Chapter 1. Introduction

Karttunen's seminal paper on *Discourse Referents* begins with the following exhortation:

> "Consider a device designed to read a text in some natural language, interpret it and store the content in some manner, say, for the purpose of being able to answer questions about it. To accomplish this task, the machine will have to fulfill at least the following basic requirement. It has to be able to build a file that consists of records of all the individuals, that is, events, objects, etc., mentioned in the text and, for each individual, record whatever is said about it."
> (Karttunen 1976: 364)

The abstract study of such a text interpreter ultimately comes down to a reconsideration of the nature of the (literal) meaning of natural language expressions. What is being reconsidered is the truth-conditional, referential theory of meaning, aptly summarized in Hintikka (1969) as follows:

> "[…] it seems to me in any case completely hopeless to try to divorce the idea of the meaning of a sentence from the idea of the *information* that that the sentence can convey to a hearer or reader, should someone truthfully address it to him. Now what is this information? Clearly it is just information to the effect that the sentence is true, that the world is such as to meet the truth-conditions of the sentence."
> (Hintikka 1969: 146)

Moreover, the truth-conditions are determined in terms of the reference (denotation) relations that hold between linguistic expressions and independent, extra-linguistic entities.

Karttunen's exhortation shifts the classical perspective on meaning in two ways. First, the central problem is not the interpretation of sentences in isolation, but the interpretation of texts, i.e. of discourse. As Kamp (2001) puts it, "discourse meanings are more than plain conjunctions of sentence meanings. And this 'more' is often the effect of interpretation principles that are an integral part of linguistic knowledge, and thus legitimate objects of linguistic study" (Kamp 2001: 57).

Thus, over and above the investigation of how natural language interpretation is context dependent, we also need to investigate how the interpretation of a natural language expression changes the context, i.e. it creates a new context out of the old one, and thereby affects how subsequent expressions are interpreted. As soon as we turn to discourse interpretation, the dynamics of meaning comes into focus and we shift from a static truth-conditional theory to a theory of meaning as information update.

The second way in which the perspective on meaning shifts is reflected in the title of Karttunen's paper: natural language interpretation crucially involves a notion of *discourse reference* which mediates between linguistic expressions and their reference in the classical sense. This is the basic requirement put forth for our abstract text interpreter: the interpreter incrementally builds a file that contains records of the individuals mentioned in the text. At any given point, the file encodes the current information state of the interpreter, i.e. the current discourse context, and we refer to actual individuals via this information state: reference in natural language is inextricably discourse reference.

The present work is part of the general project of investigating the notions of discourse reference and information state involved in natural language interpretation. In particular, I will argue that over and above the basic requirement that the information state should be able to detect when a novel individual is mentioned in discourse and "store it along with its characterization for future reference" (Karttunen 1976: 364), the information state should also be able to encode dependencies between individuals (or sets thereof) that are established and subsequently referred to in discourse.

The main proposal is that nominal and modal expressions introduce (non-singleton) sets of objects, i.e. individuals and possible worlds respectively, and that these sets are correlated in discourse: discourse reference involves two equally important components with essentially the same interpretive dynamics, namely reference to values, i.e. sets of objects, and reference to structure, i.e. the correlation / dependency between such sets, which is introduced and incrementally elaborated upon in discourse. Hence the title: *Structured Nominal and Modal Reference*.

The dissertation focuses on the development of a new dynamic system couched in classical type logic that formalizes this notion of discourse reference and in which natural language discourses involving complex descriptions of multiple related objects can be compositionally translated.

I will, therefore, have little to say about the problem of interfacing the resulting system with a more general theory of discourse structure and anaphora resolution of the kind pursued in Hobbs (1979, 1990, 1993), Kameyama (1994), Grosz et al (1995), Kehler (1995, 2002) and Asher & Lascarides (2003) among others. In particular, one of the most important assumptions I will make throughout the dissertation is that the logical form of sentences and discourses comes with the 'intended' anaphoric indexation.

