Meaning and Use of Indefinite Expressions.

Paul Dekker
ILLC/Department of Philosophy
Faculty of Humanities
University of Amsterdam
dekker@hum.uva.nl

Abstract. Sentences containing pronouns and indefinite noun phrases can be said to express open propositions, propositions which display gaps to be filled. This paper addresses the question what is the linguistic content of these expressions, what information they can be said to provide to a hearer, and in what sense the information of a speaker can be said to support their utterance. We present and motivate first order notions of content, update and support. The three notions are each defined in a compositional fashion and brought together within a single and coherent framework.

Keywords: formal semantics and pragmatics, dynamic interpretation, discourse representation, speaker's reference, multi-speaker dialogue, cross speaker anaphora.

Introduction

Pronouns, and other indefinite terms, have intrigued philosophers for centuries, linguists for decades, and logicians for years. For what can we say about the linguistic content of an indicative sentence which contains unbound pronouns, or open positions? A pretty sophisticated reply to this question, of course, is this: "That depends!" It depends on whom or what the pronoun refers to. Since, obviously, the reference of a pronoun may vary in different contexts of use, we have to generalize the reply somewhat:

OBSERVATION 1 (Received Wisdom). *The content or 'character' of a sentence with open positions is a function $\mathcal{F}$, taking as its domain contexts determining possible values of the open positions, and for any such context c, $\mathcal{F}(c)$ is the content of the sentence if the open positions were filled in as specified by c.*

We take (1) to express the received wisdom. Observation (1) is firmly rooted in the tradition of model-theoretic semantics—a tradition motivated by the founding work of, among many others, Frege, Carnap and Kripke. In a formal setting it has been made explicit and motivated in (Montague, 1974; Janssen, 1986). However, as a semantic proposition, (1) does not say anything about the *use* of sentences with open positions. If someone asserts a sentence with an open position, does she communicate such a function, and, if not, then what? In more fashionable terms, what can we say about the intended update of information which the utterance of such a sentence brings about in the information state of a hearer? And what kind of evidence can a speaker be required to have to support such an utterance? Some partial answers to these questions have been given in the literature, but we claim these to lack the required kind of generality.

This paper focuses on the meaning and use of indefinite expressions. We will be mainly concerned with indefinite noun phrases and anaphoric pronouns, but demonstrative pronouns, definite descriptions and proper names might in principle be subsumed as well. Throughout, we build upon—but also provide motivation for—the assumption that such indefinite expressions ought to be used with referential intentions.[1] The notions of update and support, like that of linguistic content, thus become functionally dependent upon intended speaker's reference, and one of the main tasks of the paper lies in giving an exact formulation of precisely that dependence.

From the outset it may have to be said that we will not take into account information which agents may have (i) about the information other agents have about individuals and (ii) about their referential intentions. For this reason we do not formally account for pragmatic features of role switches in dialogue, in particular those having to do with cross speaker anaphora. Although examples of cross speaker anaphora motivate our thinking about first order update and support in section 1, they will be beyond the scope of the formalisms developed in subsequent sections. We do claim that our notions of update and support make up a proper part of an account of cross speaker anaphora, but a full account will require a further extension of the formal system presented here.

The paper then is organized as follows. In the first section of the paper we first present further motivation for our outlook upon the use of indefinite expressions. Then, in the second section, we give a precise definition of what we conceive to be key information theoretical concepts: that of first order information, first order information containment, and first order information merge. In the next two sections we couple the developed concepts to a first order language. In the third section we present and motivate a simple semantics for a first order language modeling (the use of) indefinite noun phrases and anaphoric pronouns. In section 4 we then show that the concepts developed in section 2

[1] Such an assumption appears to be maintained in a vast body of literature, c.f., e.g., (Chastain, 1975; Donnellan, 1978; Evans, 1982; Kamp, 1990), and, more recently (van Rooy, 1997; Stalnaker, 1998; Zimmermann, 1999).

indeed supply intuitively appropriate notions of content, update and support for the language of section 3. We conclude the paper with a summary of the results, and the onset of a study of further aspects of the use of pronouns (section 5).

Before we start, however, we want to set the stage, in the remainder of this introduction in order to synchronize our basic intuitions about the notions of 'content', 'update' and 'support'. We take our cue from a fragment of natural language which can be analyzed in terms of that of propositional logic, and sketch a most simple and perspicuous model of the exchange of information by means of that language.

Let us start with formulating some straightforward assumptions about the linguistic exchange of information:

- the utterance of indicative sentences basically serves the purpose of exchanging information
- information that gets exchanged is information about the world
- information is meant to be veridical[2]

(Although these assumptions need not be generally true, they may serve to narrow down the scope of the present paper to certain idealized situations.) For a language of propositional logic, these assumptions can be made more precise in the following fairly standard fashion.

Let us assume we have a language of proposition logic built up from a set of proposition letters by means of negation and conjunction. If we want the interpretation of such a language to reflect aspects of natural language indicatives, a well-motivated notion of linguistic content is one which specifies the truth-conditions of the sentences of this language, and a convenient way to do this is in terms of possible worlds or situations.[3] The interpretation of a sentence can be equated with the set of possible worlds or situations in which the sentence is true, and an assertion of the sentence can be said to express that the actual world or situation is in that set.

So let us also assume we have a (non-empty) universe of possible worlds W, and an initial valuation function V specifying the content of the proposition letters, for each letter the set of worlds in which it is true. Relative to such a model $M = \langle W, V\rangle$ the interpretation of the sentences of the language can be specified as follows:

[2] This last assumption we judge basic as well as problematic. Although we, of course, acknowledge the existence of (deliberate and inadvertent) misinformation—and the importance of it—we think information is primarily meant to characterize the world as being a certain way.

[3] Obviously, the notions of truth and truth-conditions go back to Frege; an elegant formulation of the idea is given by Wittgenstein "Einen Satz verstehen, heißt, wissen was der Fall ist, wenn er wahr ist." (To understand a proposition means to know what is the case, if it is true. Tractatus Logico-Philosophicus, Satz 4.024.)

DEFINITION 1 (Propositional Contents).

- $[\![p]\!]_M = V(p)$
  $[\![\phi \wedge \psi]\!]_M = [\![\phi]\!] \cap [\![\psi]\!]$
  $[\![\neg\phi]\!]_M = W \setminus [\![\phi]\!]_M$
- $\phi_1 \ldots \phi_n$ *entail* ψ *iff* $\forall M$: $[\![\phi_1]\!]_M \cap \ldots \cap [\![\phi_n]\!]_M \subseteq [\![\psi]\!]_M$

Let us now extend our perspective, and see how such a language can be used to model propositional information *exchange*. Like we said, we assume that agents exchange information they have about the world, and this type of information can be modeled in terms of sets of possible worlds as well. The information which an agent a has, also called a's information state, can be characterized by the set of those possible worlds which might be the actual one for as far as the agent knows. If the agent thinks cows give milk, then cows give milk in all the worlds in his information state; if she thinks horses don't drink beer, then all worlds in which horses do drink beer are excluded from her information state.[4]

So what can such an agent a learn from the utterance of a particular sentence ϕ, or what, in more fashionable terms, is the update which such an utterance brings about in his information state t? Obviously, the update of the information which he already had with the information which he gets involves intersection as well:

DEFINITION 2 (Propositional Updates).

- *the update of* t *with* ϕ *in* M, $(t)[\phi]_M$, *is* $t \cap [\![\phi]\!]_M$

If our agent has understood ϕ well, and if he is not deceived, he will think that the actual world is as he thought it was before the update, and also as it is said to be by the utterance of ϕ.

Let us now take a look at the other party in the game of information exchange, and see how the speaker's information relates, or must relate, to what she says to be the case. One of the major goals of the game of information exchange is that agents exchange true information, so one of the main rules is that they exchange information they believe to be correct. One might formulate this as follows. A speaker is licensed to say ϕ only if her information state s supports ϕ, where:

DEFINITION 3 (Propositional Support).

- s *supports* ϕ *in* M, $s \models_M \phi$, *iff* $s \subseteq [\![\phi]\!]_M$

[4] In this paper we generally assume speakers to be female and hearers to be male. This is not a matter of political correctness, but it simply enables us to use feminine pronouns to refer to speakers, and masculine ones to hearers.

If one says that the actual world is a ϕ-world, then one ought not to conceive it possible that the actual world is not a ϕ-world. So, obviously, support involves inclusion. If one were to consider it possible that penguins don't fly, then one is not licensed to say they do.[5]

The present notions of linguistic content, update and support are closely related and they behave well together. For upon the present definition it is easily seen that the following holds:

- the information which an agent has after the update with an assertion is contained in the joint information which he and the speaker had before the assertion, provided that the assertion was supported

Formally this can be specified as follows:

OBSERVATION 2 (Supported Updates).

- *if* $s \models_M \phi$ *then* $s \cap t \subseteq (t)[\phi]_M$

This fact guarantees that a proper exchange of propositional information does not corrupt the information exchanged. If the speaker's state supports her utterance, then her information together with that of the hearer supports the update which the hearer gets from accepting the utterance. Thus, if, in that case, the speaker and hearer have correct information, then the update of the hearer is guaranteed to be correct too. Surely, this is a desirable property of any model of information exchange.[6] However, although it is completely obvious that our simple propositional model has this property, it is not at all obvious that we preserve it when we turn to (first order) extensions of this model.

Another appealing feature of our present model is that the three core notions can be defined independently, in a compositional fashion. Clearly, the content of any compound expression is already stated compositionally in terms of the contents of its parts. The notions of update and support can be specified in a compositional way too. For instance, it is easily seen that:

[5] Such a notion of support thus figures as an implementation of Grice's maxim of quality: "Try to make your contribution one that is true," and "do not say that for which you lack adequate evidence". It can also be conceived of as a specification of Hamblin's notion of commitment: a speaker can be said to be committed to having the information supporting what she says.

[6] We take this to be a prime integrity principle on any theory of first order update and support, although the principle need not be tenable under all circumstances. Since utterances are acts, they may *change* the world. Thus, the very utterance of a sentence "I have never spoken to you." may be true at the moment of the assertion, but false immediately afterwards. See (Gerbrandy and Groeneveld, 1997; Gerbrandy, 1999) for related discussion. Such changes in the actual word are, however, beyond the scope of the present enterprise.

OBSERVATION 3 (Compositional Updates).

– $(t)[\neg\phi]_M = t \setminus (t)[\phi]_M$
  $(t)[\phi \wedge \psi]_M = ((t)[\phi]_M)[\psi]_M$[7]

OBSERVATION 4 (Compositional Support).

– $s \models_M \neg\phi$ *if for no* $\emptyset \subset s' \subseteq s$: $s' \models_M \phi$
  $s \models_M \phi \wedge \psi$ *iff* $s \models_M \phi$ *and* $s \models_M \psi$

This is appealing, because these observations shows that we can study the logic of update and support separately, and in a formally transparent fashion, without having to worry about the implications of this for the philosophically well-motivated notion of linguistic content. It is this pleasant interaction between the notions of content, update and support which we seek to preserve when we study the use of a language with (expressions modeling) indefinites and pronouns. Before that, however, we have to take a look at some motivating data and the complications they raise.

1. On Indefinite Reference

It need not come as a surprise that complications show up when we start looking at the exchange of propositions by means of a (first order) language which includes referring devices such as proper names, pronouns and demonstratives, definite and indefinite descriptions. (The term 'term' will be the general label by means of which we categorize this group of noun phrases.) There is a long tradition in the literature devoted to the study of the puzzles and paradoxes which arise when such terms figure in belief contexts.[8] It can be expected that similar complications arise when we study the information which people bring to bear upon their support for the utterances they make by means of these terms.

The present paper builds upon two firm, but not indisputable, convictions. In the first place, we believe that people generally do not have information at their disposal which can be modeled by means of so-called 'singular propositions', propositions involving unique individuals only. The idea is that at any time uncertainty may arise about the identity of known individuals, so that even if one knows that John is a nerd, one may meet him one day and fail to see that that individual

[8] To mention a few highlights: (Quine, 1956; Geach, 1967; Kaplan, 1969; Lewis, 1979). Also this issue dates back to (Frege, 1892).

is a nerd.[9] In the second place we also believe that agents nevertheless use natural language terms with referential intentions. The information they have and the information they exchange is intended to relate to real individuals in the actual world. When it comes to the support of open propositions referential intentions can be seen to be at work as well.

In this section we give some further motivation for the idea that the assertion of open propositions ought to be supported by referential intentions. We will critically review a couple of examples involving indefinite reference and co-reference, and we will argue that no tools from given semantic paradigms give us a handle on the data. Most of these examples involve dialogues and cases of cross speaker anaphora, since these most clearly bring to light the issues involved in thinking about the support for the use of indefinites and pronouns. After that, we develop our own notions of first order content, update and support which play a role here, although we will not fully account for the dialogues discussed here. For a full treatment requires us take into account the information the interacting agents have about the information of one another, and, like we said, this will require a (proper) extension of the framework developed in this paper.

In what follows we will mainly be concerned with the use of pronouns and indefinite descriptions because they have dominated the academic debate about information updates. However, we think that most of our observations apply to other terms, which, from an epistemic point of view, all can be called 'indefinitely referring expressions'.

1.1. Classical notions of support

The utterance of the following two sentences may serve to raise some of the main issues:

(1) A beatle just entered.
(2) He is going to sing a love song.

If the utterance context does not indicate otherwise, an utterance of (2) will be understood to be dependent upon that of (1), in the sense that the pronoun in (2) is understood to be coreferential with the indefinite noun phrase in (1). But what information can be said to support the assertion of the two sentences? We think that an answer to this question ought to be phrased in terms of an independent statement of the support of an assertion of (1) and of that of (2). For notice that the supporting information may come from different information states:

[9] Famous examples from the literature are the attitudes of Quine's Ralph concerning Ortcutt, or the opinions of Kripke's Pierre about London.

(3) *Zag:* A beatle just entered.
(4) *Nic:* He is going to sing a love song.

In this example, support for Zag's utterance must come from Zag's information state, and support for Nic's utterance from that of Nic's. This raises a couple of puzzling questions. For instance, what information can be said to support the (open) proposition expressed by (4)? Although it is evident, of course, that what Nic says in (4) is related to what Zag says in (3), and to what support Zag may have for saying so, it seems quite inappropriate to say that the evidence which Nic may bring to bear upon his claim itself depends upon what Zag has claimed.

It has been argued that a dialogue like (3–4) is odd or incoherent, but this seems to be correct only if it is placed in the wrong context. By itself, an exchange like that in (3–4) is perfectly OK. It is a typical *Beavis and Butthead*-dialogue, in which two interlocutors are commenting upon a scene they are both perceiving. As (Groenendijk et al., 1996) observes, an exchange in such a situation can be accompanied with exclamations like "Look!", and "Yeah!", but surely a felicitous exchange as in (3–4) can do without these exclamations. Rather, the challenging question is not whether (3–4) constitutes a coherent exchange, but in which contexts it does, and in which contexts it does not.

In classical systems of interpretation, there appear to be three ways to 'tame' an open proposition like that of Nic. The proposition can be closed, universally, or existentially, and the hole can be filled by means of an individual. (Actually, there is a fourth possibility. The hole can be filled by means of an individual concept. We will come back to this possibility in due course.) It will be clear that none of these three possibilities yields an adequate characterization of the evidence which Nic can be required to have to support his utterance. Surely it would be outrageuous too require Nic to have evidence that *everyone* is going to sing a love song (universal closure), but it is also a far too weak requirement that he has evidence that *someone* will be going to (existential closure). As we already indicated, equally unrealistic would be the requirement that he has this kind of information about a particular individual.

Imagine, for instance, the following situation. Weird rumours have it that young girls get battered in Gotham city, and that city representatives like, e.g., Bernard J. Ortcutt are involved. Bor and Cor exchange the latest news from the tabloids.

(5) *Bor:* A magistrate from Gotham city has confessed battering young girls.
(6) *Cor:* They say he suspected them of sorcery. Do you know if more magistrates confessed?

(7) *Bor:* I don't know.
(8) *Cor:* Do you know who he is?
(9) *Bor:* No idea, the police didn't disclose his identity.

We think there is nothing incoherent about a dialogue like this. It is just an exchange of information between two inhabitants of Gotham city, who—perhaps naively, but not incoherently—take their tabloids serious. This example shows that the two interlocutors can exchange bits of information i about an individual d without knowing that i applies to d. Maybe, later that evening, they will read in the evening papers that the magistrate was their neighbour, the honoured Ortcutt, without anyone of them having any suspicions about Ortcutt at the time of the exchange. If anything, like the information in the tabloids, licenses their utterances, it is not necessarily information uniquely involving the actual referent, information modeled by a singular proposition about Ortcutt. In short, supporting information should not be equated with information supporting singular propositions.

We already mentioned the possibility of feeding open propositions with individual concepts, functions determining (possibly different) individuals in the worlds conceived possible. However, it is not immediately clear how to apply this idea. Support for an open proposition should obviously not reside in a support for the proposition under all possible ways in which its hole can be filled by an individual concept, nor in the support of the proposition under some such way. So, it seems, a determinate concept should be in place, but which one?

Intuitively, in the exchange (3–4), we would like to require a concept like this to relate to Zag's previous use of an indefinite, so we might equate it with the concept given by a locution like "the individual which Zag referred to before." However, if we proceed along these lines, we should definitely require that such an individual concept does not become part of the supported proposition.[10] Besides, it is not obvious how this requirement could be built into our notions of content and support in a systematic fashion.

