves incrementally as well. Ideally, one would like to see that reflected in a grammatical theory.

Chapter 4 had the task of investigating accounts of identification-based descriptive uses proposed in the literature. These fall within two broad families: heavy-handed semantic theories or heavy-handed pragmatic theories (borrowing a term from Neale 2007). The first tended to explain the data by evoking a set of linguistic constraints on interpretation. However, they fail to explain how, in a given context, a given interpretation is selected as the descriptive value of the pronoun. As they did not invoke any component of pragmatics in their explanation, these proposals were partial at best. The latter family had difficulties regarding the mechanisms by which identification-based descriptive interpretations co-vary with other expressions in discourse (e.g. *Co-variation*). Moreover, both families have difficulty explaining the *Accessibility* property of identification-based descriptive uses of pronouns, which requires some combination of pragmatic elements as well as grammatical ones. The various shortcomings motivate an alternative proposals.

Building on the conclusions of previous chapters, in chapter 5, I have presented an alternative proposal for identification-based descriptive uses of pronouns. Differently from the string-based view of grammatical processes sketched in chapter 3 and the heavy-handed semantic and pragmatic accounts discussed in chapter 4, the data seem to require (i) a Pragmatic framework different from Grice and (ii) grammar that allows pragmatically available information to interact with encoded information on the fly and describes the linguistic constraints on pragmatic opera-

tions.

The demands on the pragmatic side are met by Relevance Theory (Sperber and Wilson 1986/95). The demands on the grammar side are met by Dynamic Syntax (Kempson et al 2001, Cann et al 2005). In a nutshell, DS offer tools that model the incremental growth of tree-representations on the basis of linguistic stimuli and contextual information. It represents context, like RT, as a set of (LINKed) propositions. Identification-based descriptive uses of pronouns can be explained along the following lines. The demonstration or saliency of an object in the context provides relevant mental representations for purposes of interpretation; hence explaining *Identification Dependency*. These representations can provide implicitly available 'antecedents' (captured as LINKed trees) that can replace the metavariable introduced by the pronoun, in some cases (e.g. third person singular, plurals). Else, they provide contextual assumptions that implicate some descriptive proposition, in other cases (e.g. some first person singular uses). The availability of descriptive representations that replace the pronominal metavariable accounts for *Descriptiveness* and the fact that this process is one of inference accounts for *Connection*. We have also seen how the quantifier evaluation rules for the epsilon calculus unpack the conceptual content of the utterance in a way such that the *Co-variation* between identification-based descriptive interpretations and other representations in discourse is properly accounted for. Finally, I have argued that grammatical person encodes presuppositions about different discourse perspectives. Possible clashes between the presupposed perspective on discourse may block identification-based descriptive interpretations to provide the value for other pronominal expressions. This explains *Accessibility*. In conclusion. The combination of DS and RT is able to explain all the properties of the data without the shortcomings levelled against other accounts.

I believe that the proposal made here carries interesting consequences for a number of topics in linguistics and philosophy of language.

First, with regards to the challenge of explaining the various pronominal interpretations reviewed in chapter 1, the proposal in chapter 5 assumes that the different interpretations (e.g. deictic, bound, or donkey anaphoric) result from substituting the pronoun's metavariable by representations that become available in slightly different ways (though they would all be underpinned by the same pragmatic principles and incremental take on processing). For example, deictic uses would result from substitutions that target a mental representation that is immediately and directly available in the perceptual environment, whereas standard cases of pronominal binding and donkey anaphora result from substitutions that target conceptual

representations made available in previous discourse. The complementary distribution of bound pronouns and self anaphors ('him' vs 'himself') could be explained along the same lines as the explanation of the *Accessibility* property: 'self' would require a very local antecedent (described under LOFT terms), while a bare pronoun could be bound by non-local ones. The proposal thus provides a unified explanation rooted in a underdeterministic view of encoded meaning that is fully compositional.

Secondly and in connection with the previous point, I hope to have suggested that much of the worries raised by adepts of *Semantic Minimalism* against underdeterministic views on the proposition expressed by an utterance, namely, that it is not fully composition and overgenerates. This is not the case. As I have argued, according to the proposal, the formula value in the root node derives from the formula values of daughter nodes. Moreover, it includes a description of grammar that blocks certain representations from being proper values for substitution.

Thirdly, I believe that the data itself, reviewed in chapter 1, and the discussion in chapter 3, offer a counter argument to the standard orthodoxy that sets linguistic processes apart from other cognitive processes. The metaphor for a specialised linguistic module that 'ships off' content to other modules needs revision. According to my proposal, cognitive processing in general is an effect-driven effort-saving endeavour: linguistic and perceptual alike. Thus, the top-down expectations about what an ostensive stimuli means, such as pointing to a footprint on the ground, can make certain representations available for processing further linguistic stimuli. This, however, does not deny the existence between specialised systems.