The dissertation consists of eight chapters, the first one (i.e. the current one) being the introduction and last one the conclusion. Chapters **2**, **3** and **4** mostly review the previous literature and introduce the basic type-logical dynamic system that underlies the entire present investigation. The remaining chapters introduce the new contributions: chapters **5** and **6** are dedicated to the study of donkey anaphora and quantificational subordination respectively (i.e. structured reference in the nominal domain), while chapter **7** is dedicated to the study of modal anaphora and modal subordination (i.e. structured reference in the modal domain). A more detailed description follows.

Chapter 2. Dynamic Predicate Logic with Generalized Quantification

Chapter **2** formally explicates Karttunen's basic requirement for discourse reference. The three most important formal systems modeling the notion of discourse reference in Karttunen (1976) are Discourse Representation Theory (DRT, Kamp 1981), File Change Semantics (FCS, Heim 1982/1988) and Dynamic Predicate Logic (DPL, Groenendijk & Stokhof 1991). The empirical coverage of these classical systems is roughly the same and, notwithstanding several non-trivial differences between them, they all follow the insight in Lewis (1975) that a dynamic information state (in terms of which discourse reference is to be defined) is a *case*, which is modeled as an *assignment of values to variables*.

> "What is a case? […] [A] case may be regarded as the 'tuple of its participants; and these participants are values of the variables that occur free in the open sentence modified by the adverb. In other words, we are taking the cases to be the admissible assignments of values to these variables."
> (Lewis 1975: 5-7)

The particular version of dynamic semantics surveyed in chapter **2** is DPL – and for three reasons. First, the syntax of the system is the familiar syntax of classical first-order logic (at least in the original notation; the notation in chapter **2** is a close variant thereof); this enables us to focus on what is really new, namely the semantics.

Second, the semantics of DPL is minimally different from the standard Tarskian semantics for first-order logic: instead of interpreting a formula as a set of variable assignments (i.e. the set of variable assignments that satisfy the formula in the given model), we interpret it as a binary relation between assignments; moreover, this minimal semantic modification encodes in a transparent way the core dynamic idea that meaning is not merely truth-conditional content, but context change potential.

Third, just as classical first-order logic can be generalized to static type logic, DPL can be generalized to a dynamic version of type logic and we can thereby introduce compositionality at the sub-sentential / sub-clausal level – to which chapter **3** is dedicated.

Besides introducing it, chapter **2** extends DPL with a dynamic notion of selective generalized quantification (as opposed to the unselective generalized quantification of Lewis 1975). This dynamic notion of selective generalized quantification has been proposed in various guises by many authors: Bäuerle & Egli (1985), Root (1986) and Rooth (1987) put forth the basic proposal and van Eijck & de Vries (1992) and Chierchia (1992, 1995) were the first to formulate it in DPL terms; the proposal is also adopted in Heim (1990) and Kamp & Reyle (1993).

Selective generalized quantification enables us to solve the proportion problem and account for weak / strong donkey readings.

The proportion problem is exemplified by sentence (1) below, where, intuitively, we do not quantify over most pairs <*x*, *y*> such that *x* is a house-elf that falls in love with a

witch *y* – as unselective quantification would have it – but only over most house-elves *x* that fall in love with some witch or other.

1. Most house-elves who fall in love with a witch buy her an alligator purse.

The weak / strong ambiguity problem is posed by the donkey sentence in (2) below, which has a weak reading, in contrast to the classical donkey sentence in (3), which has a strong reading. Sentence (2) exhibits a weak reading in the following sense: its most salient interpretation is that every person who has a dime will put *some* dime s/he has in the meter (and not all her / his dimes). Sentence (3) exhibits a strong reading in the sense that its most salient interpretation is that every farmer beats *every* donkey s/he owns.

2. Every person who has a dime will put it in the meter. (Pelletier & Schubert 1989)
3. Every farmer who owns a donkey beats it.

However, the notion of selective generalized quantification introduced to account for the weak / strong donkey examples in (2) and (3) above cannot *compositionally* account for the mixed weak & strong relative-clause donkey sentences in (4) and (5) below.

4. Every person who buys a book on amazon.com and has a credit card uses it to pay for it.
5. Every man who wants to impress a woman and who has an Arabian horse teaches her how to ride it.

Consider sentence (4): it is interpreted as asserting that, for *every* book (strong) that any credit-card owner buys on amazon.com, there is *some* credit card (weak) that s/he uses to pay for the book. Note, in particular, that the credit card can vary from book to book, e.g. I can use my MasterCard to buy set theory books and my Visa to buy detective novels, which means that even the weak indefinite *a credit card* can introduce a (possibly) non-singleton set.