The account which we give of the support of open propositions below is very close in spirit to the one suggested last, except for the fact that it does not explicitly invokes individual concepts.[11] We employ a notion of an information aggregate in which information is hung onto discourse referents, or epistemic subjects, which can be taken to support

[10] For, otherwise, an utterance of "Tom believes he is going to sing a love song." could be taken to mean that Tom believes that someone is going to sing a love song and that Zag referred to him. This is quite a counterintuitive reading of the example.

[11] A sophisticated analysis which explicitly employs inidividual concepts in the setting of a dynamic semantics is offered in (Aloni, 1997).

the holes in open propositions. We do assume that such subjects are associated with individuating concepts, but these concepts, themselves, never enter the propositions which the subjects are assumed to support. Before we turn to a dynamic explication of how this is achieved, we, however, have to take a look at an alternative analysis of the support of open propositions, one which readily suggests itself, but which we eventually deem untenable.

1.2. Update Notions of Support

Systems of dynamic semantics have been developed to account, among other things, for the interpretation of (1–2) and, arguably, for the update of information which the exchange (3–4) may bring about in the information state of someone who overhears the dialogue. These systems adopt a notion of meaning richer than the classical one, and which allows for structural (anaphoric) relationships between the contents of the two sentences. Thus, the dynamic conjunction of (1) and (2) comes out to mean that a beatle who just entered is going to sing a love song, which seems adequate. Since a dynamic analysis successfully accounts for the interpretation of these sentences, it is tempting to adopt this dynamic perspective also when it comes to specifying their support. However, this turns out to be not unproblematic.

If a speaker's utterance of a sentence S' is anaphorically dependent upon a previous utterance of a sentence S, then the update notion of meaning suggests that the utterance of S' is licensed only if the speaker's update with S supports S'. This analysis does enable anaphoric connections between pronouns in S' and antecedents in S, but it turns out to be untenable, empirically speaking. Consider the following initial part of an exchange sequence:

(10) A member of parliament visited the queen yesterday.
(11) He was dead drunk.

Suppose that, as was suggested, the utterance of (11) in a sequence like this is licensed only if the update of the speaker's state with (10) supports (11). It then turns out that the speaker is only licensed to utter (11) if she has information that *all* members of parliament who visited the queen yesterday were dead drunk.[12] This does not appear to be appropriate.

[12] The standard dynamic semantic notion of support corresponds to a *strong* notion of implication, as to which the following two sentences are equivalent:

(12) If a man is in Athens, he is not in Rhodes. (from Heim)

(13) Every man who is in Athens is not in Rhodes.

On a straightforward application of dynamic semantic concepts to the analysis of the support of open propositions, support for an anaphoric proposition involves universal quantification over the possible values of its antecedent terms, and, clearly, this is too strong a requirement. For somebody may very well have information supporting the utterance of (10–11) if she has evidence that *some* member of parliament visited the queen yesterday and that *he* was dead drunk. This does not exclude that other members of parliament visited the queen yesterday who were not dead drunk, and it does not exclude that our reporter knows this.

Surely one may try out alternative notions of dynamic semantic support, such as the weak notion of support from (Beaver, 1994).[13] If we adopt the weak dynamic notion of support, support for an anaphoric proposition does not involve universal but existential quantification over the possible values of the antecedent terms. Using this notion, a speaker is said to be licensed to utter (11) in the context of (10) if he has information that some member of parliament who visited the queen yesterday was dead drunk. This seems to be much more appropriate, but, upon reflection, this weak notion of support does not seem to be tenable either.

Like we said, we think that when a speaker introduces a subject in the discourse by means of an indefinite noun phrase, she must have a particular subject in mind, and not just any arbitrary individual who happens to satisfy the properties attributed to that subject. Think of a situation where the newspapers all discuss some scandal related to a visit of a couple of MPs to the queen. Both Nel and Len have heard of it, Len happens to know that one MP who was there was dead drunk, but Nel starts out talking about another one, who was stoned as a shrimp:

(16) *Nel:* A member of parliament visited the queen yesterday.
(17) *Len:* He was dead drunk!
(18) *Nel:* No, he was not, not the one I meant. He was stoned!

The point of this example lies in Nel's correction of Len. We take it that Nel's correction can be supported by (i) the fact that she had started out talking about some MP with (16), and (ii) her information

---

[13] Beaver's notion corresponds to a *weak* notion of implication presented in (Chierchia, 1992), which renders the following two sentences equivalent:

(14) If I have a dime, I put it in the parking meter.

(15) If I have a dime, I put a dime I have in the parking meter.

that *that* MP was not dead drunk. Given this, Len's utterance has to be corrected.[14]

Facts like this are problematic if support is defined in terms of the weak dynamic notion of support. For in that case Len's utterance should be taken to assert that there was an MP with the queen yesterday who was dead drunk, and there would be no point in correcting this if Nel had no evidence for there being no such MP.[15]

In stead of discussing further, alternative dynamic notions of support, we can take examples like those in (16–18) to point at a much more fundamental problem with a dynamic semantic analysis of speaker's support. Key to the dynamic semantic analysis of anaphoric relationships lies in its capacity to integrate the meanings of utterances with indefinites with those of subsequent utterances involving pronouns. However, if we consider the support for utterances in a dialogue, it seems to be crucial to separate the contents of the uttered sentences, also if they are anaphorically related.

Consider the following exchange, reminiscent of an example due to Strawson:[16]

(21) *Toc*: A man is trying to push that oak down!

(22) *Aux*: It is a woman, and she is a jogger, stretching her hamstrings. Besides, it is a birch.

In this *Beavis and Butthead*-exchange, Aux disagrees with virtually all of what Toc says. The two only seem to agree upon the presence of an individual and a tree which they see and which they both seem to refer to. For this reason it is entirely inappropriate to state the support which Aux may bring to bear on his utterance in terms of the descriptions

---

[14] We do not deny that Len's utterance can be said to be licensed if Len had reason to suppose Nel's utterance related to the other MP, who was (known to be) drunk. Moreover, Nel may even be aware of that. Her correction, nevertheless, is independent of (her being aware of) that possibility.

[15] We submit that, like we said, Len's utterrance can be said to be licensed, if he had this other MP in mind, who was dead drunk. But even then the weak dynamic analysis of support still misses the crucial point that Len can be taken to address the individual which Nel is assumed to have had in mind, not just any arbitrary MP who satisfies a certain description which Nel used.

[16] Strawson's original example is:

(19) X: A man jumped of that bridge.

(20) Y: He didn't jump, he was pushed.

In our example (22), Aux uses a pronoun 'it' instead of 'he' as in Strawson's original. One might think this changes matters as this pronoun may have a different informational status. However, in our example the use of 'it' is simply due to the fact that what is introduced as a male is said to be female instead, so both the use of 'he' and of 'she' would be inappropriate.

which Toc has used. However, this is precisely what a dynamic semantic analysis of support seems to require. According to such an analysis, the support for Aux's reply should be stated in terms of his update with Toc's utterance, but, clearly, it is such an update which Aux sets out ot reject here! Put more generally, if the support for an open proposition is functionally dependent upon the antecedents of its pronouns, then the one who uses a pronoun must somehow agree with the descriptions associated with the antecedents. In general, this, simply, is not correct.

The preceding considerations apply with equal force to attempts to use an E-type analysis of pronouns along the lines of (Evans, 1977; Heim, 1990) in the analysis of the support of open propositions. Upon such an E-type analysis, a pronoun is short for a definite description constructed from the descriptive material associated with the pronoun's antecedent. When applied to exchange (21–22), this analysis would commit Aux to the truth of (23):

(23) A man who is trying to push that oak down is a woman jogger who is stretching her hamstrings against that birch.

Obviously, Toc does not need to have such blatantly contradictory information in order to substantiate her reply in (22).

1.3. Reference to Subjects

We have argued that the support of open propositions should not be characterized by means of a universal, existential or referential closure of the propositions involved. A dynamic semantic analysis in terms of updates which (strongly or weakly) support the open propositions has been argued to be insufficient as well. Instead, we favour an analysis upon which the support for the use of (anaphoric) pronouns relates to the speaker's conception of their referents, irrespective of the question whether the intended referents are demonstratively given, or presented by means of previous uses of (indefinite) terms. We will now discuss three more examples to substantiate this assumption in further detail.

Imagine the following situation. Liz was visited by two HiWis yesterday, who both inquired after the secretary's office. One of them (Wilburt) did, and the other one (Norbert) did not wear pink pumps. Liz is fully aware of all this. Now she starts the following exchange with Zil:

(24) *Liz:* Yesterday, a HiWi ran into my office who inquired after the secretary's office.

(25) *Zil:* Was he wearing pink pumps?

(26) *Liz:* He was indeed.

We deem this part of a dialogue perfectly felicitous if Liz had started out talking about Wilburt (otherwise, it is not of course). Notice that only upon the assumption that she had started to talk about either one of the two HiWis, a true and unconditional answer to Zil's question can be given. If it were undecided whom she was referring to when uttering (24), then there would be no such answer available.

In order to appreciate this point, consider the same situation, and suppose the conversation went as follows (this version of our scenario is due to Ede Zimmermann):

(24) *Liz:* Yesterday, a HiWi ran into my office who inquired after the secretary's office.
(25) *Zil:* Was he wearing pink pumps?
(27) *Liz:* ?I don't know. If it was Wilburt he was, if it was Norbert, he was not.

Apparently, this is a strange thing for Liz to reply. For if she had started out talking about Wilburt, then she could have replied affirmatively, and she could have replied negatively otherwise. With the present, marked, reply it seems as if she doesn't know whom she was talking about, which is odd, really. We claim it to be odd because we think Liz can be assumed to have had either Wilburt or Norbert in mind when she started out about a HiWi who ran into her office. Given that assumption, she could have given a decisive answer to Zil's question.[17]

It may be illuminating to point out that a dialogue like we find in (24–27) is not by itself incoherent. If we change the parameters a little bit the dialogue turns out to be fully felicitous again. Let us consider the same situation, but let us assume, in addition, that Liz missed her keys after Norbert and Wilburt had entered her office. Furthermore, assume that Liz is convinced that one of the two took the keys. She is just about to ask Zil if the one who took them left the keys with him when Zil interrupts her:

(24) *Liz:* Yesterday, a HiWi ran into my office who inquired after the secretary's office.
(25) *Zil:* Was he wearing pink pumps?
(27′) *Liz:* Ehmm, I don't know. If it was Wilburt he was, if it was Norbert, he was not.

This time, we deem Liz' reply OK. She started out reporting her loss of keys, and she doesn't know if the one who took them was wearing pink pumps. Like she said, if it was Wilburt he was, if it was Norbert, he was not. What is interesting in the last scenario is that Liz has *three* subjects to talk about: pink-pumpy Wilburt who ran in yesterday, non-

[17] If her first utterance of (24) could be deemed fine if she had not had anyone of the two in mind, then her actual reply would have been the most informative answer she could have given to Zil's question, and it ought to be fine, fine. But, like we said, it is odd.

pink-pumpy Norbert who also ran in yesterday, and the person who took her keys and who is either Wilburt or Norbert. The last subject is the one she had in mind in the last scenario, and relative to this subject her reply can be judged fully felicitous.

The preceding discussion can be used to substantiate two points which we suggested earlier. In the first place, it seems, a speaker must have definite individuals in mind when she talks about them, be it in definite or in indefinite terms. When Liz talks about a HiWi who ran into her office she is assumed to have a definite person in mind, be it Norbert, Wilburt, or the one who took her keys. Similarly, the examples (3–4) and (21–22) involve demonstratively present individuals, and in the exchanges in (5–9), and (16–18) the participants assume they are talking about that definite individual which they read or heard about, the individual which is deemed or assumed to be the (determinate) subject of the gossips.

In the second place, a speaker can be undecided about the true identity of this person. Liz may not know whether she refers to Norbert or Wilburt.[18] But also in this case it seems she is assumed to have a definite concept in mind of the individual she is referring to, for instance, that of the one who is supposed to have taken her keys. We think Liz' reply in the last dialogue would be pretty inappropriate if she could reasonably entertain the possibility that more than one individual had taken her keys.

1.4. (Co-)Referential Intentions

The upshot of the preceding discussion is that indefinite noun phrases, like pronouns, are generally used with referring intentions. A speaker ought to use these terms with specific subjects in mind, which she, in addition, is assumed to believe to be related to specific individuals in the actual world.[19] (Similar views upon the use of indefinites has been

[18] Likewise, in the exchanges (5–9) and (16–18) we find uncertainty about the true identity of the individuals at issue.

[19] From the speaker's point of view the support of indefinite expressions is very much like that of definite expressions. The choice to use an indefinite rather than a definite term may be given in by pragmatic considerations, for instance, because the identity of the actual referent is irrelevant, or because the addressee can be assumed not to be familiar with the referent. In either case, we assume or require there to be referential intentions at work, although this does not preclude possible uncertainty or ignorance of the speaker about who the actual referent is. Thus, a speaker may equally felicitously assert "Someone will be happy tonight," or "The winner of the race will be happy tonight," without actually knowing who in fact will win the steeplechase this afternoon.

expressed recently, in, e.g., (Aloni, 1997; van Rooy, 1997; Stalnaker, 1998; Zimmermann, 1999).)

The sketched assumptions may serve to explain why it makes sense to ask questions about the individuals which are introduced by means of indefinite noun phrases, and to make further claims about them. If a speaker has a specific subject s in mind when she uses an indefinite expression, and if she correctly supposes that a determinate individual d eventually relates to that subject s, then there is a true answer to any question about s. If the question is whether s has some property Φ, and if d has that property, then the answer is yes, otherwise it is not. Of course the agent may fail to know whether the individual has property Φ, or be miss-informed about it. The crucial point is that she assumes there to be a true answer to such a question.

Furthermore, if anybody takes the speaker serious, and believes that what she reports applies to a determinate individual, then he may set up his own subject s', which is supposed to relate to whatever individual it is that the original subject s relates to. Obviously this outlook upon the felicitous use of indefinites and pronouns fits well in the causal / intentional theory of reference (Kripke, 1972; Chastain, 1975; Donnellan, 1978; Evans, 1982), cf., also (Kamp, 1990; Zimmermann, 1999).[20]

The present outlook upon the use of indefinites expressions gives us an obvious handle on the support for their use. An agent a can be said to support the utterance of an open proposition $\phi(x)$—with x a pronoun or indefinite noun phrase—only if the utterance relates to a subject of a which is known to a as somebody who is ϕ.[21]

Before we present our notion of a subject of an information state, and that of the support of open propositions by means of subjects, it is useful to spend a few more words on the strength and limitations of the standard dynamic or update semantic paradigm.

As will appear from the discussion in the next section, systems of dynamic semantics have given rise to a notion of information fully adequate for the formulation of first order information, first order in-

[20] Notice that we predict the kind of uniqueness implications which have been associated with pronouns anaphoric upon indefinite antecedents in (Evans, 1977; Kadmon, 1987). The difference is that we think of these uniqueness implications as being pragmatic in nature, not semantic. Besides, we think they already show up with the use of the indefinite noun phrases which may serve as the antecedent of pronouns. For some more discussion, cf. (Dekker, 2000a, section 4).

[21] Surely more has to be required if x is a pronoun. In that case the speaker can be required to be justified in assuming that the intended referent is salient for her interlocutors. Some tentative remarks are made about this, complicated, issue in section 5.

formation update and first order information support. However, the prevailing systems of dynamic update semantics limit their scope to the study of first order updates, and although this apparently enhances the study of the logic of information update, it hampers the study of other aspects of use.

Like we said, standard notions of dynamic conjunction allow indefinites in one conjunct (the first one) to be structurally related with pronouns in the other (the second one). When it comes to the study of other types dependencies than those between indefinites and subsequent pronouns, this is an unwelcome limitation. For one thing, it is not just indefinite noun phrases that 'set up' discourse referents and subsequent pronouns which are anaphorically related to them. For another, pronouns do not always need explicit antecedents in order to be resolved, they can be used demonstratively as well.[22] A notion of merging first order information can be seen to be called for which is more general than the one pertaining to an update semantics.

Besides these, more or less cross-linguistic, motivations, the scope of update systems of dynamic semantics has to be extended when we start to study issues if information support. For if we set out to characterize the support for sentences with indefinites, it is simply not sufficient to be able to relate the contents of indefinites with those of subsequent pronouns. If we focus on a speaker's support for the use of indefinites, we have to relate their contents to what was there *before* their use, viz., to the subjects of the speaker supporting the use of the indefinites. Apparently, such 'backward' looking aspects of indefinites transcend the scope of rigid systems of update semantics.

[22] Besides, several (e.g., East-Asian) languages do not have any (definite or indefinite) articles at all. Even in Germanic languages there can be significant (anaphorical, rhetorical, ...) links between terms that miss articles, in newspaper headlines, for instance. The following two headlines, taken from a Dutch newspaper (the Volkskrant, June 19, 1999), constitute an interesting example:

(28) Natuurvolken moeten wijken voor de jacht op goud en hout
Primitive societies have to make way for the hunt for gold and wood

(29) Irians rijkdom is niet voor Papoea's
Irian's resources are not for the Papua's

A sensible interpretation of these two headlines, which accompany one article, associates the Papua's mentioned in (29) with the so-called 'primitive societies' of (28), and the gold and wood mentioned in (28) with the resources of Irian Jaya (29). It may be noticed that all four terms can be considered to be indefinite. Interestingly, the layout of the article does not determine a unique precedence relation between the two headlines. The second one was printed in a larger font, and thus may, but need not, be read before the first.