Finally, I believe that data raised by descriptive pronouns suggests that taking a given linguistic phenomenon and asking whether it should be captured at the explicit or implicit level of meaning is not a very interesting question. As we have seen, some descriptive uses contribute to the explicit level, whilst some others contribute to the implicit level. What seems to be required by the data is: (i) to specify the principles by which pragmatic contributions are generated, (ii) to describe the type of information encoded by specific lexical items, and (iii) to describe the interplay of information between (i) and (ii). On the one hand, many philosophers of language and formal semanticists tried to abstain themselves from a description of pragmatic principles because they took their inclusion into theory making to be a threat of more important principles, such as compositionality and systematicity, or because they took formal treatments of indexicality to be the model for context dependency in general. On the other hand, pragmatists always took the workings of pragmatic principles to be quite independent from a description of meaning (leaving this job to syntacticians and semanticists) and a description of other representational abilities.

However, encoded information might be able to constrain pragmatically available information in the comprehension process (as highlighted by the *Accessibility* property). Similarly, natural constraints in vision (echoing Marr and Pylyshyn) may shape the type of representation that emerges in cognition. These play a role in theory building and anyone interested in the study of language should take a description of representational abilities (linguistic, perceptual, etc..) into consideration.

Bibliography

Almog, J., J. Perry, and H. Wettstein (Eds.) (1989). Themes from Kaplan. New York: Oxford University Press.

Anand, P. and A. Nevins (2004). Shifty operators in changing contexts. In Proceedings of SALT, Volume 14, pp. 20–37.

Anderson, S. R. (1971). A note on one's private's. Defamatory essays presented to Jim McCawley, pp. 45–51. Edmonton: Linguistic Research. pseudonymously as 'P. R. N. Tic Douloureaux'.

Austin, J. L. (1962). How to Do Things with Words. Oxford: Clarendon.

Bach, K. (1987). Thought and Reference. Oxford University Press.

Bach, K. (1994). Semantic Slack: What is Said and More. Foundations of speech act theory: Philosophical and linguistic perspectives, 267–291.

Bach, K. (1999). The myth of Conventional Implicature. Linguistics and Philosophy 22(4), 327–366.

Bach, K. (2001). You don't say? Synthese 128(1), 15–44.

Bahrami, B. (2003). Object property encoding and change blindness in multiple object tracking. Visual cognition 10(8), 949–963.

Baillargeon, R. (1987). Object permanence in 3 1/2-and 4 1/2-month-old infants. Developmental Psychology 23(5), 655–664.

Baillargeon, R. (2008). Innate ideas revisited: For a principle of persistence in infants' physical reasoning. Perspectives on Psychological Science 3(1), 2.

Baillargeon, R., E. Spelke, S. Wasserman, et al. (1985). Object permanence in five-month-old infants. Cognition 20(3), 191–208.

Barrios, E. (2011). Meaning shift and the purity of 'I'. Talk delivered at the Meaning, Context and Implicit Content conference, Cerisy, La Salle.

Barsalou, L. (1983). Ad hoc Categories. Memory and cognition 11(3), 211–227.

Barsalou, L. (1987). The Instability of Graded Structure: Implications for the nature of concepts. Concepts and Conceptual Development: Ecological and intellectual factors in categorization, 101–140.

Barsalou, L. (1993). Flexibility, Structure, and Linguistic Vagary in Concepts: Manifestations of a Compositional System of Perceptual Symbols. Theories of memory, 29–101.

Barwise, J. and R. Cooper (1981). Generalized quantifiers and natural language. Linguistics and Philosophy 4, 159–219.

Bemis, D. K. and L. Pylkkänen (2011). Simple composition: A meg investigation into the comprehension of minimal linguistic phrases. Journal of Neuroscience 31(8)(31), 2801–2814.

Bezuidenhout, A. (1997). Pragmatics, semantic undetermination and the referential/attributive distinction. Mind 106(423)(423), 375–409.

Blackburn, P. and W. Meyer-Viol (1994). Linguistics, Logic and Finite Trees. Logic Journal of IGPL 2(1), 3–29.

Blaser, E., Z. Pylyshyn, and A. Holcombe (2000). Tracking an object through feature space. Nature 408(6809), 196–199.

Bonatti, L., E. Frot, R. Zangl, and J. Mehler (2002). The Human First Hypothesis: Identification of Conspecifics and Individuation of Objects in the Young Infant. Cognitive Psychology 44(4), 388–426.

Borg, E. (2002). Pointing at Jack, Talking about Jill: Understanding Deferred Uses of Demonstratives and Pronouns. Mind & Language 17(5), 489–512.

Borg, E. (2004). Minimal Semantics. Oxford: OUP.

Breheny, R. (1999). Context Dependence and Procedural Meaning. Ph. D. thesis, University College London, London.

Brogan, T., V. Terry, and A. Preminger (1993). The New Princeton Encyclopedia of Poetry and Poetics. Princeton University Press.

Burge, T. (2010). Origins of Objectivity. Oxford University Press.

Büring, D. (2001). Discourse-trees and dynamic updates. Unpublished Manuscript.

Büring, D. (2005). Binding Theory. Cambridge University Press.

Campbell, J. (2002). Reference and consciousness. New York: Oxford University Press.

Cann, R. (2011). Towards an account of the english auxiliary system. In R. Kempson, E. Gregoromichelaki, and C. Howes (Eds.), The Dynamics of Lexical Interfaces, pp. 279–318. CSLI.

Cann, R., R. Kempson, and L. Marten (2005). The Dynamics of Language. Elsevier.

Cappelen, H. and J. Hawthorne (2009). Relativism and Monadic Truth. Oxford: OUP.

