On First Order Information Exchange*

Paul Dekker
ILLC/University of Amsterdam

1 Introduction

Our main interest in this paper lies with the felicitous exchange of first order information. We are concerned with the question what licenses a speaker to convey open propositions or to utter sentences with pronouns. Although Grice's maxims of quality seem to give an appropriate answer to the question what licenses the felicitous exchange of propositional information, and although first order notions of truth and support have been developed within theories of discourse representation and dynamic interpretation, the question about first order licensing has to our opinion remained unanswered sofar.

In the course of this paper we argue for a specific formal notion of licensing which satisfies some intuitively plausible principles concerning first order information exchange. According to the view exposed in this paper, speakers may convey information about subjects which have been introduced in a discourse only if these correspond to subjects they themselves have information about. It makes no substantial difference if these subjects have been introduced in the discourse by the speakers themselves or by other agents. We will argue that indeed the last type of phenomenon, which has become known as "cross-speaker anaphora", is not structurally different from the phenomenon of ordinary, single speaker, anaphora. The main and rather trivial difference between single and cross-speaker anaphora resides in the fact that one agent may simply fail to have information about (the identity of) subjects introduced by other agents. Failing such information, it would clearly be inappropriate for him to exchange information about them.

We will proceed as follows. In section 2 we will critically discuss two preliminary sketches of the phenomenon of cross-speaker anaphora, one by Francez and Berg, and one by Groenendijk, Stokhof and Veltman, and we will argue that these analyses do not make sufficiently clear when or why cross-speaker anaphora is licensed. Then, in section 3, we will argue that, despite appearances, it is not even clear yet what licenses single speaker anaphora, and we develop a general notion of support which can be conceived of as a first order statement of Grice's maxims of quality. Then, in section 4, we will show that the principles governing single speaker anaphora also can be used to describe the licensing of cross-speaker anaphora. Section 5 concludes with some considerations of a more general nature.

*This note reports about work in progress on models of information exchange. I wish to thank Nicholas Asher, Jelle Gerbrandy, Jeroen Groenendijk, and Ede Zimmermann 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. Needless to say that all errors are mine.

2 Cross-Speaker Anaphora

In this section we will discuss some observations about cross-speaker anaphora which have been made in (Francez and Berg 1994) and (Groenendijk et al. 1997).

Multi-Agent DRT

For as far as we know, Nissim Francez and Jonathan Berg have been the first to pay substantial attention to the phenomenon appropriately dubbed "cross-speaker anaphora" (Francez and Berg 1994). Cross-speaker anaphora is a phenomenon occurring in multi-agent dialogue, "whereby one speaker refers anaphorically to a discourse entity introduced by *another* speaker." Here is an example:

(1) A: A man is walking in the park.
(2) B: He whistles.

Cross-speaker anaphora resembles so-called "donkey-anaphora", with the striking difference that the person who uses the anaphoric pronoun is different from the one who utters the antecedent (indefinite) noun phrase. This phenomenon raises a couple of questions, one which Francez and Berg call the 'commitment problem'. In their own words: "to what conditions about a discourse entity is a speaker committed when using cross-speaker anaphora?"

Francez and Berg present a first rough analysis of cross-speaker anaphora which has three distinctive aspects. First, for a proper representation of multi-agent dialogue, the various commitments of different agents are kept separate. Employing the representational format of discourse representation theory (DRT, henceforth), the various contributions of different agents are not lumped together in one cumulative representation, but each agent is assigned a separate section in the *multi-agent discourse representation* where his contributions are evaluated. The second feature of Francez and Berg's analysis is concerned with the representation of cross-speaker anaphora. If one agent B uses a pronoun y to refer back to a discourse marker x introduced by another agent A, then B's section in the multi-agent discourse representation is updated with the condition that y depends on x. The third aspect of Francez and Berg's analysis relates to the evaluation of such a dependency condition. If someone refers to a discourse entity introduced by another agent, then he is committed to (the satisfaction of) conditions "directly arising" from the introduction of the antecedent. This is called the "core condition commitment". The core condition includes the core description under which the relevant discourse entity is introduced

Our interest in this paper lies in the semantic (and pragmatic) impact of cross-speaker anaphora, and we think Francez and Berg's core condition commitment (CCC, henceforth) is both too strong and too weak in this respect. It seems the CCC inflexibly associates a descriptive content with any discourse entity, which any embedding of the entity must respect. This is too strong a requirement, because one can deny a discourse entity to have any property which other agents have ascribed it, this without contradicting oneself. Consider:

(3) A: Yesterday, Tom went to the Zoo with the youngest son of a colleague of his. The boy enjoyed it very much to spend the day with him.
(4) B: Sure, because '*he*' as a matter of fact was a nineteen year old *girl*, and because they went dancing, instead.