2. First Order Information Models

In the introduction we sketched a very simple model of the exchange of information by means of a language of propositional logic. One of the guiding assumptions, there, was that the information that gets exchanged is information about the world. When we turn to the exchange of information by means of a more complex first order language, which includes referring devices and quantification, things get more involved. In that case, it seems, the guiding assumption must be that the information that gets exchanged is information about the world, *and about the individuals (facts, events, ...) that inhabit that world.*[23]

In this section we study first order information structures. We present first order generalizations of the Boolean inclusion relation and the Boolean meet operation, which will enable us to characterize first order notions of update and support. The key concepts are that of an information aggregate, an entity containing information about the world and about certain numbers of individuals in that world, a relation of (information) containment on the domain of information aggregates, and a merging operation on sets of information aggregates.

When we have presented a satisfaction semantics for a first order language in section 3, we will show, in section 4, that the intuitive notions of content, update and support for that language are subsumed by the concepts developed in this section.

2.1. Information Aggregates

Like we said, very useful devices have already been developed in the literature for the representation of first order information. The 'discourse representation structures' from discourse representation theory (Kamp and Reyle, 1993), the 'domains and satisfaction sets' from file change semantics (Heim, 1982) and the 'information states' from dynamic semantics (Groenendijk et al., 1996), all enable us to model information about the world in conjunction with information about individuals in that world. The structures use variables ('discourse markers', 'pegs', ...) as labels for the individuals which the structures contain information about, together with constraints upon the possible values of these variables in the worlds conceived possible. By constraining the possible values of these variables, the subjects of the structures are associated with certain properties, for instance, by requiring the value

[23] When it comes to the issue of first order support, we thus aim at a first order extension of Grice's maxim: "Try to make your contribution one that is true *of the individuals (facts, events, ...) you are or intend to be speaking of*".

of one variable to be a woman, and requiring that of another to be her dog.

The kind of structures which we will use here are essentially those presented in (Heim, 1982) and studied further in (Dekker, 1996). In the present context we prefer to call these structures “information aggregates” because we do not only intend them to model the contents of the information states of epistemic agents, but also, more generally, those of texts and utterances and whatever else can be seen to contain (first order) information.

Already at this point one might wonder what the accommodation of subjects in information aggregates really adds to the contents of ordinary information states, which are sets of plain worlds. For the subjects of aggregates are present as individuals in the worlds conceived possible, and these also satisfy the properties attributed to the subjects. Technically, the difference is this. Specific subjects in aggregates, and not the individuals in the corresponding worlds, are accessible and readdressable. Thus, pronouns can be related to a subject in an aggregate representing the content of a sentence, and, likewise, such a subject can be linked to a subject in the aggregate of a speaker under the support relation. Conceptually, subjects in aggregates, and not the individuals in corresponding worlds, flesh out the assumption of having referential intentions. The presence of a subject in an information aggregate of an agent indicates that the agent assumes there to be a definite individual which has the properties attributed to the subject. Which individual that is may depend on a causal / intentional chain, which escapes the epistemic reach of the agent. And, as may well be, such a chain may just as well be broken or even absent. See (Dekker, 2000a) for some more discussion.

Information aggregates consist of a domain of variables and a set of variable assignments, which specify which valuations of these variables are conceived possible in that aggregate, and which ones are ruled out. In formal terms, aggregates are defined relative to a set of possible worlds W, a domain of individuals D and a set of variables V.[24] Variables are used to represent individuals in the epistemic states of epistemic agents or to label the discourse referents introduced in a discourse. Although it is not strictly necessary, we have chosen to use one distinguished variable v to represent the world.

[24] We now limit ourselves to information about individuals, although nothing hampers the inclusion of information about facts, events, etc. For the sake of simplicity, we also assume that all worlds have the same domain of individuals. Nothing really hinges upon these assumptions though.

For a finite set of individual variables $X \subseteq V$, with $v \notin X$, the set of possible variable assignments is $\{f \cup g \mid f \in W^{\{v\}} \ \& \ g \in D^X\}$, which we abbreviate as $D_W^{X_v}$, the space for X. The assignments g in the space for X each specify a world $g(v)$ and an individual $g(x)$ for any variable x in X. Information aggregates consists of a domain X and a subset of the space for X:

DEFINITION 4 (Information Aggregates). *An information aggregate is a pair $\langle X, \sigma \rangle$ with $v \notin X \subseteq V$ and $\sigma \subseteq D_W^{X_v}$*

Intuitively, in an aggregate $\langle X, \sigma \rangle$ the variables in X figure as names or labels for its subjects, and σ constrains the values of these variables, ruling out certain valuations if (and only if) they are not in σ. Thus, an aggregate $\langle X, \sigma \rangle$ can be said to contain information about the values of the variables in X. In what follows we will use Σ^X to refer to the set of information aggregates about the values of X. Furthermore, we will systematically use $\langle X, \sigma \rangle$ and σ interchangeably. This allows us to say that some assignment g is an element of an aggregate σ if, really, g is an element of the second projection of the pair $\langle X, \sigma \rangle$.[25] Finally, for each aggregate $\sigma \in \Sigma^X$ we will use $d(\sigma)$ to refer to the domain of σ, which is X. More in general we use $d(f)$ to indicate the domain of any function or relation f.

The notion of a subject of an information aggregate coincides with the pattern characterizing what possible values the associated variable x may have, given the possible valuations of the other variables.[26] By the constraints upon the possible valuations of the variables in the domain of an aggregate, the aggregate can be said to contain information about the corresponding subjects. For instance, an aggregate σ has the information that a subject has some property P if (and only if) all possible valuations of the corresponding variable yield an individual who is P. More precisely, if such a subject is labelled x, then σ dresses it with the property P iff each assignment $g \in \sigma$ assigns an individual d to x and a world w to v in which d has the property P. Thus, the aggregate determines that whatever the world is like, the value of x will always be something which is P. Similarly, an information aggregate σ may

[25] Notice that we can reconstruct X from σ as long as $\sigma \neq \emptyset$. We have made the domain information explicit in our definition of an information aggregate because this enables us to distinguish two empty aggregates σ and τ which do not belong to the same domain.

[26] Typically, such a subject is a partial object, in that it may have different possible values, and it is essentially related to other objects, in that the values of one variable may depend on those of others. See (Dekker, 1996, § 2.3) for a study of this formal notion of a subject.

contain the information that subjects x and z stand in some relation R, viz., if each $g \in \sigma$ assigns a pair $\langle g(x), g(z)\rangle$ to x and z which stand in the relation R in the world $g(v)$ according to g.

There is a close correspondence between the notion of an information aggregate and that of a discourse representation structure (*DRS*). Information aggregates can often be very neatly and perspicuously represented by means of these *DRS*s. Thus, a *DRS* K of the form:

$$(30)\quad K = \begin{array}{|c|}\hline x_1, \ldots, x_i \\ \hline \phi_1 \\ \vdots \\ \phi_n \\ \hline \end{array}$$

represents the information contained in the aggregate:

(31) $\langle X, \{g \in D_W^{X_v} \mid\ g(v), g \models \phi_1\ \&\ \ldots\ \&\ g(v), g \models \phi_n\}\rangle$

where $X = \{x_1, \ldots, x_i\}$ and $\models$ indicates the usual satisfaction relation between a formula ϕ and a model $g(v)$ relative to an assignment g. Conversely, information aggregates can be taken to model the intuitive contents of *DRS*s.

Before we carry on, it is worthwhile to point out a particular class of aggregates which we will be concerned with later in this paper. These are the aggregates with a linearly ordered domain of variables. If the domain $\{x_1, \ldots, x_i\}$ of an aggregate is linearly ordered, then we can represent this domain by the natural number: $i = \{0, 1, \ldots, i{-}1\}$.[27] Such an aggregate can then be neatly specified as a set of sequences we consisting of a world w and a sequence of i individuals e, viz., as the set $\{g(w)g(0)\ldots g(i{-}1) \in W \times D^i \mid g \in \sigma\}$. In the remainder of this section we will be concerned with aggregates of the most general type, that is with those with any domain of variables. However, in section 4, aggregates of this more particular type turn out to be very useful, technically speaking.

2.2. Information Containment

The reader may realize that the specific variables which an aggregate uses to label its subjects is of no substantial importance. If the information in an aggregate is centered around i subjects, they may be labeled by $x_1, \ldots, x_i$ or $y_1, \ldots, y_i$ or whatever set of i variables. The various aggregates count as equivalent, for the same reason that two

[27] We here adopt the von Neumann representation of the naturals according to which $0 = \emptyset$ and $n + 1 = n \cup \{n\}$, so that $n = \{0, 1, \ldots, n{-}1\}$.

discourse representation structures or two first order formulas count as equivalent if they are alphabetical variants of each other. However, if we want to relate two information aggregates, we do have to take their domains into consideration.

Consider, for instance, (the information aggregate pictured by) representation K, associated with a sentence "A man sees a woman":

$$(32)\ \ K = \begin{array}{|c|}\hline x, z \\ \hline man(x) \\ woman(z) \\ x\ see\ z \\ \hline \end{array}$$

One may wonder whether K contains the information in (the aggregate pictured by) representation K':

$$(33)\ \ K' = \begin{array}{|c|}\hline u, w \\ \hline u\ see\ w \\ \hline \end{array}$$

Clearly, the answer should be: "That depends!" Representation K contains K' if we relate u with x and w with z. Such a *resolution* would be appropriate if K' was associated with a sentence "He sees her," which appears to be a logical consequence of "A man sees a woman." Alternatively, K does not contain K' if u is related to z and w to x. Such a resolution would be appropriate if K' derived from "She sees him," which indeed does not follow from the original sentence.

The upshot of this, obviously artificial, example, is that certain substantial relations between aggregates ought to be specified relative to links between the subjects of these aggregates. Thus, a suitable notion of information containment must be defined relative to a linking function which relates the subjects of the contained state to that of the containing one. In this way we can say that an aggregate σ contains an aggregate τ with respect to such a linking function l iff σ has at least as much information about the subjects in the range of l as τ has about the corresponding subjects in l's domain. Let us use $g \bar{\circ} l$ to indicate the assignment $h \in D_W^{d(l)_v}$ such that $h(v) = g(v)$ and $\forall x \in d(l)$: $h(x) = g(l(x))$. The containment relation then can be defined as follows:

DEFINITION 5 (Information Containment).

- *σ contains τ wrt l, $\sigma \sqsubseteq_l \tau$, iff $l \in d(\sigma)^{d(\tau)}$ and $\forall g \in \sigma$: $g \bar{\circ} l \in \tau$*

An aggregate σ contains τ wrt l iff l is a function from the domain of τ to that of σ and every possible assignment in σ comes from one in τ via l. Thus, we make sure that if τ excludes a certain value of a variable x, then σ excludes it as a possible value of $l(x)$ as well. In other words, if an aggregate σ contains τ under l, then for all subjects $y_1, \ldots, y_j$ in τ there are subjects $l(y_1), \ldots, l(y_j)$ in σ such that σ has more information about $l(y_1), \ldots, l(y_j)$ then τ has about $y_1, \ldots, y_j$. In that case it is appropriate to say that σ contains at least as much information about at least as many subjects as τ.

By way of illustration, let $d(\tau) = \{x, z\}$, $d(\sigma) = \{u, w\}$, and let $l(x) = u$ and $l(y) = w$. Then:

(34) $\sigma \sqsubseteq_l \tau$ iff σ has more information about the values of u and w than τ has about those of x and z, that is, iff

– $\{\langle g(v), g(x), g(z)\rangle \mid g \in \sigma\} \subseteq \{\langle h(v), h(u), h(w)\rangle \mid h \in \tau\}$

So if x and z stand in some relation R according to τ, then $l(x)$ and $l(z)$ stand in that relation according to σ as well.

The information containment relation can be reduced to the subset relation, under an appropriate translation of the contained aggregate. Such a translation transforms the domain of that aggregate into that of the intended aggregate, thereby respecting the linking function l:

DEFINITION 6 (Translating Aggregates).

– $[l]_X\tau = \{g \in D_W^{(r(l)\cup X)_v} \mid g \mathbin{\bar{\circ}} l \in \tau\}$

(Here, and elsewhere, $r(f)$, indicates the range of any function or relation f.) Relative to a link l and a set of variables X, $[l]_X\tau$ is an aggregate with domain $r(l) \cup X$, in which the information of τ about its subjects $y_1, \ldots, y_j$ in τ is hung upon the corresponding subjects $l(y_1), \ldots, l(y_j)$. About the values of the variables in X which are not in the range of l, the translated aggregate has no information at all.[28] An aggregate now can be seen to contain another one if it is a subset of a translation of the other one:

OBSERVATION 5 (Containment as Subset).

– $\sigma \sqsubseteq_l \tau$ *iff* $\sigma \subseteq [l]_{d(\sigma)}\tau$

Using observation (5) it is relatively easily seen that $\sqsubseteq$ is reflexive under the identity relation, and transitive modulo composition of links:

OBSERVATION 6 (Reflexivity and Transitivity).

– $\sigma \sqsubseteq_i \sigma$ *if* $\forall x \in d(\sigma)$: $i(x) = x$;
– *if* $\sigma \sqsubseteq_l \tau$ *and* $\tau \sqsubseteq_{l'} \upsilon$, *then* $\sigma \sqsubseteq_{l\circ l'} \upsilon$

[28] Notice that $[l]_X\tau$ may contain strictly more information than τ, since l may involve the identification of different subjects of τ with one another.

(Whenever we use i to refer to a linking function in the sequel it will always indicate such an identity function.) The relation $\sqsubseteq$ does not partially order the space of information aggregates because, as one might expect, the containment relation is not antisymmetric. Two aggregates can be substantially equivalent without being identical, simply because they use different variables for corresponding subjects. However, since the relation is transitive and reflexive, it does partially order the structure consisting of equivalence classes of aggregates which mutually contain each other.

2.3. Information Products

In the simple model of information exchange which we presented in the introduction, the merge of two information states amounts to the intersection of sets of possible worlds. It may be clear that the merge of any two information aggregates must be somewhat more involved.[29] For different aggregates may hang their information about individuals on different sets of variables and if we merge the aggregates we may want to equate or *resolve* subjects named by different variables. So, also if we speak of the product of information of a set Π of information aggregates, we refer to linking functions indicating which variables in the aggregates in Π correspond to which variables in the product.

The product $\prod_{\varpi} \Pi$ of such a set of information aggregates is defined relative to a dependent linking function ϖ. This function ϖ assigns to each aggregate π in Π a linking function $\varpi(\pi) \in V^{d(\pi)}$ from the domain of π to the set of variables V. The product of Π under such a function ϖ is an aggregate with domain $r(\varpi) = \bigcup_{\pi \in \Pi} r(\varpi(\pi)) = \{\varpi(\pi)(x) \mid \pi \in \Pi \ \& \ x \in d(\pi)\}$. It is defined as follows:

DEFINITION 7 (Information Product).

– $\prod_{\varpi} \Pi = \bigcap_{\pi \in \Pi} \{g \in D_W^{r(\varpi)_v} \mid g \mathbin{\bar{\circ}} \varpi(\pi) \in \pi\}$

[29] The following discussion can be seen as a response to (Kamp and Reyle, 1993, p. 74*ff*), where it is argued that the contents of related sentences ought to be captured by single representations. Consider their own example of a sequence of two sentences:

(1.28) Jones owns a Porsche. It fascinates him.

Kamp and Reyle note (p. 78) "(...) *there is* some Porsche to which both the predication of the first sentence and the predication of the second sentence of (1.28) apply (...). Such an interpretation can only be captured by a single representation that captures the joint content of the two sentences together, not by a pair of unconnected representations, one for each of the sentences of its own." It seems Kamp and Reyle here overlook the possibility of having connected representations, which is what we are going to use.

Relative to a dependent linking function ϖ, the information product of a set Π of aggregates collects all the information of each one of the aggregates π in Π, by hanging the information π has about the variables in $d(\pi)$ onto corresponding variables given by the link $\varpi(\pi)$. Thus, if any such π excludes a certain valuation of the variables $x_1, \ldots, x_j$ in $d(\pi)$, then $\prod_\varpi \Pi$ excludes an analogous valuation of the corresponding variables $\varpi(\pi)(x_1), \ldots, \varpi(\pi)(x_j)$.[30]

Let us consider an example. Assume π_1 contains the information that x is a bird and π_2 that z is a nerd:

(36) $\pi_1 = \{f \in D_W^{\{x\}_v} \mid f(x) \text{ a bird in } f(v)\}$

(37) $\pi_2 = \{h \in D_W^{\{z\}_v} \mid h(z) \text{ a nerd in } h(v)\}$

Furthermore, suppose that ϖ is a linking function such that $\varpi(\pi_1)(x) = u$ and $\varpi(\pi_2)(z) = w$. Then the information product of the two under ϖ contains the information that u is a bird and that w is a nerd:

(38) $\prod_\varpi \{\pi_1, \pi_2\} = \{g \in D_W^{\{u,w\}_v} \mid g(u) \text{ a bird \& } g(w) \text{ a nerd in } g(v)\}$

Of course, the two aggregates can also be merged under a linking function under which the two subjects are identified. If $\varpi'(\pi_1)(x) = \varpi'(\pi_2)(z) = y$, then their product under ϖ' contains the information that y is a bird and a nerd:

(39) $\prod_{\varpi'} \{\pi_1, \pi_2\} = \{g \in D_W^{\{y\}_v} \mid g(y) \text{ a bird and a nerd in } g(v)\}$

Actually, this type of merging information is the semantic correlate of the conjunction of a sentence containing a pronoun with a sentence containing an antecedent term. We can think of π_1 and π_2 as conveying the information of the examples (40) and (41), respectively:

(40) There is a bird.