Cappelen, H. and E. Lepore (2005). Insensitive Semantics. Oxford: Blackwell.

Carey, S. (2009). The Origins of Concepts. New York: Oxford University Press.

Carey, S. and F. Xu (2001). Infants' knowledge of objects: Beyond object files and object tracking. Cognition 80(1-2), 179–213.

Carlson, G. (1977). Reference to Kinds in English. Ph. D. thesis, University of Massachusetts, Amherst.

Carlson, G. and F. Pelletier (1995). The Generic Book. University of Chicago Press.

Carnap, R. (1947). Meaning and Necessity. University of Chicago Press.

Carston, R. (1988). Implicature, explicature and truth-theoretic semantics. In R. Kempson (Ed.), Mental Representations: The Interface between Language and Reality, pp. 155–182. Cambridge: Cambridge University Press.

Carston, R. (1995). Quantity Maxims and Generalised Implicature. Lingua 96(4), 213–244.

Carston, R. (1997). Enrichment and Loosening: Complementary Processes in Deriving the Proposition Expressed. Linguistische Berichte 8, 103–127.

Carston, R. (2002). Thoughts and Utterances: The Pragmatics of Explicit Communication. Blackwell.

Carston, R. (2004). Explicature and semantics. In B. S. G. Steven Davis (Ed.), Semantics: A reader, pp. 817–845. Oxford University Press.

Carston, R. (2007). How Many Pragmatic Systems are There? In Saying, meaning, referring: Essays on the philosophy of Francois Recanati, pp. 18–48. Palgrave Macmillan.

Carston, R. (2010). Metaphor: ad hoc concepts, literal meaning and mental images. Proceedings of the Aristotelian Society 110(3)(110), 295–321.

Carston, R. and A. Hall (2011). Implicature and Explicature. In H. D. Schmid and D. Geeraerts (Eds.), Cognitive Pragmatics: Handbooks in Pragmatics vol. 4. Mouton de Gruyter.

Casati, R. (2004). Is the object concept formal? Dialectica 58(3), 383–394.

Chatzikyriakidis, S. (2010). Clitics in four dialects of Modern Greek: A dynamic Account. Ph. D. thesis, King's College London.

Chierchia, G. (1992). Anaphora and Dynamic Binding. Linguistics and Philosophy 15(2), 111–183.

Chomsky, N. (1957). Syntactic structures. The Hague: Mouton and Co.

Chomsky, N. (1965). Aspects of the Theory of Syntax. Cambridge, MA: MIT Press.

Chomsky, N. (1981). Lectures on Government and Binding. Dordrecht: Foris.

Chomsky, N. (1986). Knowledge of language: Its nature, origin, and use. Praeger Publishers.

Chomsky, N. (1995). The minimalist program. MIT Press.

Clark, A. (2004). Feature-placing and proto-objects. Philosophical Psychology 17(4), 443–469.

Clements, G. (1975). The logophoric pronoun in Ewe: Its role in discourse. Journal of West African Languages 10, 141–77.

Coffa, J. (1993). The semantic tradition from Kant to Carnap: a train to the Vienna station. Cambridge University Press.

Cohen, L. (1971). Some Remarks on Grice's Views about The Logical Particles of Natural Language. In Y. Bar-Hillel (Ed.), Pragmatics of Natural Languages, pp. 50–68. Dordrecht: Reidel.

Coleridge, S. (2007). Biographia literaria. Echo Library.

Cooper, R. (1979). The Interpretation of Pronouns. Syntax and semantics 10, 61–92.

Dehaene, S. (1997). The number sense. New York: Oxford University Press.

Dennis, J. and Z. Pylyshyn (2002). Effect of object discriminability on multiple object tracking. Journal of Vision 2(7), 241.

Dewar, K. and F. Xu (2007). Do 9-month-old infants expect distinct words to refer to kinds? Developmental Psychology (43), 1227–1238.

Donnellan, K. (1966). Reference and Definite Descriptions. The Philosophical Review 75(3), 281–304.

Elbourne, P. (2001). E-type anaphora as NP-deletion. Natural Language Semantics 9(3), 241–288.

Elbourne, P. (2005). Situations and Individuals. MIT Press.

Elbourne, P. (2008). Demonstratives as individual concepts. Linguistics and philosophy 31(4), 409–466.

Elugardo, R. and R. Stainton (2003). Grasping objects and contents. Epistemology of Language, 257–302.

Enç, M. (1986). Towards a referential analysis of temporal expressions. Linguistics and Philosophy 9(4), 405–426.

Evans, G. (1977). Pronouns, quantifiers and relative clauses. Canadian Journal of Philosophy 7, 467–536.

Evans, G. (1982). The Varieties of Reference. Oxford University Press, USA.

Falkum, I. L. (2010). The Semantics and Pragmatics of Polysemy: A Relevance-Theoretic Account. Ph. D. thesis, University College London, London.

Fauconnier, G. (1985). Mental spaces, Volume 1994. Cambridge, MA: MIT Press.

Feigenson, L. and S. Carey (2005). On the limits of infants' quantification of small object arrays. Cognition 97(3), 295–313.

Feigenson, L., S. Carey, and M. Hauser (2002). The representations underlying infants' choice of more: Object files versus analog magnitudes. Psychological Science 13(2), 150.