In the first example A introduces an entity in the discourse as a young son of a certain colleague of Tom. B apparently objects to this description, stating that it was in fact a nineteen year old girl whom Tom went out with. We think B may successfully refer to the entity introduced by A, but, clearly, he, thus, has not committed himself to the condition that the entity is both a young son of a colleague of A and a nineteen year old girl.

One might, of course, object that in the present example the core condition is not "x is the youngest son of a colleague of Tom", but "Tom went to the Zoo with x", which B is committed too. This might serve as an alternative explanation of the above example, but it won't do in general. Consider N's reply (6) to R's utterance of (5):

(5) R: A man is sleeping on a park bench over there.
(6) N: It is not a man, it is a woman, and she is not asleep, she is just sunbathing. Besides, it is not a parkbench.

N here denies almost *all* conditions on the discourse entity introduced by R. If we strip of the attributes which N denies it to have, the entity introduced by R is seen to be no more than just an x such that x is on a y over there. It seems this can hardly count as *the* core condition by means of which sentence (5) introduces its subject. Again, one might argue that it is not sentences that inflexibly associate their subjects with core conditions, but that it is utterances that do, but then it seems we are left without any analysis. Notice, for instance, that N might have *thought* "besides, it is not a parkbench", and *not* have *said* so. Thus N does not need to be committed to conditions on a discourse entity imposed by other agents, also if he doesn't explicitly deny them.

Besides, we think the main problem with the CCC is that it does not at all do justice to the intuition that cross-speaker anaphora involve some notion like talking about the same thing. In the little dialogues above A and B, and R and N seem to be speaking about the same thing, and it even seems that the corrections made by B and N only make sense if they are assumed to be speaking of the same entity as A and R are. However, nothing in the CCC requires that the two are speaking of the same individual. According to the CCC, reference to the same discourse entity by different agents does not at all constrain these to invoke the same individual, it at most requires the referents to be a member of the set of entities satisfying the core condition. We think Francez and Berg thus fail to account for an important characteristic of cross-speaker anaphora.

Multi-Speaker Update Semantics

Jeroen Groenendijk, Martin Stokhof and Frank Veltman have presented a further analysis of the phenomenon of cross-speaker anaphora, and their main conclusion is that it is not at all such an unconstrained phenomenon as Francez and Berg's analysis seems to suggest (Groenendijk et al. 1997). GSV discuss examples like the following:

(1) A man is walking in the park.
(7) He is wearing blue suede shoes.

Everybody will agree that there is nothing strange with an utterance of (7) after an utterance of (1) by one single speaker (or in a multi speaker *monologue*). However, GSV observe, there seems to be something marked if, without further context, one agent, A, utters (1) and another, B, replies with (7). If it is not clear whom exactly A is talking about, or which story he is about to tell, then B's reaction indeed is pretty odd.

This is not to say that B is never allowed to refer back to A's subject. For instance, if (1) is understood as a report about a directly visible situation, then A may be referring to a particular individual. By replying with (7) B then may observe that this man is wearing blue suede shoes. And also if B doesn't know whom A is talking about, or which story he is about to tell, B may ask questions about this subject, and draw conclusions about him on the basis of what was said.

So, after A has uttered (1), a reaction with any of the following two sentences would be perfectly alright:

(8) Is he wearing blue suede shoes?

(9) He must be wearing blue suede shoes then.

In short, cross-speaker anaphora appears to be possible if the relevant antecedents are 'definite' enough or when the pronouns are used with non-factual or non-assertoric force (and also in the case of multi-speaker monologue, when different speakers are said to act as a single agent).

GSV go some way in analyzing the second kind of cross-speaker anaphora in terms of a distinction between "direct" and "indirect" information. If people have direct information then it may be expressed straightforwardly, but if it is indirect information then it should be expressed using appropriate (modal) qualifications. Thus, when B doesn't know which individual A's utterance of (1) is about, then he at best has indirect information that that individual is wearing blue suede shoes. In that case he would not be licensed to utter (7) and he would only be licensed to utter (9) then. There is a certain appeal to this analysis, although it raises a couple of questions.