(41) It is a nerd.

In such a case $\prod_{\varpi'} \{\pi_1, \pi_2\}$ can be seen to convey the information contained in the conjunction of the two examples, under a resolution of the pronoun "it" with the noun phrase "a bird".

The intimate relation between our notion of information containment and that of an information product can be judged from the following observation:

[30] It can be observed that the information product of a set of aggregates Π under ϖ equals the intersection of the translations of the aggregates $\pi \in \Pi$ under $\varpi(\pi)$. For:

(35) $\prod_\varpi \Pi = \bigcap_{\pi \in \Pi} ([\varpi(\pi)]_{r(\varpi)} \pi)$

Thus understood, it may be clear that the product operation can be represented by an obvious operation on the representational level. If $\mathcal{K}$ is the set of corresponding *DRS*s, such an operation $\prod_\varpi \mathcal{K}$ substitutes any variable x of a *DRS* $K \in \mathcal{K}$ by the variable $\varpi(K)(x)$ and then combines the domain and the conditions of the resulting *DRS*s.

OBSERVATION 7 (Supported Products).

– $\sigma \sqsubseteq_l \prod_\varpi \Pi$ *iff* $\forall\pi \in \Pi$: $\sigma \sqsubseteq_{l\circ\varpi(\pi)} \pi$

Clearly this is as it should be. One chunk of information contains the sum of a whole set of chunks if and only if it contains each one of them. An immediate consequence of the observations (6) and (7) is that:

OBSERVATION 8 (Lower Bounds).

– $\forall\pi \in \Pi$: $\prod_\varpi \Pi \sqsubseteq_{\varpi(\pi)} \pi$

We see that, on the one hand, the information product of a set of aggregates contains the information of each aggregate in that set, and, on the other hand, if anything contains the information of each aggregate in that set, then it contains the information in the product. Products, in other words, constitute the infimum of a set of aggregates under dependent linking functions.

2.4. 'Dynamic' Merges

The notion of an information product is a very general one, which allows renaming of all the subjects of the combined aggregates. However, in most cases in which we take the product of two aggregates we may employ more specific products. A case in point is, for instance, the operation known as the 'Zeevat Merge'. The Zeevat Merge combines the domains of two *DRS*s K and K' and their conditions without any further translation. At the informational level this merge corresponds to the product of the corresponding aggregates π_1 and π_2 under a dependent identity function: $\prod_\varpi\{\pi_1, \pi_2\}$ where $\varpi(\pi_1)$ and $\varpi(\pi_2)$ are the identity functions on the domains of π_1 and π_2, respectively.

The Zeevat merge is a symmetric operation which does not assign any special status to one of the two aggregates or representations which it combines. Of course, we can also take the product of two information aggregates in a more dynamic fashion. One aggregate is taken as basic then, and the information of the other is injected into the basic aggregate under some linking function. Let us use $\sqcap$ to indicate such a 'dynamic' merge. If $l \in V^{d(\pi_2)}$ is any function from the domain of π_2, and $i_{d(\pi_1)}$ the identity function over $d(\pi_1)$, then:

DEFINITION 8 (Dynamic Merge).

– $\pi_1 \sqcap_l \pi_2 = \prod_\varpi\{\pi_1, \pi_2\}$, *with* $\varpi(\pi_1) = i_{d(\pi_1)}$ *and* $\varpi(\pi_2) = l$

If we think in terms of translations, it can be seen that:

OBSERVATION 9 (Merge and Intersection).

– $\pi_1 \sqcap_l \pi_2 = [i_{d(\pi_1)}]_{r(l)}\pi_1 \cap [l]_{d(\pi_1)}\pi_2$

A dynamic merge of two aggregates can be seen to capture the dynamic aspects of interpretation implemented in *DRT*. We can think of the initial *DRS* or information aggregate as the context of interpretation, possibly built up on the basis of preceding discourse. This *DRS* or aggregate is then updated with the information conveyed by a new sentence. The variables in the (initial representation for, or interpretation of) this new sentence can be matched with familiar ones, or with new ones. The result is an updated representation or aggregate, possibly containing new subjects.

The dynamic way of merging information can also be looked upon from a different angle. Instead of taking the initial context as basic, we can take a sentence's interpretation at the moment of utterance as basic. Dynamic conjunction then can be seen to involve the translation and addition of elements from the context required to resolve the free variables of the sentence.

Since our interest lies with the exchange (update and support) of first order information, it is expedient to study the underlying notions of containment and merge somewhat further. As might have been expected, the two behave well relative to one another:

OBSERVATION 10 (Support and Merge).

– $\sigma \sqsubseteq_l \tau$ *iff* $\sigma = \sigma \sqcap_l \tau$

Proof. First notice that if $r(l) \not\subseteq d(\sigma)$, then none of the two clauses can be true, and that if $r(l) \subseteq d(\sigma)$, then, and only then, $\sigma \sqcap_l \tau = \sigma \cap [l]_{d(\sigma)}\tau$ (observation 9). So, $\sigma = \sigma \sqcap_l \tau$ iff $\sigma = \sigma \cap [l]_{d(\sigma)}\tau$ iff $\sigma \subseteq [l]_{d(\sigma)}\tau$ iff $\sigma \sqsubseteq_l \tau$ (observation 5). End of proof.

An aggregate σ contains an aggregate τ under l if and only if its dynamic merge with τ under l doesn't add anything of substance. This is a welcome result, reminiscent of the Boolean relation between $\subseteq$ and $\cap$.

Observation (7) can also be given a further specification when it concerns dynamic merges:

OBSERVATION 11 (Supported Merge).

– $\sigma \sqsubseteq_m (\pi_1 \sqcap_n \pi_2)$ *iff* $\sigma \sqsubseteq_{m \restriction d(\pi_1)} \pi_1$ *and* $\sigma \sqsubseteq_{m \circ n} \pi_2$

An information state σ contains the merge of two aggregates π_1 and π_2 if and only if it contains both aggregates, all of this relative to appropriate linking functions, of course. By way of illustration, consider again the aggregates π_1 and π_2 which we discussed above, and their product under the linking function ϖ' which identifies the subject x of π_1 with the subject z of π_2. An aggregate σ can be seen to contain this product if it contains π_1 and π_2 under links l' and l'' which identify x with z. In this case we can again conceive of π_1 and π_2 as the interpretation of two sentences, and of σ as the information state of a speaker supporting the two sentences.

The notions of support and merge can also be related in a more subtle way, useful for the characterization of a first order notion of safe information exchange later in this paper:

OBSERVATION 12 (Distribution).

– *if* $\sigma \sqsubseteq_l \pi$, *then* $(\sigma \sqcap_m \tau) \sqsubseteq_{l'} (\pi \sqcap_n \tau)$, *provided that*

1. $d(m) = d(n) = d(\tau)$ *and* $d(l') = d(\pi) \cup r(n)$

2. $l \subseteq l'$ *and* $m = l' \circ n$

Proof. By the observations (5) and (9), we have $\sigma \sqsubseteq_l \pi$ iff $\sigma \subseteq [l]_{d(\sigma)}\pi$, and $(\sigma \sqcap_m \tau) \sqsubseteq_{l'} (\pi \sqcap_n \tau)$ iff $([i]_{r(m)}\sigma \cap [m]_{d(\sigma)}\tau) \subseteq [l']_{(d(\sigma)\cup r(m))}([i]_{r(n)}\pi \cap [n]_{d(\pi)}\tau)$. The latter inclusion holds iff:

(a) $([i]_{r(m)}\sigma \cap [m]_{d(\sigma)}\tau) \subseteq [l']_{(d(\sigma)\cup r(m))}[i]_{r(n)}\pi$ and
(b) $([i]_{r(m)}\sigma \cap [m]_{d(\sigma)}\tau) \subseteq [l']_{(d(\sigma)\cup r(m))}[n]_{d(\pi)}\tau$

Now if $\sigma \subseteq [l]_{d(\sigma)}\pi$, and $l \subseteq l'$, it can be seen that (a) holds, since then $[i]_{r(m)}\sigma \subseteq [l']_{(d(\sigma)\cup r(m))}[i]_{r(n)}\pi$.[31] Likewise, if $\sigma \subseteq [l]_{d(\sigma)}\pi$ and $m = l' \circ n$, then $[m]_{d(\sigma)}\tau \subseteq [l']_{(d(\sigma)\cup r(m))}[n]_{d(\pi)}\tau$, so (b) holds as well.[32] So if $\sigma \sqsubseteq_l \pi$, then $\sigma \subseteq [l]_{d(\sigma)}\pi$, whence, under the stated proviso's, (a) and (b) hold, and, hence, $(\sigma \sqcap_m \tau) \sqsubseteq_{l'} (\pi \sqcap_n \tau)$. End of proof.

This result can be appreciated in the following way. The aggregate π can be conceived of as the interpretation of an utterance supported by an information aggregate σ, and the merge of π with τ as the information state resulting from the update of an information state τ with π. Observation (12) then tells us that the resulting state is contained in the joint information which was already there in σ and τ. Notice that in the 'update' from τ to $\pi \sqcap_n \tau$ resolutions may be made by means of the linking function n. These, however, must correspond to resolutions in the product $\sigma \sqcap_m \tau$. If the resolutions by m are justified, then the update is safe. As we will see in section 4, observation (12) is key to the first order generalization of the result in example (2) concerning safe information exchange.

[31] Suppose $g \in [i]_{r(m)}\sigma$, i.e., $d(g) = d(\sigma) \cup r(m)$ and $g \,\bar{\circ}\, i_{d(\sigma)} \in \sigma$. Since $\sigma \subseteq [l]_{d(\sigma)}\pi$, we find that $g \,\bar{\circ}\, i_{d(\sigma)} \,\bar{\circ}\, l \in \pi$. Since $d(l) = d(\pi)$, $l \subseteq l'$, and $r(l) \subseteq d(\sigma)$, this is equivalent with $g \,\bar{\circ}\, l' \,\bar{\circ}\, i_{d(\pi)} \in \pi$. In conjunction with the observation that $d(g) = d(\sigma) \cup r(m)$ iff $d(g) = r(l') \cup d(\sigma) \cup r(m)$ this serves to show that $g \in [l']_{(d(\sigma)\cup r(m))}[i]_{r(n)}\pi$.

[32] By a line of reasoning similar to that in the previous footnote.

Although it is not pertinent to the purposes of this paper, we like to point out that the space of information aggregates also accommodates a notion of a common ground, the dual of our notion of an information product. Given this, it can be seen that the information space has the structure of what may be called a relativized lattice, with joins and meets defined relative to dependent linking functions.

3. Predicate Logic with Anaphora

In the preceding section we have developed and studied a general notion of first order information, which can be put to use when we study the meaning and use of a fragment of natural language with the expressive power of a first order logic. Information aggregates can be used to model the contents of sentences of that language, as well as the information which linguistic agents have about (individuals in) their world. Likewise, the notions of information containment and (dynamic) merge can be used to characterize suitable (linguistic) notions of first order update and support. All of this is, however, not done before the fourth section of this paper.

First, in this section, we will focus in on the dynamics of indefinites and anaphoric coreference in a simple extensional setting. We will present and motivate a system of interpretation here for a language of first order predicate logic with anaphora. We thus hope to show that the relevant aspects of dynamic interpretation can be motivated independent from the technical machinery we have developed in the preceding section. The semantics of the system in this section is stated as a Tarskian satisfaction relation between sequences of individuals and sentences. These sequences must be understood to provide the intended referents of the terms in the sentences (indefinite noun phrases or existentially quantified variables and anaphoric pronouns).

Like in most of the existing dynamic interpretations of predicate logic, indefinite terms have a progressive dynamic potential in the sense that their intended referents are accessible for subsequent anaphoric pronouns. In the system presented below, however, they are also visible from the past, so to speak. That is to say, they can be related to the subjects of the agents who have used them. Precisely this feature enables us to characterize, in a uniform and systematic manner, important features of the use of indefinites and pronouns from both an update and a support perspective.

In what follows we make some simplifying assumptions. The only terms which we deal with are indefinite noun phrases (modeled by means of existential quantifiers) and pronouns (introduced as a separate category of terms $\mathtt{p}_1, \mathtt{p}_2, \ldots$). We adopt the following (rather unrealistic, but commonly made) assumptions, too:

1. Only indefinites contribute (discourse) referents;
2. Pronouns can only be resolved by previously used indefinites;
3. Pronouns carry a disambiguating index.

More in particular, indefinites are assumed to contribute their intended referents ('discourse referents') one after the other; an index i on a pronoun $\mathtt{p}_i$ is used to indicate that it is coreferential with the i-th indefinite which can be found by going back in the discourse from the place where the pronoun occurs.[33]

3.1. Syntax and Terms

The system of interpretation which we present in this section is referred to as *PLA*, *(Static) Predicate Logic with Anaphora*.[34] The language of *PLA* is like that of first order predicate logic except for the fact that it also contains a category of pronouns $P = \{\mathtt{p}_1, \mathtt{p}_2, \ldots\}$. In order to keep things as simple as possible, we do not consider individual constants and identity, but this is not in any way essential.

Given a set U of variables and sets of m-ary relational constants $\mathcal{R}^m$, *PLA*-formulas are constructed as follows:

DEFINITION 9 (*PLA* Syntax).

- *the set of terms T is $P \cup U$*
- *the set of well-formed formulas is the smallest set $\mathcal{L}$ such that:*

  if $R \in \mathcal{R}^m$, and $t_1, \ldots, t_m \in T$, then $Rt_1 \ldots t_m \in \mathcal{L}$

  if $\phi \in \mathcal{L}$, then $\neg(\phi) \in \mathcal{L}$

  if $\phi \in \mathcal{L}$ and $x \in V$, then $\exists x(\phi) \in \mathcal{L}$

  if $\phi, \psi \in \mathcal{L}$, then $\phi \wedge \psi \in \mathcal{L}$

the sentences of $\mathcal{L}$ are those formulas of $\mathcal{L}$ which do not contain free variables

The notion of a free variable here is the standard one. Less standard are our bracketing devices. Since conjunction is associative, a conjunction does not come with brackets. Thus, a series of conjunctions can be displayed as $\phi_1 \wedge \ldots \wedge \phi_n$. (This will prove convenient when we

[33] We here thus exclusively deal with *anaphoric* pronouns. Bound pronouns, as in *Every man loves his car.*, have to be treated as bound variables, as usual.

[34] In an earlier paper (Dekker, 1994) a closely related system was presented as an update semantics. Although the system in that paper was called *Predicate Logic with Anaphora* as well, we prefer to call the system presented here *PLA*, and re-baptize the system in (Dekker, 1994) *DPLA*, *Dynamic Predicate Logic with Anaphora*.

present a normalization procedure below.) For this reason, however, we do need brackets to indicate the scope of other operators, as in $\neg(\phi)$ and $\exists x(\phi)$. As is usual, however, brackets are left out in case no confusion can arise. Material implication and universal quantification are defined in a standard way in terms of negation, conjunction, and existential quantification.

Like we said, existential quantifiers are used to model the interpretation of indefinite noun phrases, and they are associated with open places, to be filled by possibly intended referents. Sequences of indefinites thus will be matched with sequences of individuals in the semantics. However, the very same sequences are used to determine the possible referents of pronouns. These sequences not only reflect what are the possible referents of indefinites in a sentences, but also what might be possible referents of indefinites in preceding discourse. In order to be able to separate these two uses we will now define what is called the 'length' and the 'scope' of a formula. The length of a formula equals the number of indefinites in it which are used with referential intentions.[35] The scope of a formula is its length plus the number of indefinites which pronouns require to be there in preceding discourse. The two numbers are defined as follows:

DEFINITION 10 (Length and Scope of a Formula).

- $n(Rt_1 \ldots t_m) = 0 \qquad n(\neg\phi) = 0$
  $n(\exists x\phi) = n(\phi) + 1 \qquad n(\phi \wedge \psi) = n(\phi) + n(\psi)$
- $s(Rt_1 \ldots t_m) = \bigcup\{j \mid \exists i(1 \leq i \leq m \ \&\ t_i \equiv \mathtt{p}_j)\}$
  $s(\neg\phi) = s(\phi) - n(\phi)$
  $s(\exists x\phi) = s(\phi) + 1$
  $s(\phi \wedge \psi) = (s(\phi) + n(\psi)) \cup s(\psi)$

Since indefinites are associated with existential quantifiers, and since we neglect names or individual constants, atomic formulas do not bring up new referents. (Pronouns refer to referents from previous discourse.) An atomic formula thus has a length zero. As usual a negation is assumed to block the anaphoric or referential potential of existentials in its scope and therefore it has length zero as well. As an existential quantifier is associated with an indefinite noun phrase, it contributes one referent next to the ones contributed by the formula which it binds. The length of a conjunction of course is the sum of the lengths of its conjuncts.

[35] The fact that a sentence or discourse has length n can be conceived of as a 'fact about the discourse' or 'discourse information'. As observed by (Groenendijk et al., 1996; Stalnaker, 1998), this type of information, which is not strictly relevant to the contents of a discourse itself, is relevant to interpretation.