Fiengo, R. and R. May (1994). Indices and identity. MIT Press.

Filik, R., A. Sanford, and H. Leuthold (2008). Processing pronouns without antecedents: Evidence from event-related brain potentials. Journal of Cognitive Neuroscience 20(7), 1315–1326.

Fodor, J. A. (1975). The Language of Thought. Harvard University Press.

Fodor, J. A. (1983). The modularity of mind: An essay on faculty psychology. MIT Press.

Fox, D. (2000). Economy and Semantic Interpretation. MIT Press.

Franconeri, S., Z. Pylyshyn, and B. Scholl (2006). Spatiotemporal cues for tracking multiple objects through occlusion. In Object Perception, Attention, and Memory 2005 Conference Report, Volume 14, pp. 74–128.

Frege, G. (1948). Sense and reference. The Philosophical Review 57(3), 209–230. Originally published in 1892.

Frege, G. (1950). The foundations of arithmetic. Blackwell. Trans. Austin, J.L.

Frege, G. (1956). The thought: a logical inquiry. Mind 65(259)(1), 289–311.

Frege, G. (1967). Begriffsschrift, a formula language, modelled upon that of arithmetic, for pure thought. In From Frege to Gödel: A Source Book in Mathematical Logic, Volume 1, pp. 1–82. Harvard University Press. Originally published in 1879.

Frege, G. (1980). Philosophical and Mathematical Correspondence. Chicago: University of Chicago Press. ed. by G. Gabriel, H. Hermes, F. Kambartel, C. Thiel, and A. Veraart, trans. by H. Kaal.

Gabbay, D. (1996). Labelled Deductive Systems. Oxford University Press.

Galery, T. (2008). Singular content and deferred uses of indexicals. UCL Working Papers in Linguistics 20, 157–194.

Geach, P. (1962). Reference and Generality. Ithaca, NY: Cornell University Press.

Gerrig, R. (1989). The Time Course of Sense Creation. Memory & Cognition 17(2), 194.

Gibbs, R. and J. Moise (1997). Pragmatics in understanding what is said. Cognition 62(1), 51–74.

Gordon, R. and D. Irwin (1996). What's in an object file? Evidence from priming studies. Perception and Psychophysics 58(8), 1260–1277.

Gordon, R. and D. Irwin (2000). The role of physical and conceptual properties in preserving object continuity. Journal of Experimental Psychology: Learning, Memory, and Cognition 26(1), 136–150.

Grandy, R. (1989). On Grice on Language. The Journal of Philosophy 86(10), 514–525.

Green, M. (1998). Direct Reference and Implicature. Philosophical Studies 91(1), 61–90.

Gregoromichelaki, E. (2006). Conditionals in Dynamic Syntax. Ph. D. thesis, King's College London, London.

Gregoromichelaki, E., R. Kempson, and G. Mills (2011). NL grammars as models for language use: incrementality and context integration. Presented at the Meaning and Context Conference, UWE.

Grice, H. P. (1957). Meaning. The Philosophical Review 66, 377–88.

Grice, H. P. (1967). Logic and conversation. William James Lectures. Published in Grice (1975).

Grice, H. P. (1969). Vacuous names. In D. Davidson and J. Hintikka (Eds.), Words and Objections, pp. 118–45. D. Reidel.

Grice, H. P. (1975). Logic and Conversation. Syntax and Semantics 3: Speech Acts, 41–58.

Grice, H. P. (1978). Further Notes on Logic and Conversation. Syntax and Semantics 9: Pragmatics, 113–127.

Grice, H. P. (1981). Presupposition and conversational implicature. In Radical Pragmatics, pp. 183–198. New York: Academic Press.

Grice, H. P. (1989a). Retrospective epilogue. Grice, 1989b, 339–386.

Grice, H. P. (1989b). Studies in the Way of Words. Cambridge: Harvard University Press.

Grimberg, M.-L. (1994). On Nunberg on Indexicality and Deixis. UCL Working Papers in Linguistics 6, 341–378.

Grimberg, M.-L. (1996). Against Rigidity. Ph. D. thesis, University College London.

Haegeman, L. (1994). Introduction to government and binding theory. Oxford: Blackwell.

Hagège, C. (1974). Les pronoms logophoriques. Bulletin de la Société de Linguistique de Paris 69(1), 287–310.

Hankamer, J. and I. Sag (1976). Deep and surface anaphora. Linguistic Inquiry 7(3), 391–428.

Hawthorne, J. (2006). Testing for Context-Dependence. Philosophy and Phenomenological Research 73(2), 443–450.

Heim, I. (1990). E-type pronouns and donkey anaphora. Linguistics and Philosophy 13, 137–78.

Heim, I. and A. Kratzer (1998). Semantics in Generative Grammar. Oxford: Blackwell Publishers.

Higginbotham, J. (1994). Priorities in the Philosophy of Thought. Proceedings of the Aristotelian Society, Supplementary Volumes 68, 85–130.

Higginbotham, J. (2002). Competence with demonstratives. Noûs 36, 1–16.

Hilbert, D. and P. Bernays (1939). Die Grundlagen der Mathematik II. New York: Springer. Second Edition, 1970 reprint.