For each buyer, the two sets of objects, i.e. all the books purchased on amazon.com and some of the credit cards that the buyer has, are correlated and the dependency between these sets – left implicit in the restrictor of the quantification – is specified in the

nuclear scope: each book is correlated with the credit card that was used to pay for it. The above paraphrase of the meaning of sentence (4) is formalized in classical (static) first-order logic as shown in (6) below.

6. $\forall x(person(x) \wedge \exists y(book(y) \wedge buy\_on\_amazon(x, y)) \wedge \exists z(c.card(z) \wedge have(x, z))$
$\rightarrow \forall y'(book(y') \wedge buy\_on\_amazon(x, y')$
$\rightarrow \exists z'(c.card(z') \wedge have(x, z') \wedge use\_to\_pay(x, z', y')))$

As the first-order translation in (6) explicitly shows, the challenge posed by sentence (4) is to *compositionally* derive its interpretation while allowing for: (**i**) the fact that the two donkey indefinites in the restrictor of the quantification receive two distinct readings (strong and weak respectively) and (**ii**) the fact that the value of the weak indefinite *a credit card* co-varies with / is dependent on the value of the strong indefinite *a book* although, by the Coordinate Structure Constraint, the strong indefinite cannot syntactically scope over the weak one since both DP's are trapped in their respective VP-conjuncts. Example (5) is a variation on the same theme.

The fact that DPL with selective quantification cannot compositionally account for the mixed reading relative-clause donkey sentences in (4) and (5) above provides the basic plot and motivation for the next three chapters of the dissertation, namely chapters **3**, **4** and **5**. In particular, chapters **3** and **4** endeavor to recast DPL and its extension with selective generalized quantification in classical (many-sorted) type logic, which will automatically enable us to define a compositional interpretation procedure of the kind available in Montague semantics.

Chapter 3. Compositional DRT

Chapter **3** is the last chapter that is almost entirely a review of the previous literature. In particular, I review Compositional DRT (CDRT, Muskens 1996), which generalizes DPL to a dynamic version of type logic (just as static type logic generalizes classical first-order logic). This move enables us to introduce compositionality at the sub-sentential / sub-clausal level in the tradition of Montague semantics.

The first part of the chapter is dedicated to the definition of Dynamic Ty2 (which is Muskens' Logic of Change with negligible modifications) and to the translation of DPL in Dynamic Ty2. The Dynamic Ty2 translation is shown to preserve the DPL account of cross-sentential anaphora, relative-clause donkey sentences and conditional donkey sentences.

The second part defines a type-drive translation procedure based on a rough-and-ready syntax for a fragment of English. The resulting CDRT system effectively unifies DPL and Montague semantics and enables us to compositionally account for a variety of anaphoric and quantificational phenomena, including bound variable anaphora, quantifier scope ambiguities and donkey anaphora.

Chapter 4. Compositional DRT with Generalized Quantification

Chapter **4** is the first one that adds to the previous literature in a more substantial way by translating in Dynamic Ty2 the DPL-style definitions of unselective and selective generalized quantification introduced in chapter **2**.

CDRT is then extended with these two notions of dynamic generalized quantification. The resulting system, which I label CDRT+GQ, provides a fully compositional account of the proportion problem and of the simple (non-mixed) examples of weak / strong donkey sentences.

The chapter also introduces the analysis of the interaction between anaphora and generalized coordination in Muskens (1996). I show that this analysis successfully generalizes to account for DP-conjunction donkey sentences like *Every boy who has a dog and every girl who has a cat must feed it* due to Chierchia (1995).

Such examples are interesting for two reasons. First, Chierchia (1995) uses DP-conjunction donkey sentences of this kind to argue in favor of an approach to natural language interpretation that builds (part of) the dynamics into the semantic value of natural language expressions and against approaches that build the dynamics of the interpretation into syntactic operations at the level of Logical Form (LF).