The definition of the scope of a formula is slightly more involved. If an atomic formula contains an anaphoric pronoun $\mathtt{p}_i$, it requires there to be (at least) i indefinites in preceding discourse. The scope of an atomic formula therefore equals the index of the most demanding pronouns among its terms.[36] If there are no pronouns around, its scope is 0, of course. The scope of a negation is that of the negated formula minus the length of that formula. Like we said, negation is assumed to block the referential or anaphoric potential of indefinites, but, of course, it leaves pronominal claims on preceding discourse intact. Since the scope of ϕ equals its length plus the number of previously required antecedents, the scope of $\neg\phi$ equals the scope of ϕ minus its length, that is, the number of previously required antecedents.

An existential quantifier again adds another element, since the scope of a formula includes its length. The definition of the scope of a conjunction is the most involved one. Surely, some of the demands of the second conjunct ψ on preceding discourse $(s(\phi) - n(\psi))$ can be met by the first conjunct, so that ψ requires there to be $s(\psi) - n(\psi)$ *minus* $n(\phi)$ indefinites before the whole conjunction. If this demand exceeds the demands of ϕ (which is given by $s(\phi) - n(\phi)$), then these are the demands of the whole conjunction, and its scope is this demanded number plus the length of the conjunction, i.e., the scope $s(\psi)$ of ψ itself. Otherwise, the scope of the conjunction equals the demand of the first conjunct plus the length of the conjunction, i.e., $s(\phi) - n(\phi) + n(\phi \wedge \psi) = s(\phi) + n(\psi)$. In any case, the scope of the conjunction is the largest number among $s(\phi) + n(\psi)$ and $s(\psi)$.[37]

By way of illustration, consider the following two sentences, with associated (schematic) translations:

(42) A bird hurt a nerd. $[\exists x(Bx \wedge \exists y(Ny \wedge Hxy))]$
(43) She deferred his concert. $[\exists z(Cz \wedge OFz\mathtt{p}_2 \wedge D\mathtt{p}_1 z)]$

The length of example [42] is two terms, and that of [43], under the associated translation, is one. The scope of the first is also two: since it contains no pronouns, its scope equals its length. The scope of the second is three. One pronoun invokes further restrictions on a term which has occurred two terms before the sentence, and the sentence also adds a term itself. The length of the conjunction of the two sentences, $n([42 \wedge 43])$, as well as its scope $s([42 \wedge 43])$, is three. For although the second conjunct refers two terms back, these terms are provided by the

[36] Notice that upon the von Neumann representation of the naturals the union of a set of numbers selects the highest number in that set.

[37] Since we assume unbracketed conjunctions, it has to be shown of course that $s(\phi \wedge \psi \wedge \chi)$ yields one and the same number no matter how it is calculated. Indeed $s((\phi \wedge \psi) \wedge \chi) = s(\phi \wedge (\psi \wedge \chi))$ is the largest number among $(s(\phi) + n(\psi) + n(\chi))$, $(s(\psi) + n(\chi))$ and $s(\chi)$, as is appropriate of course.

first conjunct. Generally, if the scope of a formula ϕ equals its length we will say that it is anaphorically closed, or resolved. In that case the formula contributes $n(\phi)$ intended referents but it does not impose (trans-sentential) conditions upon referents from previous discourse.

3.2. Satisfaction Semantics

Like we said, the semantics of *PLA* is stated as a satisfaction relation. *PLA*-satisfaction is defined relative to the usual first order models $M = \langle D, E\rangle$ which consist of a domain of individuals D and an interpretation E for the non-logical constants, and we employ variable assignments in a standard way. New is our use of sequences of individuals. These are intended to supply the intended referents of indefinite noun phrases and pronouns. We should mention here that these sequences model the intended referents in reversed order. The first element in a sequence corresponds to the last, 'most prominent', indefinite in a formula ϕ; the first $n(\phi)$ elements correspond to the ones contributed by ϕ, and the other ones are the targets of unresolved pronouns in ϕ.

We will write $M, g, e \models \phi$ if, relative to a model M and an assignment g, formula ϕ is satisfied by a sequence e of at least $s(\phi)$ individuals. The relation is defined as follows:

DEFINITION 11 (Interpretation of Terms; Satisfaction).

– $[x]_{g,e} = g(x)$ $\quad$ $[\mathtt{p}_i]_{M,g,e} = e_i$

– $M, g, e \models Rt_1 \dots t_m$ *iff* $\langle [t_1]_{g,e}, \dots, [t_m]_{g,e}\rangle \in E(R)$

$M, g, e \models \exists x\phi$ *iff* $M, g[x/e_1], e{-}1 \models \phi$

$M, g, e \models \neg\phi$ *iff* $\neg\exists c \in D^{n(\phi)}$: $M, g, ce \models \phi$

$M, g, e \models \phi \wedge \psi$ *iff* $M, g, e \models \psi$ *and* $M, g, e{-}n(\psi) \models \phi$

where $e - m$ *is the sequence* $e_{m+1}, e_{m+2}, \dots$

Atomic formulas are evaluated in a Tarskian way, but relative to both sequences of individuals and variable assignments. A pronoun $\mathtt{p}_i$ here selects the i-th individual in the sequence e, 'aiming' at a conjunction with preceding discourse which contributes at least i referents.

An existentially quantified formula $\exists x\phi$ behaves like an ordinary quantifier in the sense that it binds free occurrences of the variable x in its scope, but it also behaves like a free variable itself.[38] The quantifier addresses the first element in the sequence relative to which the

[38] Like existentially bound variables do in systems of dynamic semantics like *DPL*, one might say. A dynamically quantified variable in *DPL* behaves like a free one in the sense that it can be (re-)bound by other quantifiers by means of a technique called 'existential disclosure' (Dekker, 1993). For example, if x is bound by a dynamic existential quantifier in ψ in *DPL*, then it turns out to be universally quantified in $\forall y(\psi \wedge x = y)$. Thus, if $\psi \equiv \exists x Fx$, then $\forall y(\psi \wedge x = y)$ is equivalent to $\forall y Fy$.

quantified formula is evaluated. Intuitively, this element is a possibly intended referent, which, by being mentioned, can be said to contribute a discourse referent. Thus, if e is a sequence satisfying ϕ under an assignment of d to x, then de satisfies $\exists x\phi$. Notice that this referent d is put in the front of the sequence e. This reflects the fact that the initial referents of a sequence are the ones most recently introduced.[39]

The negation of a formula ϕ tells us that there is no way of finding referents for the indefinites in ϕ, or, rather, that there is no way to fill its existential holes with a sequence c. In this way a negation binds the existential holes in its scope, so that, e.g., $\neg\exists xFx$, as usual, means that no individual is F. Likewise, $\forall x(Fx \rightarrow Gx) \equiv \neg\exists x(Fx \wedge \neg Gx)$ will be true only if every F is G.[40]

If we evaluate a conjunction $\phi \wedge \psi$ relative to a sequence e, we evaluate the first conjunct ϕ relative to $e{-}n(\psi)$, which is e with the contribution of ψ stripped of. Intuitively, this says that ϕ is evaluated *before* ψ has contributed its referents. This notion of conjunction may be more easily understood if it is read in a 'constructive' way. If e satisfies ϕ, and ce satisfies ψ, where c fits the discourse referents contributed by ψ (so the length of c must be $n(\phi)$), then ce satisfies $\phi \wedge \psi$ as well.

Before we illustrate the above definitions by means of two examples, let us first spend a word on the PLA notions of truth, validity and entailment. Truth of a formula ϕ is specified relative to a model M, assignment g and a sequence e in M of at least $s(\phi) - n(\phi)$ individuals. The formula is said to be true iff there is a sequence $c \in D^{n(\phi)}$ of $n(\phi)$ individuals such that $M, g, ce \models \phi$; the formula is false otherwise. If a sequence e consists of less then $n(\phi)$ individuals, then ϕ is said to be undefined relative to any M, g and that sequence e.

A formula is valid iff it is true, if defined, relative to all models M, assignments g and sequences e. Finally, we say that ϕ entails ψ, iff ψ is true, if defined, relative to all models M, assignments g and sequences e, such that $M, g, e \models \phi$. We note that PLA thus accomodates a dynamic notion of entailment. Pronouns in a conclusion may relate back to indefinites in premise, as is one of the hallmarks of dynamic semantics.

[39] Thus, the last referent contributed by a formula $\exists x(Fx \wedge \exists y(Dy \wedge Bxy))$—i.e., the first element in a satisfying sequence e—must be an F who B-s a D; the one but last—the second in the sequence—must be a D who is B-ed by that F.

[40] Although, as we already suggested above, such a closing effect of negation is assumed in most systems of dynamic semantics, including DRT, one may wonder why this should be so. If, as we think, indefinites are generally used with referential intentions, then why should these disappear under negation? In section 4.4 we will given an explanation for this phenomenon by taking to heart the assumption that referential intentions are of a pragmatic nature.

3.3. Two Illustrations

Let us now see in some detail how *PLA* deals with anaphoric dependencies. First consider example (42) again:

(42) A bird hurt a nerd. $[\exists x(Bx \wedge \exists y(Ny \wedge Hxy))]$

As we have seen, the associated formula contributes two referents and it is anaphorically resolved: $n([42]) = s([42]) = 2$. The formula is satisfied by a sequence cde which consists of at least two individuals cd, and such that c is a bird who hurts nerd d (e may be of arbitrary length):

(44) $M, g, cde \models \exists x(Bx \wedge \exists y(Ny \wedge Hxy))$ iff
$M, g[x/c], de \models Bx \wedge \exists y(Ny \wedge Hxy)$ iff
$M, g[x/c], de \models Bx$ and $M, g[x/c], de \models \exists y(Ny \wedge Hxy)$ iff
$c \in E(B)$ and $M, g[x/c][y/d], e \models Ny \wedge Hxy$ iff
$c \in E(B)$, $d \in E(N)$ and $\langle c, d\rangle \in E(H)$

Now consider the follow-up (43):

(43) She deferred his concert. $[\exists z(Cz \wedge OFz\mathtt{p}_2 \wedge D\mathtt{p}_1 z)]$

(In order to speed up the present discussion, we neglect the uniqueness implications associated with the demonstrative "his", and simply assume that "his concert" is read as "a concert of him".) As we have already seen, formula [43] is not resolved. Its length $n([43])$ is 1 whereas its scope $n([43])$ is 3. The formula is satisfied by a sequence $bcde$ which consists of at least three individuals and such that it contributes b which is a concert of d which c deferred:

(45) $M, g, bcde \models \exists z(Cz \wedge OFz\mathtt{p}_2 \wedge D\mathtt{p}_1 z)$ iff
$M, g[z/b], cde \models Cz \wedge OFz\mathtt{p}_2 \wedge D\mathtt{p}_1 z$ iff
$b \in E(C)$, $\langle b, d\rangle \in E(OF)$ and $\langle c, b\rangle \in E(D)$

Finally let us look at the conjunction of the two:

(46) A bird hurt a nerd. She deferred his concert. $[42 \wedge 43]$

The length $n([46])$ of the conjunction is 3, as is its scope $s([46])$, so this conjunction is anaphorically resolved again.

(47) $M, g, bcde \models [42 \wedge 43]$ iff
$M, g, bcde \models [43]$ and $M, g, cde \models [42]$ iff
$b \in E(C)$, $\langle b, d\rangle \in E(OF)$, $\langle c, b\rangle \in E(D)$,
$c \in E(B)$, $d \in E(N)$ and $\langle c, d\rangle \in E(H)$

We see that the sequence $bcde$ satisfies the conjunction $[42 \wedge 43]$ iff b is a concert of nerd d which is deferred by bird c who hurt d.

Let us now turn to a so-called 'donkey sentence,' a conditional sentence with an anaphoric pronoun in the consequent clause which relates back to an indefinite noun phrase in its antecedent:

(48) If a bird hurt a nerd she deferred his concert. $[42 \rightarrow 43]$

Like we said, the PLA notion of implication is defined in the usual way in terms of negation and conjunction:

(49) $M, g, e \models \phi \rightarrow \psi$ iff $M, g, e \models \neg(\phi \wedge \neg\psi)$

Since negation in PLA binds existentials in its scope, also the existential holes in an implication are closed. However, the indefinites in the antecedent clause may resolve pronouns in the consequent clause, and, by the definition of negation, this has universal, instead of existential impact. If we work through the definitions, satisfaction of an implication can be spelled out further as follows:

(50) $M, g, e \models \phi \rightarrow \psi$ iff $M, g, e \models \neg(\phi \wedge \neg\psi)$ iff
$\forall c \in D^{n(\phi)}$: if $M, g, ce \models \phi$ then $\exists a \in D^{n(\psi)}$: $M, g, ace \models \psi$

For the example under consideration this amounts to the following:

(51) $M, g, e \models [42 \rightarrow 43]$ iff
$\forall cd \in D^2$: if $M, g, cde \models [42]$ then $\exists b \in D$: $M, g, bcde \models [43]$ iff
$\forall cd \in D^2$: if $c \in E(B)$, $d \in E(N)$ and $\langle c, d\rangle \in E(H)$
then $\exists b \in D$: $b \in E(C)$, $\langle b, d\rangle \in E(OF)$ and $\langle c, b\rangle \in E(D)$

In short, the implication [48] will be satisfied by any sequence e provided that every bird deferred a concert of every nerd which she hurt. It may be observed that these are the standard truth conditions associated with implications of the form $[42 \rightarrow 43]$ in systems of dynamic semantics.

The examples from this section may serve to show that the system of PLA, indeed, accounts for the phenomena dealt with in most systems of dynamic semantics, those involving anaphoric relationships across conjunctions and in donkey sentences. Obviously, this raises the question whether PLA, thus, involves a (radical) departure from classical, static systems of interpretation. As the next subsection sets out to show, it does not.

3.4. Normal Forms

In this subsection we compare the system of PLA with that of first order predicate logic (PL). It is shown that PLA really is a *proper* extension of PL with pronouns, and, in addition, that this addition does not affect expressive power. We present a meaning preserving normalization algorithm by means of which resolved pronouns can be eliminated. In the following discussion we assume an interpretation for PLA and PL relative to one and the same model M.

It is easily seen that PLA and PL agree on the interpretation of all formulas without pronouns:

OBSERVATION 13 (*PLA* and *PL*). *for all formulas* ϕ *of* PL*:*

– $M, g \models_{\mathrm{PL}} \phi$ *iff* $\exists e$*:* $M, g, e \models_{\mathrm{PLA}} \phi$

PLA, in other words, is a proper extension of *PL* with pronouns. So what can we say about *PLA*-formulas which do contain pronouns? Two things. Firstly, if a pronoun in a *PLA*-formula ϕ is unresolved, then it behaves like a free variable, of course. Secondly, and more interestingly, if a pronoun is resolved in ϕ, then we can remove it from the formula by an algorithm which produces so-called 'normal binding forms'.

In the following definition we use the term "closed" for all formulas which are either atomic or a negation. The normal binding form $(\phi)^\bullet$ of a *PLA*-formula ϕ then is defined as follows:

DEFINITION 12 (Normal Binding Form).

1. $(Rt_1 \dots t_m)^\bullet = Rt_1 \dots t_m \qquad (\neg\phi)^\bullet = \neg(\phi)^\bullet \qquad (\exists x\phi)^\bullet = \exists x(\phi)^\bullet$
2. $(\phi_1 \wedge \dots \wedge \phi_n)^\bullet = (\phi_1)^\bullet \wedge \dots \wedge (\phi_n)^\bullet$ *if* ϕ_j *is closed (for* $0 < j \leq n$*)*
3. $(\phi_1 \wedge \dots\ \exists x(\phi_i) \wedge \phi_{i+1} \dots \wedge \phi_n)^\bullet = \exists x(\phi_1 \wedge\ \ \wedge \phi_i \wedge \phi_{i+1} \dots \wedge \phi'_n)^\bullet$
   if a) ϕ_j *is closed (for* $i < j \leq n$*)*
   b) $\phi'_j = [x/1]\phi_j$ *(for* $i < j \leq n$*)*
   c) x *is not free in* ϕ_k *(for* $0 < k \leq n$ *and* $k \neq i$*)*

where $[x/i]\mathrm{p}_i = x$, $[x/i]\mathrm{p}_j = \mathrm{p}_{j-1}$ *if* $i < j$, *and* $[x/i]t = t$ *otherwise*
and $[x/i]Rt_1 \dots t_m = R[x/i]t_1 \dots [x/i]t_m\ \ [x/i]\neg\phi = \neg[x/i]\phi$
$[x/i](\phi \wedge \psi) = [x/i]\phi \wedge [x/i{+}n(\phi)]\psi\ \ [x/i]\exists x\phi = \exists x[x/i]\phi$

Explanation. Clause 3 is the one where the action takes place.[41] If a formula cannot be dealt with by the previous two clauses, it is a conjunction of a series of formulas which includes at least one existentially quantified formula. In that case the procedure raises an existential quantifier to the upper sentence level. If there is more than one, it first deals with the last existentially quantified formula $\exists x(\phi_i)$ (condition a), and adjusts the pronouns in the formulas following $\exists x(\phi_i)$ to the new scope configuration (clause b).[42] Of course (clause c) a quantifier

[41] The first two clauses are pretty straightforward. The normalization algorithm leaves atomic formulas unaffected, and when it deals with a negated formula $\neg\phi$ or with an existentially quantified formula $\exists x\phi$, then it only normalizes the embedded formula ϕ (clause 1). When we are dealing with a conjunction of a series of closed formulas, then the algorithm only normalizes each of these (clause 2).