Hugly, P. and C. Sayward (1979). A problem about conversational implicature. Linguistics and Philosophy 3(1), 19–25.

Huntley-Fenner, G., S. Carey, and A. Solimando (2002). Objects are individuals but stuff doesn't count: Perceived rigidity and cohesiveness influence infants' representations of small groups of discrete entities. Cognition 85(3), 203–221.

Itti, L. and C. Koch (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. Vision research 40(10-12), 1489–1506.

Jacobson, P. (2000). Paycheck pronouns, Bach-Peters sentences, and variable-free semantics. Natural Language Semantics 8(2), 77–155.

Kadmon, N. (1987). On the unique and non-unique reference and asymmetric quantification. Ph. D. thesis, University of Massachusetts, Amherst.

Kahneman, D., A. Treisman, and B. J. Gibbs (1992). The Reviewing of Object Files: Object-Specific Integration of Information. Cognitive Psychology 24(2), 175–219.

Kamp, H. (1971). Formal properties of 'now'. Theoria 37, 227–73.

Kamp, H. (1981). A theory of truth and semantic representation. In J. A. Groenendijk, T. Janssen, and M. Stokhof (Eds.), Formal Methods in the Study of Language, pp. 277–322. Dordrecht: Foris.

Kaplan, D. (1977). Demonstratives. Manuscript, references to the reprint in Almog et al. (1989), pp. 481-563.

Kaplan, D. (1978). Dthat. Syntax and Semantics 9: Pragmatics, 221–43.

Kaplan, D. (1989). Afterthoughts. See Almog et al. (1989), pp. 565 – 614.

Kempson, R. (2010). Language as systems for on-line concept construction. Talk given at the CSMN workshop on word meaning.

Kempson, R., E. Gregoromichelaki, M. Purver, A. Eshghi, R. Cann, C. Howes, W. Meyer-Viol, J. Hough, G. White, and P. Healey (2011). On the processing of elliptical fragments. A preliminary report.

Kempson, R., W. Meyer-Viol, and D. Gabbay (2001). Dynamic Syntax: The Flow of Language Understanding. Wiley-Blackwell.

Kinzler, K. and E. Spelke (2007). Core systems in human cognition. Progress in Brain Research 164, 257–264.

Kjøll, G. (2010). Words, Concepts, and Ontology from the Perspective of Radical Pragmatics and Semantic Externalism. Ph. D. thesis, University of Oslo.

Korta, K. (1997). Implicitures: cancellability non-detachability. Technical report, Institute for Logic, Cognition, Language and Information (ILCLI), San Sebastian.

Kratzer, A. (1989). An investigation of the lumps of thought. Linguistics and Philosophy 12(6), 607–53.

Krifka, M. (1987). An Outline of Genericity. Partly in collaboration with C. Gerstner. SNS-Bericht, Universität Tübingen.

Kripke, S. (1972). Naming and necessity. In G. Harman and D. Davidson (Eds.), Semantics of Natural Language, pp. 253–355. Dordrecht: D. Reidel Publishing Company.

Lakoff, G. and M. Johnson (1980). Metaphors We Live by. University of Chicago Press.

Lakoff, G. and M. Turner (1989). More than Cool Reason: A Field Guide to Poetic Metaphor. University of Chicago Press.

Lasersohn, P. (2005). Context Dependence, Disagreement, and Predicates of Personal Taste. Linguistics and Philosophy 28(6), 643–686.

Levinson, S. (2000). Presumptive Meanings: The Theory of Generalized Conversational Implicature. Cambridge, MA: MIT Press.

Lewis, D. (1969). Convention. Cambridge, MA: Harvard University Press.

Lewis, D. (1983). Languages and Language. In Philosophical Papers, Volume I, pp. 163–188. Oxford: Oxford University Press.

Link, G. (1983). The Logical Analysis of Plural and Mass Nouns: a Lattice Theoretic Approach. Meaning, use and interpretation of language, de Gruyter, Berlin, 302–323.

MacFarlane, J. (2003). Future contingents and relative truth. Philosophical Quarterly 212(7), 321–36.

MacFarlane, J. (2005). Making sense of relative truth. Proceedings of the Aristotelian Society 105, 321–339.

MacFarlane, J. (2007). Relativism and disagreement. Philosophical Studies 132(1), 17–31.

MacFarlane, J. (2009). Nonindexical contextualism. Synthese 166, 231–250.

Macnamara, J. (1972). Cognitive basis of language learning in infants. Psychological review 79(1), 1–13.

Macnamara, J. (1982). <u>Names for things: A study in human learning</u>. Cambridge, MA: MIT Press.

Macnamara, J. (1986). <u>A border dispute: The place of logic in psychology</u>. Cambridge, MA: MIT Press.

Marr, D. (1982). <u>Vision: A computational investigation into the human representation and processing of visual information</u>. New York: Henry Holt and Co.

Martí, L. (2006). Unarticulated constituents revisited. <u>Linguistics and Philosophy</u> <u>29</u>(2), 135–166.

Martí, L. (2011). Implicit indefinite objects: grammar, not pragmatics. Manuscript.

Matsui, T. (1998). Pragmatic criteria for reference assignment: A relevance-theoretic account of the acceptability of bridging. <u>Pragmatics and cognition</u> <u>6</u>(1), 47–97.