The argument is, in a nutshell, that the same donkey pronoun is semantically bound by two distinct donkey indefinites, which can be naturally accounted for in a dynamic type-logical system with generalized conjunction (generalized to arbitrary types in the sense of Partee & Rooth 1983 among others). This kind of 'double binding', however, presents difficulties for approaches that require a particular syntactic configuration at the level of LF for the donkey pronouns to be semantically bound – because the same pronoun cannot enter two such distinct syntactic configurations.

The second reason for examining DP-conjunction donkey sentences is that, in the following chapter (chapter **5**), they will help us distinguish between different accounts of mixed reading donkey sentences.

Chapter **4** concludes with the somewhat surprising observation that CDRT with generalized quantification (CDRT+GQ) inherits the problem of DPL with generalized quantification: CDRT+GQ is not compositional enough to account for the mixed weak & strong relative-clause donkey sentences in (4) (and (5)) above. The main difficulty is due to the fact that, in (4), the weak indefinite *a credit card* co-varies with / is dependent on the value of the strong indefinite *a book*, although the strong indefinite cannot syntactically scope over the weak one (recall that they are both trapped in a coordination island).

It will be the task of the following chapter to modify the notion of information state employed in CDRT+GQ (and inherited from DRT / FCS / DPL) and thereby provide a compositional account of mixed weak & strong relative-clause donkey sentences.

Chapter 5. Structured Nominal Reference: Donkey Anaphora

This chapter incrementally introduces a new dynamic system that extends CDRT+GQ and within which we can give a compositional account of the donkey sentences in (4) and (5) above.

The main proposal is that discourse reference involves two equally important components with essentially the same interpretive dynamics, namely reference to values, i.e. sets of objects, and reference to structure, i.e. the correlation / dependency between such sets, which is introduced and incrementally elaborated upon in discourse.

The main technical innovation is that, just as in the Dynamic Plural Logic of van den Berg (1994, 1996a), information states are modeled as *sets* of variable assignments. Such *plural* info states can be represented as matrices with assignments (sequences) as rows, as shown in (7) below.

Plural info states encode discourse reference to both values and structure: the values are the sets of objects that are stored in the *columns* of the matrix, e.g. a discourse referent (dref) *u* for individuals stores a set of individuals relative to a plural info state, since *u* is assigned an individual by each assignment (i.e. row). The structure is *distributively* encoded in the *rows* of the matrix: for each assignment / row in the plural info state, the individual assigned to a dref *u* by that assignment is structurally correlated with the individual assigned to some other dref *u'* by the same assignment. The resulting system is dubbed Plural CDRT.

7. Info State *I*

	…	*u*		*u'*		…
i_1	…	x_1	(i.e. ui_1)	y_1	(i.e. $u'i_1$)	…
i_2	…	x_2	(i.e. ui_2)	y_2	(i.e. $u'i_2$)	…
i_3	…	x_3	(i.e. ui_3)	y_3	(i.e. $u'i_3$)	…
…	…	…		…		…

Values – sets: $\{x_1, x_2, x_3, \ldots\}, \{y_1, y_2, y_3, \ldots\}$ **Structure – relations:** $\{\langle x_1, y_1\rangle, \langle x_2, y_2\rangle, \langle x_3, y_3\rangle, \ldots\}$

In Plural CDRT (PCDRT), sentences denote *relations* between an input and an output plural info state. Indefinites non-deterministically introduce both values and structure, i.e. they introduce *structured sets* of individuals; pronouns are anaphoric to such structured sets. Quantification is defined in terms of matrices instead of single assignments and the semantics of the non-quantificational part becomes rules for how to fill out a matrix.

The PCDRT analysis of sentence (4) is as follows. First, the weak / strong donkey ambiguity is attributed to the indefinite articles. This is the first step towards a compositional account because we *locally* decide for each indefinite article whether it receives a weak or a strong reading. The two basic meanings have the format provided in (8) below.

8. **weak indefinites:** $a^{\mathbf{wk}:u} \rightsquigarrow \lambda P'_{\mathbf{et}}.\, \lambda P_{\mathbf{et}}.\, [u]; P'(u); P(u)$

   strong indefinites: $a^{\mathbf{str}:u} \rightsquigarrow \lambda P'_{\mathbf{et}}.\, \lambda P_{\mathbf{et}}.\, \mathbf{max}^{u}(P'(u); P(u))$

The only difference between a weak and a strong indefinite article is the presence vs. absence of a maximization (**max**) operator. The **max** operator ensures that, after we process a strong indefinite, the output plural info state stores with respect to the dref *u* the *maximal* set of individuals satisfying both the restrictor dynamic property *P′* and the nuclear scope dynamic property *P*. In contrast, a weak indefinite will non-deterministically store *some* set of individuals satisfying its restrictor and nuclear scope.