[42] This adjustment may take three forms. Pronouns which refer back to the raised quantifier are replaced by the variable x and they are thus bound in the new configuration. Pronouns p_j which seek an antecedent before $\exists x(\phi_i)$ get a reduced index $j-1$, because we have removed a possible antecedent between their place of occurrence and that of their intended antecedent. Pronouns which are resolved *after* $\exists x(\phi_i)$ are left untouched.

$\exists x$ may only be raised if x is not free in the other formulas. Needless to say that, in such a case, we can always use an alphabetical variant of $\exists x(\phi_i)$. End of explanation.

By means of this normalization procedure all the (open) existentials in a formula ϕ are raised to the upper sentence level, and the net result is a formula of the form $\exists x_1 \ldots \exists x_i(\phi_1 \wedge \ldots \wedge \phi_n)$, in which all the ϕ_j are closed (atomic formulas or negated ones) and in normal binding form. The result is much like a discourse representation structure, in which all resolved pronouns reappear as bound variables.

To show how the algorithm works, let us consider the normalization of the donkey conjunction (46) discussed above:

(46) A bird hurt a nerd. She deferred his concert.
$$\exists x(Bx \wedge \exists y(Ny \wedge Hxy)) \wedge \exists z(Cz \wedge OFz\mathrm{p}_2 \wedge D\mathrm{p}_1 z)$$

Normalization of this formula proceeds as follows:

(52)
$$\begin{aligned}
&(\exists x(Bx \wedge \exists y(Ny \wedge Hxy)) \wedge \exists z(Cz \wedge OFz\mathrm{p}_2 \wedge D\mathrm{p}_1 z))^\bullet = \\
&\exists z(\exists x(Bx \wedge \exists y(Ny \wedge Hxy)) \wedge Cz \wedge OFz\mathrm{p}_2 \wedge D\mathrm{p}_1 z)^\bullet = \\
&\exists z\exists x(Bx \wedge \exists y(Ny \wedge Hxy) \wedge Cz \wedge OFz\mathrm{p}_1 \wedge Dxz)^\bullet = \\
&\exists z\exists x\exists y(Bx \wedge Ny \wedge Hxy \wedge Cz \wedge OFzy \wedge Dxz)^\bullet = \\
&\exists z\exists x\exists y(Bx \wedge Ny \wedge Hxy \wedge Cz \wedge OFzy \wedge Dxz)
\end{aligned}$$

The crucial point is of course that the normalization procedure is meaning preserving:

OBSERVATION 14 ($^\bullet$ Correctness). *for all formulas ϕ of* PLA*:*

– $M, g, e \models \phi$ *iff* $M, g, e \models \phi^\bullet$

If a formula ϕ is resolved, then pronouns in ϕ are removed in $(\phi)^\bullet$. Therefore, the observations (13) and (14) now tell us that all resolved formulas ϕ in *PLA* are equivalent to the translations of their normal forms in *PL*. Since, besides, unresolved pronouns behave like free variables, we, thus, have established a functional equivalence between *PLA* and *PL*: there is no difference in expressive power. The difference lies only in the fact that *PLA* accounts for intersentential anaphoric relationships by means of a compositional interpretation procedure based upon a standard notion of meaning.[43]

Since *PLA* has an arguably dynamic notion of interpretation, it make sense to compare it with the system of dynamic predicate logic

[43] It may be noticed here that *PLA* can be easily transformed into a higher order logic, if we allow information states which contain information about subjects of all types, and if the quantifiers are made to range over variables of all types. We may as well add λ-abstraction and functional application, so as to define the interpretation of a fragment of natural language in a fully compositional fashion, but this raises some technical complications, as the interpretation of a sentence can be a relation of any type. Such complications are not insurmountable, though.

(*DPL*) (Groenendijk and Stokhof, 1991), too. Briefly, the two systems relate in the following way. *PLA* and *DPL* are both systems with the expressive power of *PL*, which, in addition, account for the same empirical data concerned with intersentential anaphoric relationships. The difference is that *PLA* is arguably more conservative. *DPL* is not a proper extension of *PL* since its notions of scope (traditional) and binding (dynamic) are different, and, hence, it has a slightly different logic. *PLA*, however, is a proper extension of *PL* (with anaphoric pronouns). By the syntactic distinction between pronouns and free variables, we are able to treat anaphors without giving up the standard notions of scope and binding. Moreover, as will be shown in the next section, *PLA*'s notion of content is also classically static. The dynamics of *PLA* interpretation derives from a dynamic notion of conjunction, which can be explained as a pragmatically adjusted form of intersection. And as will be shown in more detail in the next section too, the *PLA* notion of content can be suitably related to appropriate notions of first update and support as well. As we argued in section 1, it is not clear what a proper notion of first order support should have to be like in order to correspond to updates of systems like *DPL*.

In the next section we will recast the findings of the present section in terms of the notion of an information aggregate which we developed in the previous section. Before that, however, it may be interesting to point out that the system of *PLA* could do without any variables as well. The possibility of eliminating resolved pronouns in favour of bound variables of course corresponds to the possibility of eliminating bound variables in suitably adjusted formulas containing pronouns only. The interpretation of an existentially quantified formula $\exists x\phi$ can be mimicked by a trivially true existential quantification $\exists x(x = x)$ in conjunction with a suitable adjustment of ϕ, in which occurrences of x are replaced by appropriate pronouns, and in which unresolved pronouns p_i are replaced by a pronoun p_{i+1}. All bound variables can, thus, be removed.[44] Surely all *free* variables can be dispensed with (by replacing them by pronouns with high enough indices). Thus, we see that the system of *PLA* not only accommodates an ordinary first order logic, but also a variable-free one, like the one presented in (Quine, 1960), or, more recently, in (Jacobson, 1999; van Eijck, 2000). With this observation we conclude our exposition of the system of *PLA*.

[44] Since the variable x in $\exists x(x = x)$ is arbitrary, and since the condition on x is vacuous, we can just as well write this as a variable free 'formula' $\exists$.

4. First Order Information Exchange

In the first section of this paper we have spotted some theoretical problems concerned with the use of indefinite expressions in discourse, with the information they can be said to convey to a hearer, and with the information of a speaker that can be said to support their use. In the second section we have studied formal notions of first order information, of first order information containment and information merge, that could be brought to bear upon the meaning and use of sentences with indefinite expressions in discourse. In the third section we have set out to give independent motivation for the interpretation of a first order language (PLA) with indefinite expressions and pronouns. This section is dedicated to combine and harvest the results.

We will first show that the satisfaction conditions for PLA-formulas from section 3 can be captured by the information aggregates of section 2. The dynamic (PLA) notion of conjunction then can be seen to be a merge of information which is basically a form of intersection that takes the obvious fact to heart that the information of one conjunct is presented before that of the other. We then show that the update potential of the sentences of our language can be stated by means of the dynamic merge operation. Support for the use of these sentences is next shown to be definable in terms of the information containment relation. We finally define the conditions under which first order information exchange remains safe, thereby generalizing the results of the introductory section from the propositional to the first order level.

4.1. Contents and Updates

The system of PLA was presented in the previous section as a case study showing how intersentential anaphoric relationships can be dealt with in a compositional and intuitively motivated fashion. We will now show that the information aggregates from section 2 constitute an appropriate means for specifying PLA's satisfaction conditions. Recall, from the introduction, that sets of possible worlds are an adequate means to specify the truth conditions of the formulas of a propositional language. Likewise, the first order satisfaction conditions of PLA-formulas can be specified as sets of pairs consisting of a world and a sequence satisfying that formula in that world. As indicated in section 2, sets of such sequences are information aggregates with linearly ordered domains.

Let us be more precise. We define the content of a PLA-formula relative to an intensional model $\mathcal{M} = \langle W, D, I \rangle$, which consists of a set of worlds W, a domain of individuals D and an interpretation function I, and such that for all $w \in W$: $\mathcal{M}_w = \langle D, I_w \rangle$ is an (extensional) model for *PLA*. The content $[\![\phi]\!]_{\mathcal{M},g}$ of a *PLA*-formula ϕ can be defined relative to such a model $\mathcal{M}$ and an assignment g in the following way:

DEFINITION 13 (PLA Contents).

– $[\![\phi]\!]_{\mathcal{M},g} = \{we \in W \times D^{s(\phi)} \mid \mathcal{M}_w, g, e \models \phi\}$

The content of a formula ϕ is specified as a set of worlds and sequences of individuals which satisfy the conditions imposed by ϕ in the associated worlds. Like we said before, a set of sequences consisting of a world and a sequence of n individuals corresponds to an information aggregate with domain $n = \{0, \ldots, n{-}1\}$.[45] Understood thus, the content of a *PLA*-formula really is an information aggregate.

Conceiving of the contents of formulas as information aggregates, the *PLA*-notion of conjunction from section 3 can now be seen to be an instance of the merging operation from section 2, for:

OBSERVATION 15 (Conjunction and Merge).

– $[\![\phi \wedge \psi]\!]_{\mathcal{M},g} = [\![\psi]\!]_{\mathcal{M},g} \sqcap_{+n(\psi)} [\![\phi]\!]_{\mathcal{M},g}$

(where $+m$ denotes the addition function f: $f(n) = n + m$)

This observation shows that the dynamics of interpretation in a system like *PLA* really resides in a dynamic notion of conjunction. Whereas the contents of sentences are stated in a static way, as sets of sequences of worlds and individuals, they are merged in a dynamic way, by means of which pronouns in one conjunct (ψ) are matched with indefinites in another (ϕ).

Observation (15) also helps to show what the dynamics of the *PLA* notion of conjunction actually consists in. For, also drawing from observation (9), we find that:

OBSERVATION 16 (Conjunction as Intersection).

– $[\![\phi \wedge \psi]\!]_{\mathcal{M},g} = [\![\psi]\!]_{\mathcal{M},g} \cap [+n(\psi)]_m [\![\phi]\!]_{\mathcal{M},g}$ *(where $m = d([\![\psi]\!]_{\mathcal{M},g})$)*

Skipping over the 'sugaring' subscript m, this equation says that the conjunction of ϕ with ψ involves the intersection of the contents of ψ with that of ϕ *after a translation* by $+n(\psi)$. An intuitive interpretation of this translation is this. The translation serves to add, to the contents of ϕ, the discourse information that in the conjunction of ϕ with ψ, $n(\psi)$ more referents have been mentioned after ϕ.[46] The referent mentioned

---

[45] A set σ of such sequences corresponds to the aggregate $\{g \in D_W^{n_v} \mid g(v)g(0)\ldots g(n{-}1) \in \sigma\}$.

[46] This form of dynamic conjunction thus reflects a change in (discourse) time, which can be said to comply with Frege's 'donkey'-imperative: 'Wenn jemand heute dasselbe sagen will, was er gestern das Wort "heute" gebrauchend ausgedrückt hat, so wird er dieses Wort durch "gestern" ersetzen." (Frege, 1918, p. 64) Thus, the 'character' of the conjunction of an utterance "Es regnet heute" on Wednesday and that of "Es regnet heute nicht" on Thursday is given by "Es regnete gestern und es regnet heute nicht" on Thursday.

last in an utterance of ϕ, is the referent mentioned $n(\psi)$ but last after a subsequent utterance of ψ. Observation (16) thus shows that the PLA-treatment of anaphoric relationships is really based on classical notions of meaning and conjunction, except that the last takes the fact into account that intended referents are mentioned by indefinite noun phrases in a discourse in a particular order of occurrence.

Indirectly, observation (15) serves to show another merit of the PLA-notion of conjunction. As we have already said above, PLA only deals with pretty simplistic types of anaphoric relationships. There is a one-way binding of terms only. Pronouns readdress the values of indefinites in preceding discourse and their indices rigidly determine which ones these are. However, nothing in the present set-up prevents us to incorporate more flexible resolution mechanisms. For, as a matter of fact, we can drop (or neglect!) the indices on pronouns, and incorporate a resolution algorithm which tells us which pronominal elements in one conjunct get matched with which open places in an other. Since the dynamic merge employed in (15) is only one of a whole variety of ways of conjoining information, all of these other ways of merging the information of two conjuncts are possible.[47]

Now we have a suitable characterization at our disposal of the contents of PLA-formulas as information aggregates, we can generalize the picture and take basic aspects of use into account. In the remainder of this subsection we study the update which the use of (a sentence modeled by) a PLA-formula may bring about in the information state of a hearer. We will show that such an update, too, consists of a dynamic merge of the information state of the hearer with the contents of the formula. In the next subsection we show that also the support which a speaker can be required to have for an utterance, is adequately rendered by means of the relation of information containment between information aggregates.

We first give a recursive formal definition of the update which an utterance of a PLA-formula brings about in the information state of a hearer who accepts the utterance. We assume that the (first order) information of the hearer is rendered by an information aggregate, and, for the sake of exposition it is convenient to assume that the domains of the states of linguistic agents are linearly ordered. (Nothing essential really hinges upon this assumption, though.) The update $(\tau)[\phi]_{\mathcal{M},g}$ of an information state τ with a formula ϕ is defined relative to a model

[47] Kataphoric and so-called 'Bach-Peters' dependencies can in principle be accounted for in terms of the general notion of an information product which we presented in section 2. Notice that such alternative notions of conjunction can be implemented in the system of PLA, without this requiring a change of the general architecture of the system.

$\mathcal{M}$ and an ordinary variable assignment g. Throughout we assume that τ has enough subjects to resolve pronouns in ϕ, that is, that $d(\tau) \geq s(\phi)-n(\phi)$ (otherwise, the update process is assumed to halt). Updates are defined as follows:

DEFINITION 14 (Update Calculus).

- $[x]_{g,e} = e_{g(x)} \qquad [\mathtt{p}_i]_{g,e} = e_i$
- $(\tau)[Rt_1 \dots t_m]_{\mathcal{M},g} = \{we \in \tau \mid \langle [t_1]_{g,e}, \dots, [t_m]_{g,e} \rangle \in I_w(R)\}$

  $(\tau)[\neg\phi]_{\mathcal{M},g} = \{we \in \tau \mid \neg\exists c \in D^{n(\phi)}\text{: } wce \in (\tau)[\phi]_{\mathcal{M},g}\}$

  $(\tau)[\exists x\phi]_{\mathcal{M},g} = \{wde \mid we \in (\tau)[\phi]_{\mathcal{M},g[x/d]}\}$

  $(\tau)[\phi \wedge \psi]_{\mathcal{M},g} = ((\tau)[\phi]_{\mathcal{M},g})[\psi]_{\mathcal{M},g}$

The update calculus defined here is almost identical with the update semantics presented in (Dekker, 1994) and close in spirit to those presented in (Groenendijk et al., 1996) and (Dekker, 1996). Let us quickly review the four clauses. If a formula is atomic or a negation, then the update of an information state τ with that formula preserves the sequences in τ which (*PLA*-)satisfy the formula and it removes the others. The update with an existentially quantified formula $\exists x\phi$ involves the extension of sequences e with an individual d, if e can be found in the update of τ with ϕ under an assignment of the witness d to x. As usual, conjunction in an update semantics involves function composition. The update brought about by $\phi\wedge\psi$ in τ equals the update brought about by ψ in the state that results from updating τ with ϕ.

It is easily seen that if an existentially quantified formula $\exists x\phi$ introduces a subject on top of a state τ, then a subsequent formula with a pronoun $\mathtt{p}_1$ refers back to that subject, so that, e.g., $((\tau)[\exists xFx])[G\mathtt{p}_1] = (\tau)[\exists x(Fx \wedge Gx)]$. Updates, in other words, not only may add information about the world, but they may also add subjects (by means of $\exists$) and further information about these subjects (by means of $\mathtt{p}$).

Updates like the ones defined above have been studied and motivated at various places, so we will not elaborate upon them here. (For further discussion the reader may consult the above references.) Interestingly, the update which an utterance of a formula brings about is closely related to its content. More in particular, the update of a state τ with a formula ϕ can just as well be obtained by merging the contents of ϕ with τ, for it can be shown that:

OBSERVATION 17 (Update and Merge). *for all sentences ϕ:*

- $(\tau)[\phi]_{\mathcal{M},g} = [\![\phi]\!]_{\mathcal{M},g} \sqcap_{+n(\phi)} \tau$

The link $+n(\phi)$ here makes sure that the $n(\phi)$ subjects contributed by ϕ are added as subjects to τ, whereas unresolved pronouns are matched with subjects of τ, (hopefully) introduced by discourse preceding ϕ.[48] Observation (17) shows that a fairly standard update semantics fits in the first order information models as we have presented them in section 2. More importantly, it shows that first order updates can be defined in terms of a static notion of content and a suitable dynamic merging operation.

As a dual of the previous observation we also find that:

OBSERVATION 18 (Content and Update). *for all sentences ϕ:*

– $[\![\phi]\!]_{\mathcal{M},g} = (W \times D^{(s(\phi)-n(\phi))})[\phi]_{\mathcal{M},g}$

In this observation, $(W \times D^{(s(\phi)-n(\phi))})$ indicates the aggregate which contains no information about $s(\phi) - n(\phi)$ individuals.[49] This observation shows that we can also define the contents of sentences in terms of their update potential. Together, the observations (17) and (18) show that the *PLA*-notions of content and update are interdefinable. This implies, e.g., that we can safely study the logic of first order updates without having to worry about the implications of this for our (intuitive understanding of the) notion of content.

4.2. Supported Contents

Let us now come back to the issue which we started out with in the opening sections of this paper. When can somebody's information state be said to support the utterance of an open proposition? According to the picture we sketched in section 2, all terms used in a piece of discourse must be rooted in the subjects of the information state of the utterer. The idea is that these are supposed to stand in for the real individuals which the information is originally drawn from and which these open propositions are intended to be about.