Matsui, T. (2000). <u>Bridging and relevance</u>. John Benjamins.

Montague, R. (1970a). English as a formal language. In B. Visentini (Ed.), <u>Linguaggi nella societa e nella tecnica</u>, pp. 189–223. Milan: Edizioni di Comunita.

Montague, R. (1970b). Universal grammar. <u>Theoria</u> <u>36</u>, 373–98.

Montague, R. (1973). The proper treatment of quantification in ordinary english. In K. J. M. P. Suppes (Ed.), <u>Approaches to Natural Language</u>, pp. 221–242. Dordrecht: Reidel.

Moore, G. E. (1918). Some judgments of perception. <u>Proceedings of the Aristotelian Society</u> <u>19</u>, 1–29.

Moreno, R. (2007). <u>Creativity and Convention: The Pragmatics of Everyday Figurative Speech</u>. John Benjamins Publishing Co.

Mount, A. (2008). The impurity of pure indexicals. <u>Philosophical Studies</u> <u>138</u>(2), 193–209.

Neale, S. (1990). <u>Descriptions</u>. Cambridge, Mass.: MIT Press.

Neale, S. (1992). Paul Grice and the Philosophy of Language. <u>Linguistics and Philosophy</u> <u>15</u>(5), 509–559.

Neale, S. (1999). Colouring and composition. In K. Murasugi and R. Stainton (Eds.), <u>Philosophy and Linguistics</u>, pp. 35–82. Boulder: Westview Press.

Needham, A. and R. Baillargeon (1998). Effects of prior experience on 4.5-month-old infants' object segregation. <u>Infant Behavior and Development</u> <u>21</u>, 1–24.

Ness, A. (2011). Thematic unity in the phenomenology of thinking. <u>Philosophical Quarterly</u>, 1–22.

Nicolle, S. and B. Clark (1999). Experimental Pragmatics and What Is Said: A Response to Gibbs and Moise. <u>Cognition</u> <u>69</u>(3), 337–54.

Noveck, I. (2004). Pragmatic inferences related to logical terms. In I. Noveck and D. Sperber (Eds.), <u>Experimental Pragmatics</u>, pp. 301–321. Palgrave Macmillan.

Noveck, I. and D. Sperber (2007). The why and how of experimental pragmatics: The case of 'scalar inferences'. In N. Burton-Roberts (Ed.), <u>Pragmatics</u>, pp. 184–212. Basingstoke: Palgrave.

Nunberg, G. (1978). <u>The Pragmatics of Reference</u>. Indiana University Linguistics Club.

Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. <u>Linguistics and Philosophy</u> (3), 143–184.

Nunberg, G. (1990). Indexicality in contexts. Paper delivered at the conference on Philosophy and Cognitive Science, Cerisy-la-Salle, France.

Nunberg, G. (1993). Indexicality and Deixis. <u>Linguistics and Philosophy</u> <u>16</u>(1), 1–43.

Nunberg, G. (1995). Transfers of meaning. <u>Journal of semantics</u> <u>12</u>(2), 109.

Nunberg, G. (2004a). Indexical descriptions and descriptive indexicals. In M. Reimer and A. Bezuidenhout (Eds.), <u>Descriptions and Beyond: An Interdisciplinary Collection of Essays on Definite and Indefinite Descriptions</u>, pp. 261–279. Oxford University Press.

Nunberg, G. (2004b). The pragmatics of deferred interpretation. In <u>Handbook of pragmatics</u>, pp. 344–364.

Olshausen, B. and C. Koch (1995). Selective visual attention. In M. A. Arbib (Ed.), <u>The handbook of brain theory and neural networks</u>. MIT Press.

Palmer, S. (1999). Vision Science: Photons to Phenomenlogy. MIT Press.

Papafragou, A. (1996). On metonymy. Lingua 99(4), 169–195.

Papafragou, A. (2000). Modality: issues in the semantics-pragmatics interface. Elsevier, Amsterdam; New York.

Partee, B. and V. Borschev (1998). Integrating lexical and formal semantics: Genitives, relational nouns, and type-shifting. In Proceedings of the Second Tbilisi Symposium on Language, Logic, and Computation, pp. 229–241. Citeseer.

Patel, P., P. Grosz, E. Fedorenko, and T. Gibson (2009). Experimental evidence against a strict version of the formal link condition on e-type pronouns. Poster presented at the 22nd CUNY conference on human sentence processing.

Perry, J. (1977). Frege on demonstratives. The Philosophical Review 86(4), 474–97.

Perry, J. (1979). The problem of the Essential Indexical. Nous, 3–21.

Perry, J. (1986). Thought without representation. Supplementary Proceedings of the Aristotelian Society 60, 137–52.

Perry, J. (2001). Reference and reflexivity. CSLI Publications.

Pollard, C. and I. Sag (1992). Anaphors in English and The Scope of Binding Theory. Linguistic Inquiry 23(2), 261–303.

Pollard, C. and I. Sag (1994). Head-driven phrase structure grammar. University of Chicago Press.

Postal, P. (1969). Anaphoric Islands. In Chicago Linguistic Society, Volume 5, pp. 205–239.

Powell, G. (1998). The deferred interpretation of indexicals and proper names. UCL Working Papers in Linguistics 10, 143–172.