In sentence (4), the indefinite $a^{\mathbf{str}:u_2}$ *book* in is strong and the indefinite $a^{\mathbf{wk}:u_3}$ *credit card* is weak. Thus, by the time we are done processing the restrictor of the quantification in (4), we will be in an info state that stores: (**i**) the *maximal set* of books with respect to the dref u_2; (**ii**) *some* (non-deterministically introduced) *set* of credit cards with respect to the dref u_3 (the weak indefinite) and (**iii**) *some* (non-deterministically introduced) *structure* correlating the values of u_2 and u_3.

The nuclear scope of the quantification in (4) is anaphoric to both values and structure: we test that the non-deterministically introduced value for u_3 and the non-deterministically introduced structure associating u_3 and u_2 satisfy the nuclear scope condition, i.e., for each assignment in the info state, the u_3-card stored in that assignment is used to pay for the u_2-book stored in the same assignment. That is, the nuclear scope elaborates on the unspecified dependency between u_3 and u_2 introduced in the restrictor of the quantification – and, crucially, introducing such a dependency does not require the strong indefinite to take scope over the weak one.

The PCDRT account successfully generalizes to the mixed reading DP-conjunction donkey sentences in (9) and (10) below, where the same pronoun is intuitively interpreted as having two distinct indefinites as antecedents – and the two indefinites have different readings (one is weak and the other is strong). These examples will be used to distinguish between PCDRT and D-/E-type approaches to donkey anaphora.

9. (Today's newspaper claims that, based on the most recent statistics:)
   Every u_1 company who hired a **str:**u_2 Moldavian man, but no u_3 company who hired a **wk:**u_2 Transylvanian man promoted him u_2 within two weeks of hiring.
10. (Imagine a Sunday fair where people come to sell their young puppies before they get too old and where the entrance fee is one dollar. The fair has two strict rules: all the puppies need to be checked for fleas at the gate and, at the same time, the one dollar bills also need to be checked for authenticity because of the many faux-monnayeurs in the area. So:)
    Everyone u_1 who has a **str:**u_2 puppy and everyone u_3 who has a **wk:**u_2 dollar brings it u_2 to the gate to be checked.

Finally, chapter **5** shows that PCDRT preserves all the previously obtained results, including the analysis of bound variable anaphora, the analysis of quantifier scope ambiguities and the compositional account of the proportion problem and of the simple (non-mixed) examples of weak / strong ambiguities.

Chapter 6. Structured Nominal Reference: Quantificational Subordination

Chapter **6** extends the PCDRT system introduced in chapter **5** with a notion of dynamic generalized quantification that enables us to give a compositional account of quantificational subordination, specifically of the contrast between the interpretations of the following two discourses from Karttunen (1976):

11. **a.** Harvey courts a^u girl at every convention.
    b. She$_u$ is very pretty.
12. **a.** Harvey courts a^u girl at every convention.
    b. She$_u$ always comes to the banquet with him.

The initial sentence (11a/12a) is ambiguous between two quantifier scopings: Harvey courts the same girl at every convention (*a^u girl*>>every convention) vs. at every convention, Harvey courts a (possibly) different girl (*every convention*>>*a^u girl*).

However, discourse (11) as a whole allows only for the "same girl" reading, while discourse (12) allows for both readings.

Using plural information states, we can capture the cross-sentential interaction between quantifier scope and anaphora exhibited by (11), in particular, the fact that the *singular* pronoun in sentence (11b) can disambiguate between the two readings of sentence (11a) (to see that the disambiguation is due to the singular pronoun, replace (11b) with *They$_u$ are very pretty*).