In the first place, it is not clear what is exactly meant with the terms "direct" and "indirect". The use of these terms invokes the suggestion that what is at stake here is a distinction between information for which we have direct, perceptual, evidence, as opposed to information obtained by other means, from hearsay, readwrite, watchtelevision, inference and so on. But this cannot be the relevant distinction it seems. Consider again B's reaction to A's remark that Tom went to the Zoo with a colleague's son. Clearly B's source may be other than perceptual, for instance plain gossip, tabloids, or just good thinking. Still there seems to be nothing basically wrong with his referring to A's subject without any modal qualification.

We think, that what *is* crucial is that one's utterances are supported by one's own information, not that of the other participants. If one's utterance, like B's (9) above, depends on information provided by other participants, or even refers to subjects introduced by other participants, then one has to add a modal qualification to indicate that one does not have independent support oneself. Thus, B's reply (9) should be understood as referring to the information expressed by A's utterance, and as saying that if indeed a man is walking in the park, then he is wearing blue suede shoes. In other words, anyone who buys A's claim that a man is walking in the park—B says—may buy from him that he wears blue suede shoes.

In the second place, the distinction between information that has, and information that has not to be modally qualified, at least involves a distinction between the information which one agent has from the information expressed by others in a dialogue (cf., Groenendijk et al. 1997, p. 24). But this seems to force us to revaluate the update semantic notion of meaning, for this distinction can only be maintained, it seems, if people do *not* in general update their information with what other agents say. An alternative dynamic conception of meaning which may be appropriate in this respect is one presented in (Beaver 1995), where the interpretation of a sentence is spelled out in terms of updates of what may be the speaker's picture of the common ground. In this paper we want to adopt a less involved view. We will stick to the fairly standard practice of assigning contents to information states and sentences alike. In our overall view, the dynamic resides in the combinatorics of these 'information containers'. Thus it is a straightforward matter to keep (the contents of) the sentences uttered by one agent, as well as his own information, separate from those of others.

The Definiteness Constraint

GSV also considered cross-speaker anaphora involving definite antecedents. When protagonist A introduce a definite subject, respondent B may be able to identify him. Because the respondent may have and convey information he himself has concerning the individual which (he thinks) his protagonist has introduced, the use of anaphoric pronouns in such cases can be seen to be licensed without requiring modal qualification. *GSV* also think of indefinite noun phrases as sometimes supplying definite antecedents. That is, they propose an analysis of sentence (1) as (10):

(1) A man is walking in the park.
(10) T is a man who is walking in the park.

where 'T' is a demonstratively referring phrase or a definite denoting phrase, which then may serve to license a pronoun by other agents.

Again we think there is both something promising and something problematic about *GSV*'s proposal. We agree that if an indefinite licenses non-modally qualified cross-speaker anaphora, this has to do with a definite or specific interpretation of the indefinite. However, we don't think this should be spelled out by saying that the indefinite noun phrase should be understood as a definite noun phrase. This seems to be a rather ad hoc move which we think is not necessary, because a specific analysis of indefinites can be built upon their dynamic semantic analysis by pragmatic means. As we will see in section 4, such a specific interpretation of indefinites may serve to explain the licensing of cross-speaker anaphora.

We also think it would be inappropriate to require a re-analysis of indefinite noun phrases as definite noun phrases, in order for them to license cross-speaker anaphora. Consider the following dialogue:

(11) H: A magistrate from Gotham village has confessed battering young girls.
N: They say he suspected them of sorcery. Do you know if more magistrates confessed?
H: I don't know.
N: Do you know who he is?
H: No idea, he preferred to remain anonymous.

We think that if H and N talk like this, they would probably agree that they talk about one and the same subject, perhaps one they read about in the (news-)paper. And probably they can come up with some identifying description when asked *whom* they talk about, e.g., "the magistrate reported about in *The Daily Planet* two days ago." So, as we said, it makes sense to say that the above dialogue involves an exchange of information about one and the same subject. However it is not at all clear from the dialogue which individual is concerned, or which definite must be used to refer to it. (There is a sense in which even H and N can be said not to know which individual it is.) But clearly we can understand H and N's utterances, without knowing which definite term, if any, is to replace the indefinite term "a magistrate of Gotham village". In other words, we can make perfectly good sense of the dialogue, without knowing how the example should be (re-)analyzed on *GSV*'s account.