Powell, G. (2001). Compositionality and semantic theory. 75th annual meeting of the Pacific Division of the American Philosophical Association.

Powell, G. (2003). Language, Thought and Reference. Ph. D. thesis, University College London, London.

Purver, M., E. Gregoromichelaki, W. Meyer-Viol, and R. Cann (2010). Splitting the 'I's and crossing the 'you's: Context, speech acts and grammar. In Aspects of Semantics and Pragmatics of Dialogue. SemDial 2010, 14th Workshop on the Semantics and Pragmatics of Dialogue.

Pylyshyn, Z. (2001). Visual indexes, preconceptual objects, and situated vision. Cognition 80(1-2), 127–158.

Pylyshyn, Z. (2007). Things and places: How the mind connects with the world (Jean Nicod Lectures Series).

Pylyshyn, Z. and R. Storm (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. Spatial Vision 3(3), 179–197.

Quine, W. V. (1960). Word and Object. Cambridge: Technology Press of the Massachusetts Institute of Technology.

Quine, W. V. (1968). Ontological relativity. The Journal of Philosophy 65(7), 185–212.

Recanati, F. (1989). The pragmatics of what is said. Mind and Language (4), 295–329.

Recanati, F. (1993). Direct Reference: From Language to Thought. Oxford: Blackwell.

Recanati, F. (1995). The alleged priority of literal interpretation. Cognitive Science 19(2), 207–232.

Recanati, F. (1999). Situations and the structure of content. In K. Murasugi and R. Stainton (Eds.), Philosophy and Linguistics, pp. 113–165. Boulder: Westview Press.

Recanati, F. (2004). Literal meaning. Cambridge University Press.

Recanati, F. (2007). Perspectival Thought. New York: Oxford University Press.

Recanati, F. (2009). Compositionality, Flexibility, and Context-Dependence. In The Oxford Handbook of Compositionality, pp. 175–191. Oxford University Press.

Reichenbach, H. (1947). Elements of Symbolic Logic. New York: MacMillan.

Reinhart, T. (1976). The Syntactic Domain of Anaphora. Ph. D. thesis, MIT, Cambridge, MA.

Reinhart, T. and E. Reuland (1993). Reflexivity. Linguistic Inquiry 24(4), 657–720.

Rose, S., J. Feldman, and J. Jankowski (2001). Attention and recognition memory in the 1st year of life: a longitudinal study of preterm and full-term infants. Developmental Psychology 37(1), 135–151.

Ross-sheehy, S., L. Oakes, and S. Luck (2003). The Development of Visual Short-Term Memory Capacity in Infants. Child Development 74(6), 1807–1822.

Rubio Fernandez, P. (2005). Pragmatic Processes and Cognitive Mechanisms in Lexical Interpretation. Ph. D. thesis, University of Cambridge, Cambridge.

Russell, B. (1904). Letter to Frege. published in (Frege, 1980, 166-70).

Russell, B. (1905). On denoting. Mind 14, 479–493.

Russell, B. (1910). Knowledge by acquaintance and knowledge by description. Proceedings of the Aristotelian Society 11, 108–128.

Sag, I. (1981). Formal semantics and extralinguistic context. Radical pragmatics, 273–94.

Salmon, N. (1986). Frege's Puzzle. MIT Press/Bradford Books.

Sandburg, C. (1950). Complete poems. Harcourt, Brace.

Sanford, A., R. Filik, C. Emmott, and L. Morrow (2008). They're digging up the road again: The processing cost of Institutional 'They'. The Quarterly Journal of Experimental Psychology 61(3), 372–380.

Sauerland, U. (2008). The semantic markedness of phi-features. In Phi Theory, pp. 57–82. Oxford University Press.

Schiffer, S. (1972). Meaning. Oxford: Oxford University Press.

Schiffer, S. (1981). Indexicals and the Theory of Reference. Synthese 49(1), 43–100.

Schlenker, P. (2003). A plea for monsters. Linguistics and Philosophy 26, 29–120.

Schlenker, P. (2006). Ontological symmetry in language: A brief manifesto. Mind & language 21(4), 504–539.

Scholl, B. and Z. Pylyshyn (1999). Tracking multiple items through occlusion: Clues to visual objecthood. Cognitive Psychology 38(2), 259–290.

Scholl, B., Z. Pylyshyn, and J. Feldman (2001). What is a visual object? Evidence from target merging in multiple object tracking. Cognition 80(1-2), 159–177.

Scholl, B., Z. Pylyshyn, and S. Franconeri (1999). When are spatiotemporal and featural properties encoded as a result of attentional allocation. Paper presented at the meeting of the association of investigative ophthalmology and visual science.

Searle, J. (1969). Speech Acts: An Essay in the Philosophy of Language. Cambridge: Cambridge University Press.

Sells, P. (1987). Aspects of logophoricity. Linguistic Inquiry 18, 445–79.

Smith, E. and D. Medin (1981). Categories and Concepts. Harvard University Press Cambridge, MA.

Spelke, E. (1990). Principles of object perception. Cognitive science 14(1), 29–56.

Spelke, E., R. Kestenbaum, D. Simons, and D. Wein (1995). Spatiotemporal continuity, smoothness of motion and object identity in infancy. British Journal of Developmental Psychology 13, 113–142.