The basic idea is that plural info states enable us to store both quantifier domains (i.e. values) and quantificational dependencies (i.e. structure), pass them across sentential boundaries and further elaborate on them, e.g. by letting a pronoun constrain the cardinality of a previously introduced quantifier domain. More precisely, after processing the update contributed by sentence (11a), the dref u will store the set of all girls that Harvey courts at some convention or other. The singular pronoun *it$_u$* in (11b) will then constrain this set to be a singleton set; hence, the only available reading for discourse (11) as a whole is "wide-scope indefinite" reading.

The fact that discourse (12) is also compatible with the "narrow-scope indefinite" reading is attributed to the presence of the quantificational adverb *always* in (12b), which can take scope over the singular pronoun *she$_u$* and thereby neutralize the effect that singular number morphology has on the cardinality of the previously introduced set of girls.

PCDRT derives the contrast between the two Karttunen examples with minimal stipulations: the dynamic meanings for generalized quantifiers and singular number morphology are basically reformulations of their independently motivated static meanings that incorporate the notion of structured discourse reference argued for in the previous chapters.

Chapter 7. Structured Modal Reference: Modal Anaphora and Subordination

Chapter **7** shows that PCDRT successfully generalizes to other phenomena independent of (yet interacting with) structured donkey anaphora and quantificational subordination, while still preserving both the intuitive appeal of the DRT account of anaphora and the compositional (Montagovian) character of the analyses.

In particular, the fact that we work with finer-grained meanings (given our plural info states) enables us to analyze complex interactions between individual-level and modal anaphora by simply adding discourse referents (dref's) *p*, *p′* etc. for possible worlds: a dref *p* stores a set of worlds, i.e. a *proposition*, relative to a plural info state and it can be structurally correlated with some other modal dref *p′* and / or with individual-level dref's *u*, *u′* etc.

The resulting Intensional PCDRT (IP-CDRT) system generalizes the notion of dynamic quantification introduced in chapter **6** to the modal domain, thereby enabling us to provide an account of the modal subordination that is completely parallel to the account of quantificational subordination. In particular, the analysis of the modal subordination discourse in (13) below is point-for-point parallel to the analysis of the quantificational subordination discourse in (12) above (provided in chapter **6**).

13. **a.** A^{u} wolf might come in. **b.** It_{u} would eat Harvey first.

Thus, IP-CDRT allows us to systematically capture the anaphoric and quantificational parallels between the individual and modal domains argued for in Stone (1999), Bittner (2001) and Schlenker (2005b) among others.

Moreover, chapter **7** also shows that IP-CDRT successfully generalizes to more complex interactions between modal and invidual-level anaphora exhibited by naturally occurring discourses like (14) below.

14. **a.** [A] man cannot live without joy. **b.** Therefore, when he is deprived of true spiritual joys, it is necessary that he become addicted to carnal pleasures. (Thomas Aquinas)

In particular, we are interested in the entailment relation established by *therefore* between the modal premise in (14a) and the modal conclusion in (14b) – and, to capture this, we need to account for several interrelated phenomena.

First, we want to capture the meaning of the entailment particle *therefore*, which relates the content of the premise (14a) and the content of the conclusion (14b) and requires the latter to be entailed by the former. I take the content of a sentence to be truth-

conditional in nature, i.e. to be the set of possible worlds in which the sentence is true, and entailment to be content inclusion, i.e. (14a) entails (14b) iff for any world *w*, if (14a) is true in *w*, so is (14b));

Second, we want to capture the meanings of (14a) and (14b). I take meaning to be context-change potential, i.e. to encode both content (truth-conditions) and anaphoric potential.

Thus, on the one hand, we are interested in the contents of (14a) and (14b). They are both modal quantifications: (14a) involves a circumstantial modal base (to use the terminology introduced in Kratzer 1981) and asserts that, in view of the circumstances, i.e. given that God created man in a particular way, as long as a man is alive, he must find some thing or other pleasurable; (14b) involves the same modal base and elaborates on the preceding modal quantification: in view of the circumstances, if a man is alive and has no spiritual pleasure, he must have a carnal pleasure. Note that we need to make the contents of (14a) and (14b) accessible in discourse so that the entailment particle *therefore* can relate them.

On the other hand, we are interested in the anaphoric potential of (14a) and (14b), i.e. in the anaphoric connections between them. These connections are explicitly represented in discourse (15) below, which is intuitively equivalent to (14) albeit more awkwardly phrased.