These observations can be spelled out in detail in terms of a support calculus for the language of *PLA*. The following definition gives a compositional specification of a support relation $\models_{\mathcal{M},g,l}$ between information states σ and formulas ϕ. This relation is defined relative to assignments g and links l which assign subjects of σ to variables and terms in ϕ. Since assignments relate variables with subjects of σ, the value of a variable x relative to an assignment g is also related to a sequence $we \in \sigma$, and it is the $g(x)$-th individual in the sequence e

[48] There is no guarantee, of course, that if someone updates with a sentence "He walks", then he has just before updated with a sentence with an indefinite noun phrase. We will return to this issue in section 4.3.

[49] The product $W \times D^{(s(\phi)-n(\phi))}$ corresponds to the full function space $D_W^{(s(\phi)-n(\phi))_v}$ which we employed in definition (4).

then. Links l are conceived of as lists of subjects $l_1 \dots l_{s(\phi)} \dots$ of $d(\sigma)$, so that l_i is the subject in σ supporting the i-th term in ϕ. Relative to a sequence we in σ the value of that term is e_{l_i}. The definition runs as follows:

DEFINITION 15 (Support Calculus).

- $[x]_{g,e,l} = e_{g(x)} \qquad [\mathtt{p}_i]_{g,e,l} = e_{l_i}$
- $\sigma \models_{\mathcal{M},g,l} Rt_1 \dots t_m$ *iff* $\forall we \in \sigma$: $\langle [t_1]_{g,e,l}, \dots, [t_m]_{g,e,l} \rangle \in I_w(R)$

  $\sigma \models_{\mathcal{M},g,l} \neg\phi$ *iff for no* $\sigma^+, l + n(\phi)$: $\sigma^+ \models_{\mathcal{M},g,l+n(\phi)} \phi$

  $\sigma \models_{\mathcal{M},g,l} \exists x\phi$ *iff* $\sigma \models_{\mathcal{M},g[x/l_1],l-1} \phi$

  $\sigma \models_{\mathcal{M},g,l} \phi \wedge \psi$ *iff* $\sigma \models_{\mathcal{M},g,l} \psi$ *and* $\sigma \models_{\mathcal{M},g,l-n(\psi)} \phi$

where, σ^+ *is any state* σ' *such that for some* n: $\emptyset \subset \sigma' \subseteq (\sigma \times D^n)$ *and* $l + n$ *is any sequence* l' *such that* $l' - n = l$

An information state σ supports an atomic formula if the formula is satisfied by all sequences considered possible in σ. A negated formula $\neg\phi$ is supported if there is no non-empty extension σ^+ of σ and an extended link $l+n(\phi)$ such that σ^+ supports ϕ. This definition is rather complicated, but, as we will see in subsection 4.4, it can be greatly simplified. Support of an existentially quantified formula $\exists x\phi$ involves the support of ϕ under an association of l_1 to x and relative to the link $l - 1$ which is l without l_1. Finally, a conjunction $\phi \wedge \psi$ is supported under a link l if ψ is supported under l and ϕ under the link obtained from l by removing the subjects linked to existential quantifiers in ψ. Observe that this definition of the support relation is close in form and spirit to that of the satisfaction relation. The main difference is that sequences of satisfying referents have been replaced by sequences of supporting subjects.

In order to see the support calculus at work, let us take a quick look at example (46), which we discussed before:

(46) A bird hurt a nerd. She deferred his concert.
$[\exists x(Bx \wedge \exists y(Ny \wedge Hxy)) \wedge \exists z(Cz \wedge OFz\mathtt{p}_2 \wedge D\mathtt{p}_1 z)]$

Abbreviating this formula as $\phi \wedge \psi$, we find that:

(53) $\sigma \models_{\mathcal{M},g,l} \phi \wedge \psi$ iff $\sigma \models_{\mathcal{M},g,l} \psi$ and $\sigma \models_{\mathcal{M},g,l-1} \phi$

Let us assume that $l = ijk$, the list of subjects labeled i, j and k in σ. The second conjunct of [46] is supported iff:

(54) $\sigma \models_{\mathcal{M},g,ijk} \exists z(Cz \wedge OFz\mathtt{p}_2 \wedge D\mathtt{p}_1 z)$ iff
$\sigma \models_{\mathcal{M},g[z/i],jk} Cz \wedge OFz\mathtt{p}_2 \wedge D\mathtt{p}_1 z$ iff
$\forall we \in \sigma$: $e_i \in I_w(C)$, $\langle e_i, e_k \rangle \in I_w(OF)$ and $\langle e_j, e_i \rangle \in I_w(D)$

The first conjunct then is supported iff:

(55) $\sigma \models_{\mathcal{M},g,jk} \exists x(Bx \wedge \exists y(Ny \wedge Hxy))$ iff
$\sigma \models_{\mathcal{M},g[x/j],k} Bx \wedge \exists y(Ny \wedge Hxy)$ iff
$\sigma \models_{\mathcal{M},g[x/j],\emptyset} Bx$ and $\sigma \models_{\mathcal{M},g[x/j],k} \exists y(Ny \wedge Hxy)$ iff
$\sigma \models_{\mathcal{M},g[x/j],\emptyset} Bx$, $\sigma \models_{\mathcal{M},g[x/j][y/k],\emptyset} Ny$ and Hxy iff
$\forall we \in \sigma$: $e_j \in I_w(B)$, $e_k \in I_w(N)$ and $\langle e_j, e_k\rangle \in I_w(H)$

In short, σ supports the conjunction [46] iff, for all sequences we in σ we find that, in w, e_i is a concert of nerd e_k which is deferred by bird e_j who hurt e_k.

As appears from our treatment of example (46), the support relation adequately deals with conjunctions in which one conjunct contains a pronoun anaphorically dependent upon an indefinite in another. More importantly, the support for such conjunctions has been specified in a fully compositional fashion. Although the support for a conjunction is defined in terms of the support of each of the conjuncts, this does not hamper the support for an indefinite in one conjunct to be corelated with that of a pronoun in another. The example also shows our notion of support to differ from both the notions of strong and weak dynamic entailment discussed in section 1. Support for the second conjunct of [46] does not require σ to have the information that all birds who hurt a nerd deferred his concert, nor that just some did. Rather, it requires that, in the individual labelled j in σ is known as a bird who hurt a nerd labelled k whose concert labelled i she deferred.

As can be seen from the analogies between the definitions of support and satisfaction, the support of a formula is intimately related to its content. Formally, it can be seen that an information state σ supports an utterance of ϕ iff σ contains the contents of ϕ. For:

OBSERVATION 19 (Support and Containment). *for any sentence ϕ and link $l \in d(\sigma)^{n(\phi)}$:*

– $\sigma \models_{\mathcal{M},g,l} \phi$ *iff* $\sigma \sqsubseteq_l [\![\phi]\!]_{\mathcal{M},g}$

This is an interesting result. In the introductory section we observed a close connection between propositional content and propositional support. This connection has now been generalized to the first order level.

As a dual to the last observation, it can also be seen that:

OBSERVATION 20 (Content and Support). *for all sentences ϕ:*

– $[\![\phi]\!]_{\mathcal{M},g} = \{we \in W \times D^{s(\phi)} \mid \{we\} \models_{\mathcal{M},g,i} \phi\}$

(Recall that i is the identity function.) The content of a formula consists of those sequences which, as a singleton, support it. This observation need not come as a surprise, as it is easily seen that $\{we\} \models_{\mathcal{M},g,i} \phi$ iff

$\mathcal{M}_w, g, e \models \phi$. Satisfaction corresponds to support by a singleton aggregate. Together, the two observations (19) and (20) show our notions of content and support to be interdefinable. Among other things, this implies that we can safely study the logic of first order support without having to worry about the implications of this for our (intuitive) notion of content.

4.3. Safe Information Exchange

The results obtained in the last two subsections, viz., those stated in the observations (15–20), can all be seen to be first order generalizations of similar results obtaining at the propositional level. As we observed in the introduction, the notions of propositional content, conjunction, update and support are closely related and these close connections have now been transposed to the first order level. As may be clear, this has been achieved by conceiving of information states and linguistic contents as information aggregates. Sets of possible worlds have been replaced by sets of worlds plus sequences of individuals in those worlds. Precisely this shift has enabled us to come up with suitable linking relations between (the subjects of) the aggregates, and, relative to these links, to define intuitive notions of dynamic conjunction, update and support.

However, one of our main interests in the opening section, and at the end of section 2, was concerned with the relation between update and support. We are now in a position to address the issue how safe updates can be if they are supported at a first order level.

When it comes to the exchange of information by means of assertions, there are, in principle, at least two actors involved. One is the speaker, responsible for an assertion, and the other ones are the addressees, who have to take the assertion at face value. Under idealized circumstances, the speaker's own information supports what she asserts, and an addressee, or hearer, updates his information state with the content of the speaker's assertion. As we indicated in the introduction, such an update ought to be safe, in the sense that the updated information state of the hearer should contain no information beyond (i) the information he already had before the assertion, and (ii) the information which the speaker may be required to have in support for her utterance. Since such an update thus depends upon the joint information contained in the separate aggregates of the speaker and the hearer, it is expedient to see if, or when, updates with supported utterances are indeed safe. Like we said, this seems to be an integrity principle for any theory of update and support.

The, theoretical, question is, thus, how do our notions of update and support relate? Under what conditions is the update of an information state τ with a formula ϕ supported by the information contained in τ and in that of a state σ which supports the utterance of ϕ? It must be clear by now that such a question can only be answered relative to appropriate linking functions (in the following observation, we use $m \circ -n$, which is the function m' with domain $\{i + n \mid i \in d(m)\}$ and such that $m'(i) = m(i - n)$):

OBSERVATION 21 (Update and Support). *for all sentences ϕ:*

– *if $\sigma \models_{\mathcal{M},g,l} \phi$, then $(\sigma \sqcap_m \tau) \models_{l \cup (m \circ -n(\phi))} (\tau)[\phi]_{\mathcal{M},g}$*

provided that $d(m) = d(\tau)$ and $l \cup (m \circ -n(\phi))$ is a function

Proof. Observation (19) tells us that $\sigma \models_{\mathcal{M},g,l} \phi$ iff $\sigma \sqsubseteq_l [\![\phi]\!]_{\mathcal{M},g}$ and observation (17) tells us that $(\tau)[\phi]_{\mathcal{M},g} = [\![\phi]\!]_{\mathcal{M},g} \sqcap_{+n(\phi)} \tau$, so with observation (12) we find:

(56) if $\sigma \models_{\mathcal{M},g,l} \phi$, then $(\sigma \sqcap_m \tau) \models_{l'} (\tau)[\phi]_{\mathcal{M},g}$ provided that $l \subseteq l'$ and $m = l' \circ +n(\phi)$

If $l' = l \cup (m \circ -n(\phi))$, clearly $l \subseteq l'$. We also have that if this l' is indeed a function, then $l'(j) = (m \circ -n(\phi))(j) = m(j - n(\phi))$ if $j > n(\phi)$ and $j \in d(l')$. So $(l' \circ +n(\phi))(i) = l'(i + n(\phi)) = m(i + n(\phi) - n(\phi)) = m(i)$ if $i + n(\phi) > n(\phi)$ and $i + n(\phi) \in d(l')$. Assuming definedness, $i + n(\phi) \in d((\tau)[\phi]_{\mathcal{M},g})$ iff $i \in d(\tau)$ iff $i \in d(m)$ iff $i + n(\phi) \in d(l')$. Surely, $i + n(\phi) > n(\phi)$, and, hence, $m = l' \circ +n(\phi)$ (if l' exists). End of proof.

What does this observation say, intuitively, or, more in particular, what does the requirement that "$l \cup (m \circ -n(\phi))$ is a function" really mean? Formally, it imposes two requirements upon the safe exchange of supported information. Firstly, it requires that if a subject in the updated state $(\tau)[\phi]$ derives from a discourse referent in ϕ, then it is traced back, by l, to the subject which supported this discourse referent in σ. Secondly, it ensures that every subject i in τ reappears as a subject $i + n(\phi)$ in $(\tau)[\phi]$, which can be traced back to the subject $m \circ -n(\phi)$ in the merge $\sigma \sqcap_m \tau$ of σ with the original state τ. Intuitively, the requirement that $l \cup (m \circ -n(\phi))$ is a function is that if, in the update of τ with ϕ resolutions take place between terms in ϕ and subjects of τ, then these have to correspond to a resolution of subjects in τ with subjects of σ (by m).

Surely, no such resolution between the subjects of τ and σ is needed when ϕ is resolved, because then $r(m)$ can be taken to be disjoint from $d(\sigma)$ and $l \cup (m \circ -n(\phi))$ is automatically a function. This means that, factually, updates with resolved sentences are generally safe. More interesting, then, are the cases in which the sentences are not resolved, and where the updates demand such resolutions.

Consider an utterance of an unresolved sentence like "A journalist has seen her." The update with such a sentence is supported by the information of the speaker and the initial information of the hearer under a link equating the subject of the speaker which supports her use of "her", with the one the hearer associates with that use of the pronoun. This can, but need not be, dubious business. Obviously, such an equation is not at all dubious if the utterance has been preceded by the assertion of, e.g., "A member of parliament visited the prime minister today." If this previous assertion was supported by a subject of the speaker, and if the hearer has accepted it, then his information state accommodates a subject which he can associate with the pronoun. All of his information about this subject then is supported by the information which the speaker has, provided that the speaker's support for the use of both "A member of parliament" in the first assertion, and that of "her" in the second, derives from one and the same subject.

Sure enough, pronouns may be used without there being any other individuals mentioned in preceding discourse, or in disregard of preceding discourse. For at any time, in or out of any discourse context, pronouns can be used to refer to individuals demonstratively or intentionally present. Imagine a bar, where a man enters, silently. One of the customers then may say to another "she just died", thereby referring to the man's cat whom both guests know was very sick. In cases like this, updates are still safe if the hearer is right in presupposing that he knows whom the speaker is talking about, and if the speaker's utterances are supported.

In general, then, the exchange of supported first order information is safe, but when it concerns the exchange of information by means of unresolved sentences additional care has to be given to the resolution of the involved pronouns. When it concerns anaphoric pronouns, with disambiguating indices, safety is preserved by means of the *PLA*-interpretation procedures, provided that the pronouns have indeed been preceded by sentences with sufficiently many antecedents. When it concerns non-anaphoric pronouns, extra-linguistic aspects of the context have to be taken into account. These, however, fall beyond the scope of the present paper.[50]

[50] Needless to say that we can have a safe exchange of information also if the exchanged information is actually incorrect. A speaker can be deceived, and she can even be blamed for being deceived, but she cannot be blamed for deceiving a hearer when she believes what she asserts while what she asserts is actually not true.

Whereas the meaning of sentences is deemed to be concerned with (objective) truth and reference, the felicity conditions of their use are at best stated in relation to the situations which the interlocutors believe they are in, and with respect to the individuals they intend to talk about. Since, of course, we want our notion of safe

4.4. Agents, Games and Witnesses

In the preceding subsections we have developed a notion of linguistic content, and notions of update and support for a first order language. Talk about update and support has drawn into the picture (abstract representatives of) the agents which actually use and interpret sentences of this language. These agents can be conceived of as players in a game of information exchange, which is regulated, among others, by the support and update rules which we specified above.

Thinking of update and support as rules regulating a game of information exchange is revealing in a number of respects. We already indicated that the support for a negated formula is defined in a relatively obscure way and, besides, that, as a matter of fact, it is a mystery why negations invoke an existential closure of the open places associated with indefinites noun phrases in their scope.[51] Facts like these can be explained when we conceive of informational language use as a game involving players with distinct roles.

Let us first return to the clause concerned with the support for negated formulas in definition (15). This clause is pretty complicated and it may be hard to grasp. Basically, it says that σ supports $\neg\phi$ if no non-absurd extension of σ supports ϕ. Actually, this requirement can be stated in a much more transparent manner, for:

OBSERVATION 22 (Negative Support). *for all resolved formulas ϕ:*

- $\sigma \models_{\mathcal{M},g} \neg\phi$ *iff* $(\sigma)[\phi]_{\mathcal{M},g} = \emptyset$

The support for $\neg\phi$ is most perspicuously defined in terms of the absurdity of an update with ϕ.[52] Notices that negation, thus, invokes a switch of roles. A *speaker*'s state supports an utterance of $\neg\phi$ if she as a *hearer* would go nuts upon hearing that ϕ.[53]

information exchange to be applicable to the use of linguistic expressions, we have stated it relative to the information states of the agents involved, not with respect to the world they actually find themselves in.

[51] Although not invariably. We want to come back to the empirical and theoretical aspects of this issue at another occasion.

[52] The reason why we have not adopted this easier formulation of the support for $\neg\phi$ in the original definition is that we set out to demonstrate that the support-relation can be given an independent compositional formulation.

[53] The idea that negation is a role switcher is of course familiar from systems of game-theoretical semantics (Hintikka and Sandu, 1997; Sandu, 1997). In game-theoretical semantics truth is defined in terms of the availability of a winning strategy for a verifier to verify a formula, and falsity in terms of the availability of a winning strategy for her opponent to falsify it. When it comes to playing games with negations the roles of the two players switch. For a verifier to verify $\neg\phi$, she has to take over the role of her opponent and set out to falsify ϕ. Conversely, the

Understanding negation as involving a switch of roles between (virtual) players is quite illuminating when we consider indefinites in the scope of a negation in our framework. Like we said, a speaker is generally assumed or required to use indefinite noun phrases with referential intentions. As we observed earlier, these referential requirements seem to disappear when these indefinites occur under a negation. Although such a closure effect of a negation seems to be empirically motivated, and although it is taken up in most systems of dynamic semantics, it has been a mystery why, indeed, negation behaves this way. Thinking of negation as a role switcher, however, much of the mystery dissolves.