Spelke, E. and K. Kinzler (2007). Core knowledge. Developmental Science 10(1), 89–96.

Sperber, D., F. Clément, C. Heintz, O. Mascaro, H. Mercier, G. Origgi, and D. Wilson (2010). Epistemic vigilance. Mind & Language 25(4), 359–393.

Sperber, D. and D. Wilson (1986). Loose talk. Proceedings of the Aristotelian Society 86, 153–171.

Sperber, D. and D. Wilson (1987). Precis of Relevance: Communication and cognition. Behavioral and Brain Sciences 10(04), 697–710.

Sperber, D. and D. Wilson (1995). Relevance: Communication and Cognition (2nd ed.). Blackwell.

Sperber, D. and D. Wilson (1996). Fodor's frame problem and relevance theory (reply to Chiappe & Kukla). Behavioral and Brain Sciences 19(3), 530–532.

Sperber, D. and D. Wilson (2008). A Deflationary Account of Metaphor. In R. Gibbs (Ed.), The Cambridge Handbook of Metaphor and Thought, pp. 84–109. Cambridge: Cambridge University Press.

Stanley, J. (2000). Context and logical form. Linguistics and Philosophy 23(4), 391–434. All citations to the reprint in Stanley (2007), 30–68.

Stanley, J. and Z. Szabó (2000). On Quantifier Domain Restriction. Mind & Language 15(2&3), 219–261.

Stokke, A. (2008). Descriptive uses of indexicals as implicatures. Unpublished Manuscript.

Stokke, A. (2010). Indexicality and Presupposition. Ph. D. thesis, University of St. Andrews, St. Andrews.

Strawson, P. F. (1950). On Referring. Mind 59(235), 320–344.

Strawson, P. F. (1959). Individuals: An essay in Descriptive Metaphysics. London: Methuen.

Strawson, P. F. (1964). Intention and Convention in Speech Acts. The Philosophical Review 73(4), 439–460.

Surian, L., S. Caldi, and E. Piredda (2004). Agents' motion and the development of object individuation. Poster presented at the 9th biennial meeting of the International Conference on Infant Studies.

Treisman, A. and G. Gelade (1980). A feature-integration theory of attention. Cognitive psychology 12(1), 97–136.

Trick, L. and Z. Pylyshyn (1994a). Cueing and counting: Does the position of the attentional focus affect enumeration? Visual Cognition 1(1), 67–100.

Trick, L. and Z. Pylyshyn (1994b). Why are small and large numbers enumerated differently? A limited-capacity preattentive stage in vision. Psychological Review 101, 80–80.

VanMarle, K. and B. Scholl (2003). Attentive tracking of objects vs. substances. Psychological Science 14(5), 498–504.

von Fintel, K. and A. Gillies (2008). CIA leaks. The Philosophical Review 117(1), 77–98.

Ward, G. (2008). Equatives and deferred reference. In J. Gundel and N. Hedberg (Eds.), Reference: Interdisciplinary Perspectives, pp. 73–94. Oxford: Oxford University Press.

Wettstein, H. (1986). Has Semantics Rested on a Mistake? The Journal of Philosophy 83(4), 185–209.

Wharton, T. (2003). Natural pragmatics and natural codes. Mind & language 18(5), 447–477.

Wiggins, D. (1967). Identity and Spatio-Temporal Continuity. Oxford: Blackwell.

Wiggins, D. (1980). Sameness and Substance. Oxford: Blackwell.

Wiggins, D. (2001). Sameness and Substance Renewed. Cambridge: Cambridge University Press.

Wilson, D. and R. Carston (2006). Metaphor, relevance and the 'emergent property' issue. Mind & Language 21(3), 404–433.

Wilson, D. and R. Carston (2007). A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), Pragmatics, pp. 230–259. Palgrave.

Wilson, D. and T. Matsui (2000). Recent approaches to bridging: Truth, coherence, relevance. Lengua, discurso, texto 1, 103–132.

Wilson, D. and D. Sperber (1988). Representation and Relevance. In R. Kempson (Ed.), Mental Representations: The Interface between Language and Reality, pp. 133–153.

Wilson, D. and D. Sperber (1993). Linguistic Form and Relevance. Lingua 90(2), 1–25.

Wilson, D. and D. Sperber (2002). Truthfulness and Relevance. Mind 111(443), 583–632.

Wilson, D. and D. Sperber (2006). Relevance theory. In L. Horn and G. Ward (Eds.), Handbook of pragmatics, pp. 607–632. Oxford: Wiley-Blackwell.

Wittgenstein, L. (1953). Philosophical investigations. New York: Macmillan.

Wittgenstein, L. (1969). On Certainty. New York: Harper and Row.

Woodfield, A. (1991). Conceptions. Mind 100(399), 547–72.

Xu, F. (1999). Object individuation and object identity in infancy: The role of spatiotemporal information, object property information, and language. Acta Psychologica 102(2-3), 113–136.

Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. Cognition 85(3), 223–250.

Xu, F. (2005). Categories, kinds, and object individuation in infancy. Building object categories in developmental time, 63–89.

Xu, F. and S. Carey (1996). Infants' metaphysics: The case of numerical identity. Cognitive Psychology 30(2), 111–153.