15. **a.** If a u_1 man is alive, he u_1 must find something u_2 pleasurable / he u_1 must have a u_2 pleasure.

    b. Therefore, if he u_1 doesn't have any u_3 spiritual pleasure, he u_1 must have a u_4 carnal pleasure.

Discourse (14/15) is analyzed in Intensional PCDRT (IP-CDRT) as a network of structured anaphoric connections and the meaning (and validity) of the Aquinas argument emerges as a consequence of the intertwined individual-level and modal anaphora.

In particular, note that the conditional in (15b) is modally subordinated to the antecedent of the conditional in (15a), i.e. (15b) is interpreted as: if *a man is alive and* he

doesn't have any spiritual pleasure, he must have a carnal pleasure. Modal subordination is analyzed as simultaneous modal and individual-level anaphora, i.e. the morpheme *if* in (15b) is anaphoric to the modal dref introduced by the antecedent of (15a) and the pronoun *he* u_1 is anaphoric to the individual-level dref introduced by the indefinite *a* u_1 *man* in the antecedent of (15a).

The IP-CDRT account of discourse (14/15) brings further support to the idea that the dynamic turn in natural language semantics does not require us to abandon the classical approach to meaning and reference. In fact, the analysis of (14/15) does the exact opposite: the introduction of propositional dref's in IP-CDRT enables us to recover the classical notion of truth-conditional content, which in turn enables us to analyze the Aquinas discourse in (14/15) as involving structured discourse reference to the propositional contents contributed by the premise and the conclusion of the argument.

Chapter 8. Conclusion

The last chapter contains a summary of the main results and briefly presents two future extensions of Intensional PCDRT, namely: *de se* attitudes (focusing on the Romanian subjunctive B mood) and plural anaphora and quantification.

The dissertation is located at the intersection of two major research programs in semantics that have gained substantial momentum in the last fifteen years: (**i**) the development of theories and formal systems that unify different semantic frameworks and (**ii**) the investigation of the semantic parallels between the individual, temporal and modal domains. As the dissertation shows, one of the outcomes of bringing together these two research programs is a novel compositional account of non-local (modal and individual-level) quantificational dependencies as *anaphora to structure*.

Thus, on the one hand, the present investigation takes the program in Muskens (1996) (see also Janssen 1986 and Groenendijk & Stokhof 1990 among others) of

unifying Montague semantics and dynamic semantics one step further[1]: Intensional PCDRT unifies – in classical type logic – the static Lewis (1973) / Kratzer (1981) analysis of modal quantification and van den Berg's Dynamic Plural Logic.

On the other hand, Intensional PCDRT enables us to explicitly and systematically capture the anaphoric and quantificational parallels between the individual and modal domains, in particular, between quantificational and modal subordination, thus bringing further support to the conjecture that our semantic competence is domain neutral, first put forth in Partee (1973, 1984) for the individual and temporal domains and extended to the modal domain by Stone (1997, 1999), Bittner (2001) and Schlenker (2005b) among others.

Summarizing, the dissertation can be seen as an extended investigation of a central issue raised by the dynamic turn in natural language semantics, namely: what kind of information is stored in an information state and how is this information updated in discourse? And, in particular: what can anaphora and quantification in both the nominal and the modal domain tell us about this?

The main result of the investigation is a new representation language, i.e. Intensional PCDRT, which is couched in classical type logic and in which natural language discourses involving complex descriptions of multiple related objects can be compositionally translated. Intensional PCDRT formalizes the idea that information states involve two equally important components with essentially the same interpretive dynamics, namely discourse information about values (sets of objects: individuals and possible worlds, but also times, eventualities etc.) and discourse information about structure (the correlations / quantificational dependencies between sets of objects that are introduced and elaborated upon in discourse).

---

[1] The research program of unifying Montague semantics and dynamic semantics goes back at least to Partee (1984): "I don't see how to incorporate Montague's elegant treatment of compositionality into the framework followed in this paper, nor do I see how to reproduce within Montague's theory the unified and explanatory account of nominal and temporal anaphora provided by these extensions of Kamp's, Heim's and Hinrichs' work. So the next task is to try to construct a theoretical framework which incorporates the insights of both approaches." (Partee 1984: 279).