The referential intentions associated with the use of indefinites are attributed to the speaker, who is supposed to be able to support what she says. No such requirements pertain to the hearer, for whom, generally, an indefinite introduces a new subject. Since the speaker's support for $\neg\exists x\phi$ can be characterized as a ban on an update of her state with $\exists x\phi$, the requirement of referential intentions automatically dissolves. The speaker is only required to have evidence that she, as a hearer, bans any update with $\exists x\phi$, no matter who, with whatever intentions, would try to attempt to bring about such an update.

The role switching effect of negations of course carries over to implications, since these are defined in terms negation and conjunction: $(\phi \rightarrow \psi) \equiv \neg(\phi \wedge \neg\psi)$. Using observation (22) we can establish the following equivalence:

OBSERVATION 23 (Conditional Support). *if $\phi \rightarrow \psi$ is resolved:*

– $\sigma \models_{\mathcal{M},g} (\phi \rightarrow \psi)$ *iff* $(\sigma)[\phi]_{\mathcal{M},g} \models_{\mathcal{M},g} \neg\neg\psi$

Neglecting, for the moment, the double negation on ψ, this observation shows that a state σ supports $\phi \rightarrow \psi$ iff the *update* of σ with ϕ yields a state that *supports* ψ. This time we witness a double role switch which can be rendered as follows: σ supports $\phi \rightarrow \psi$ iff, were anyone to come up with evidence for ϕ, σ's *update* with that evidence would *support* ψ. This understanding of implications resembles that of the intuitionists and game-theorists, who require evidence for an implication to consist in a method to transform evidence for the antecedent clause into evidence for the consequent clause. In our framework we can say that the update function on the supporting state σ itself turns evidence for ϕ into evidence for ψ.

opponent's attempt to falsify $\neg\phi$ must of course consist in an attempt to verify ϕ. Obviously, the framework of game-theoretical semantics is different from ours in that it invokes strictly competitive ('zero-sum') games, whereas ours is best characterized as involving games of cooperation or coordination.

The present, pragmatic, understanding of implications may serve to explain why indefinites in the antecedent of an implication are not associated with referential intentions, and do not contribute any discourse referents. They contribute virtual referents, which a possible interlocutor might come up with to support the antecedent of the implication. However, observation (23) indicates that indefinites in the consequent are also void of referential or anaphoric potential, due to its double negation. This, too, is a feature present in most systems of dynamic semantics, but one which we think can and should be remedied.

A speaker who claims support for an implication $\exists x\phi(x) \rightarrow \exists y\psi(y)$ may base this on evidence which gives her, for each x who is ϕ, a y who is ψ. That is, the support for such an implication can be taken to consist in a function assigning a ψ to each individual who is ϕ. If the speaker has such a function at her disposal, we can say that an update of her state with $\exists x\phi(x)$ really supports $\exists y\psi(y)$, without the double negation.

Actually, such a functional notion of support seems to be called for. Consider the following examples, familiar from the literature:

(57) If a book is printed with Zwijndrecht University Press, it always has an index, but usually, it is very incomplete. (adapted from Heim)

(58) Harvey courts a girl at every convention. She always comes to the banquet with him. (Karttunen)

The second pronoun in (57) is obviously dependent upon the books printed with Zwijndrecht University Press. The same goes for the pronoun "she" in (58), which appears to be functional upon the conventions which Harvey visits. Examples of functional readings of pronouns, indefinites, and *wh*-phrases are abundant in the literature. A (functional) extension of the referential and anaphoric potential of indefinites in the consequent of an implication, and in the nuclear scope of quantifiers, thus, seems to be called for, and has been presented in (Dekker, 2000b).

5. Extensions and Conclusion

In this paper we have formulated some basic intuitions about the notions of the content of a formula, the update of information which the utterance of such a formula may produce, and the support which can be used to motivate such an update. In the introduction we have seen that the three concepts can be formalized in a straightforward manner as long as they are concerned with propositional information about the world.

The paper has concentrated upon a first order extension of the three notions, which brings into the picture the exchange of information about individuals in the world. Such an extension has proved to be non-trivial, but we have succeeded in maintaining the two main results obtained at the propositional level. We have specified the satisfaction of the language of a first order logic with anaphora, an update calculus for that language and a support calculus. The three systems were shown to be coherently related in the sense that either one of them can be defined in terms of another. We have shown, furthermore, that also when it concerns the exchange of first order information, a neat characterization can be given of the situations in which these exchanges are felicitous and safe.

The main moral of our findings is, we think, twofold. In the first place we think that, when it comes to matters of use, it is crucial to extend the scope of current systems of dynamic semantics, and pay due attention to the role of the speaker. As was argued in section 1, a strict focus on the updates which utterances may bring about in the information state of a hearer does not seem to give us an (intuitive) handle on matters of support.[54] For that reason we have explicitly focused on the development of a separate notion of speaker support. In the second place, the notion of information developed in systems of dynamics has proved to be of great value. The reason seems to be that, although this notion of information is employed in only one specific way in these update systems, it does have the right structural properties to be applicable in a support calculus too. Whereas this notion of information allows for structural relationships between the interpretation of indefinites and that of subsequent pronouns, it equally allows for structural relationships between indefinites and precedent subjects of the speaker. Apparently, update and support constitute two sides of one and the same coin: the transparent exchange of first order information.

Our notions of content, update and support have been used to characterize the meaning and use of open propositions, but, by and large, we have concentrated on the use of indefinite noun phrases in *monologue*. As we have seen, the principles governing their use license subsequent uses of pronouns. Although indefinite noun phrases are not by themselves taken to be referring expressions, we claim they ought to be used with referential intentions, and it is the intended referents that subsequent pronouns may relate back to.

[54] Since content, update and support are interdefinable, it turns out to be possible to define support in terms of updates after all. A state σ thus can be seen to support ϕ under l iff it equals its merge under l with the update of the empty state with ϕ. Evidently, this is a roundabout way of stating things, which does not seem to leave much of our original intuitions about updates in place.

In the system of *PLA*, and in the update and support calculi, these links between pronouns and indefinites have been kept rather tight. The indices on pronouns unambiguously indicate which terms in preceding discourse they take as antecedents. Obviously, such an unambiguous language facilitates the study of the logic of update and support, and of the supported exchange of information. However, when it comes to the analysis of actual *dialogues* in real life, such simplifying assumptions have to be suspended and we should take a closer look at the meaning and support of the pronouns themselves, too.

When it comes to the meaning of pronouns, the system of *PLA* has not contributed much to the received wisdom formulated in observation (1). The interpretation of pronouns has been defined in conjunction with particular contexts, and in *PLA* these are linguistic contexts only. We may dub this a 'circumstantial' interpretation which we, from a compositional perspective, would prefer to reformulate in a more independent fashion. This is not a trivial issue.

It can be taken to belong to the meaning or character of a pronoun that it has to be used to refer to an individual which is salient (in some or other sense) in its immediate context of use. In other words, a specification of this character requires us to take into account these possible contexts of use, and, in relation to these contexts, something like the 'intentional space' of the interlocutors in that context. Such a specification requires us to delve into both the indexical and spatial-temporal aspects of the possible contexts of use, as well as various modal and intentional aspects. Surely, this is a task which goes beyond the scope of the present paper and we leave it as a subject of further study.

Also when it comes to the support for pronouns, much more, of course, can and should be said than what we have said above.

In the system of *PLA* an utterance of "He walks" is said to be supported if the speaker has somebody who walks in mind, and whom she has mentioned before. This is appropriate as long as we stick to monologues, but things get more involved when it comes to dialogue and cross speaker anaphora. Consider again the exchange (3–4), which initiated our search for a notion of first order support:

(3) *Zag:* A beatle just entered.
(4) *Nic:* He is going to sing a love song.

Intuitively, it may seem clear what is going on in an exchange like the above. Indefinite terms are supposed to be supported by subjects in an information state of the speaker, and these, in turn, may relate to definite individuals in the actual world. This is a fact which Nic can exploit. If he is confident that, when Zag had uttered (3), Zag had an individual in mind which Nic thinks is going to sing a love song, then Nic's utterance of (4) is properly supported. That is, Nic can be supposed to have information that the individual which he thinks Zag in (3) started out about is going to sing a love song.

In the case above, it seems to be clear what allows for Nic's use of a pronoun. It appears to be the demonstrative presence of an individual which Zag has referred to just before. However, the situation might be entirely different. Maybe no such individual is present, and maybe Nic and Zag know that *always*, when Zag talks about a beatle, he is really nagging about his little brother. If Nic then happens to know that Zag's brother is going to sing a love song, then his utterance seems fine, too. Yet another reason may be that Nic has heard a certain rumour that a beatle just entered, that he assumes Zag has heard the rumour as well, and that part of the rumour is that that beatle is going to sing a love song. Finally, Nic does not even need to refer to the beatle which Zag mentioned, if, for instance, he nods at a man who came in just after the beatle.

It may be clear, then, that although our formulation for the support for sentences with pronouns must be sharpened, it should not be too specific. Most generally, it can be required that a pronoun relates to a subject of the speaker, which represents an individual to her which can be assumed to be most salient, or which can at least be unambiguously identified as such. Stated as roughly as this, this principle aims to capture, under a single heading, the support for both single and cross speaker anaphora and demonstrative reference.

A further specification of this principle, however, seems to be a tedious affair. Spelling it out further requires us to take into account, not only the assumptions of the interlocutors about the world and the individuals in that world, but also their assumptions about the information which the other interlocutors have, as well as their assumptions about how that information relates to the individuals in the world they are discussing. Obviously, this requires more besides the first order notion of information which we assumed in this paper and this, then, is left as a subject for further study, too.

Acknowledgements

Parts and earlier versions of this paper have been presented at the first Jahrestagung of the Gesellschaft für Semantik *Sinn und Bedeutung* (Tübingen, December 1996), at the Mundial Workshop *Formal Semantics and Pragmatics of Dialogue* (München, March 1997), at the 10-th *European Summerschool in Logic, Language and Information* (ESSLLI, Saarbrücken, August 1998), and at the 21. Jahrestagung der Deutsche Gesellschaft für Sprachwissenschaft (Konstanz, February 1999). Apart from the audiences at these occasions, I wish to thank Maria Aloni, Nicholas Asher, Jelle Gerbrandy, Jeroen Groenendijk, Gerhard Jäger, Hans Kamp, Peter Krause, Barbara Partee, Massimo Poesio, Robert van Rooy, Ulrike Haas-Spohn, Martin Stokhof, Frank Veltman, Maria Wolters, Henk Zeevat, Ede Zimmermann and three anoynomous reviewers for helpful comments. The research for this work is supported by a fellowship from the Royal Netherlands Academy of Arts and Sciences (KNAW), which is gratefully acknowledged. The work forms part of the *Logic in Communication* project, which is funded by a Spinoza award from the Dutch Organization for Scientific Research (NWO).

References

Aloni, M.: 1997, Quantification in dynamic semantics, in P. Dekker, M. Stokhof, and Y. Venema (eds.), *Proceedings of the Eleventh Amsterdam Colloquium*, pp 73–87, ILLC, University of Amsterdam, Amsterdam

Beaver, D. I.: 1994, When variables don't vary enough, in L. Santelmann and M. Harvey (eds.), *Proceedings of SALT IV*, pp 35–60, DMLL Publications, Cornell University

Chastain, C.: 1975, Reference and context, in K. Gunderson (ed.), *Language, Mind, and Knowledge*, University of Minnesota Press, Minneapolis

Chierchia, G.: 1992, Anaphora and dynamic binding, *Linguistics and Philosophy* 15, 111–83

Dekker, P.: 1993, Existential disclosure, *Linguistics and Philosophy* 16(6), 561–588

Dekker, P.: 1994, Predicate logic with anaphora (seven inch version), in L. Santelmann and M. Harvey (eds.), *Proceedings of SALT IV*, pp 79–95, DMLL Publications, Cornell University

Dekker, P.: 1996, The values of variables in dynamic semantics, *Linguistics and Philosophy* 19(3), 211–257

Dekker, P.: 2000a, Coreference and representationalism, in K. von Heusinger and U. Egli (eds.), *Reference and Anaphorical Relations*, Kluwer, Dordrecht

Dekker, P.: 2000c, Support for update semantics (house version), in M. Poesio and D. Traum (eds.), *Proceedings of Götalog 2000*, pp 5–12, Department of Linguistics, Göteborg

Donnellan, K.: 1978, Speaker reference, descriptions and anaphora, in P. Cole (ed.), *Syntax and Semantics, Vol 9. Pragmatics*, pp 47–68, Academic Press, New York

Evans, G.: 1977, Pronouns, quantifiers and relative clauses (1), *The Canadian Journal of Philosophy* 7, 467–536, reprinted in: Evans, G., 1985, *Collected papers*, Dordrecht: Foris, pp 76–152

Evans, G.: 1982, *The Varieties of Reference*, Clarendon Press, Oxford, edited by John McDowell

Frege, G.: 1892, Über Sinn und Bedeutung, *Zeitschrift für Philosophie und philosophische Kritik* NF 100, 25–50

Frege, G.: 1918, Der Gedanke. Eine logische Untersuchung, *Beiträge zur Philosophie des deutschen Idealismus 2* 1, 58–77

Geach, P. T.: 1967, Intentional identity, *Journal of Philosophy* 74, 309–44

Gerbrandy, J.: 1999, Dynamic epistemic logic, in L. S. Moss, J. Ginzburg, and M. de Rijke (eds.), *Logic, Language and Computation, Vol II*, Kluwer, Stanford, CA, in print

Gerbrandy, J. and Groeneveld, W.: 1997, Reasoning about information change, *Journal of Logic, Language and Information* 6, 147–169

Groenendijk, J. and Stokhof, M.: 1991, Dynamic predicate logic, *Linguistics and Philosophy* 14(1), 39–100

Groenendijk, J., Stokhof, M., and Veltman, F.: 1996, Coreference and modality, in S. Lappin (ed.), *The Handbook of Contemporary Semantic Theory*, pp 179–213, Blackwell, Oxford

Heim, I.: 1982, *The Semantics of Definite and Indefinite Noun Phrases*, Ph.D. thesis, University of Massachusetts, Amherst, published in 1988 by Garland, New York.

Heim, I.: 1990, E-type pronouns and donkey anaphora, *Linguistics and Philosophy* 13(2), 137–78

Hintikka, J. and Sandu, G.: 1997, Game-theoretical semantics, in J. Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, pp 361–410, Elsevier, Dordrecht

Jacobson, P.: 1999, Towards a variable-free semantics, *Linguistics and Philosophy* 22, 117–84

Janssen, T. M.: 1986, *Foundations and Application of Montague Grammar. Part 1: Philosophy, Framework, Computer Science*, CWI Tract 19, CWI (Centre for Mathematics and Computer Science), Amsterdam

Kadmon, N.: 1987, *On Unique and Non-unique Reference and Asymmetric Quantification*, Ph.D. thesis, University of Massachusetts, Amherst

Kamp, H.: 1990, Prolegomena to a structural account of belief and other attitudes, in C. A. Anderson and J. Owens (eds.), *Propositional Attitudes*, pp 27–90, CSLI, Stanford

Kamp, H. and Reyle, U.: 1993, *From Discourse to Logic*, Kluwer, Dordrecht

Kaplan, D.: 1969, Quantifying in, in D. Davidson and J. Hintikka (eds.), *Words and Objections: Essays on the Work of W. V. Quine*, pp 221–243, Reidel, Dordrecht

Kripke, S.: 1972, Naming and necessity, in D. Davidson and G. Harman (eds.), *Semantics of Natural Languages*, pp 254–355, 763–769, Reidel, Dordrecht

Lewis, D.: 1979, Attitudes de dicto and de se, *The Philosophical Review* 88(4), 513–45

Montague, R.: 1974, Universal grammar, in R. Thomason (ed.), *Formal Philosophy. Selected papers of Richard Montague*, pp 222–46, Yale University Press, New Haven, originally published in 1970, *Theoria* 36, 373–98

Quine, W.: 1956, Quantifiers and propositional attitudes, *Journal of Philosophy* 53, 177–187, Reprinted in: W.V. Quine, 1966, *The Ways of Paradox and Other Essays*, Random House, New York

Quine, W. V.: 1960, Variables explained away, in *Proceedings of the American Philosophical Society*, Vol. 104, pp 343–347, Reprinted in W.V. Quine, 1966, *Selected Logic Papers*, pp. 227–235

Sandu, G.: 1997, On the theory of anaphora: dynamic predicate logic vs. game-theoretical semantics, *Linguistics and Philosophy* 20, 147–174

Stalnaker, R.: 1998, On the representation of context, *Journal of Logic, Language and Information* 7, 3–19

van Eijck, J.: 2000, Incremental dynamics, *Journal of Logic, Language and Information* 9, To appear

van Rooy, R.: 1997, *Attitudes and Changing Contexts*, Ph.D. thesis, IMS, Stuttgart

Zimmermann, E.: 1999, Remarks on the epistemic rôle of discourse referents, in L. S. Moss, J. Ginzburg, and M. de Rijke (eds.), *Logic, Language and Computation, Vol II*, Kluwer, Stanford, CA

Table of Contents