To conclude, we think with *GSV* that non-modally qualified statements, made with or without pronouns, should provide information which the speakers themselves dispose of, and we also think that this invokes a requirement that they report on definite subjects they themselves have information about. We do not, however, think this excludes subjects introduced in the discourse by other agents by means of indefinite noun phrases. As we will argue in section four, the principles governing the licensing of cross-speaker anaphora follow from principles governing the licensing of utterances more in general. But before we can give substance to that claim, we have to say what these principles are, and then it appears that the principles governing the licensing of *single* speaker anaphora are not yet sufficiently understood.

3 First Order Licensing

In the next section we will argue that the phenomenon of cross-speaker anaphora is not structurally different from that of ordinary, single speaker, anaphora. The fact that cross-speaker anaphora has appeared to be a special type of phenomenon, we will argue, derives from a lack of understanding of what licenses the use of anaphora in ordinary cases. For this reason we first study the question what licenses the exchange of first order information in this section.

Licensing Open Propositions

When talking about the licensing of utterances, or about the justification for exchanging certain information, one tends to think of Grice's maxims of quality: do not say what you believe to be false and do not say that for which you lack adequate evidence. As long as we are dealing with ordinary propositions in a possible worlds framework, it seems easy enough to spell out what this kind of licensing amounts to. If s is the set of worlds compatible with one's information, then s should be a subset of p in order for one to be licensed to state that p. However, it appears to be not too straightforward to apply this notion of licensing in a first order framework. What should we think of as licensing the utterance of open formulas, formulas containing one or more free variables, or, if we think more in terms of natural language, sentences containing pronouns?

Two options for ruling the utterance of an open sentence licensed immediately spring to mind, and we can reject them just as straightforwardly. They involve the universal and the existential closure of the proposition. According to the first, one is licensed to state that Wx, or to utter "He whistles" only if one has the information that $\forall x\, Wx$, or that everything whistles. According to the second, one is licensed only if one has the information that $\exists x\, Wx$, or that something whistles. Clearly, this doesn't make much sense. The universal requirement is far too strong, the existential far too weak.

A mediating alternative consists in requiring that one has the information that the formula 'Wx', or the sentence "He whistles", holds of the individual referred to by means of the variable 'x', cq. the pronoun "he". This clearly makes much more sense than the two previous options, but it still is not appropriate. As we have seen above, in the dialogue about the Gotham village magistrate, it need not be clear which individual an agent refers to with a pronoun, not even for him- or herself. Clearly this does not need to render the utterances infelicitous, or unlicensed. The analysis can, again, be improved by requiring the agent to refer to a specific 'individual concept', such that he has the information that whatever is the value of that individual concept has the property denoted by W, cq. that it whistles. As a matter of fact, our analysis will come close to this one. Notice, however, that this notion of licensing essentially includes reference to individual concepts which an agent may or may not have in mind. For this reason it is not clear how such an analysis of licensing an utterance ϕ relates to the meaning of ϕ, standardly conceived. If we want to spell out the licensing of utterances in terms of their meanings, then it it seems it would not be appropriate conceive of their meanings as mere sets of possible worlds.

An alternative statement of what licenses the utterance of open propositions we may seek to find in dynamic systems of interpretation. Dynamic systems of interpretation make substantial sense of the (linguistic) context dependent interpretation of variables or pronouns and, therefore, we can expect them to accommodate a sensible notion of the licensed uttering of open propositions. Indeed, GSV suggest that it is a dynamic semantic notion of support which is relevant here, although they provide us with no clues as to what that notion precisely amounts to. Two dynamic notions of support which can be found in the literature naturally suggest themselves, but none of the two seems to be adequate.

The notion of support which has been used to define entailment in systems of dynamic semantics is a too strong notion. A formula Gx is supported by an information state s if in s only valuations of x are considered possible which map x into an individual which is a G. This seems to be a sensible qualification of the universal requirement considered above, but it is still too strong a requirement. After a licensed utterance of $\exists xFx$, an utterance of Gx would only be licensed within a dynamic semantic system of information update, if one were to have the information that $\forall x(Fx \rightarrow Gx)$. But this is too strong. Consider the following example:

(12) A member of parliament visited the queen yesterday.
(13) He was dead drunk.

An agent may utter these two sentences if he has the information that the queen was visited by a member of parliament yesterday, who was dead drunk. It does not require him to know that *every* member who visited the queen yesterday was dead drunk. However, one would be required to have this kind of universal information for the utterance of the second sentence if the notion of a licensed utterance was stated in terms of (strong) dynamic support.

An alternative notion of weak entailment is conceivable in a system of dynamic semantics, which corresponds to (Chierchia 1992)'s weak notion of implication. Adopting such a notion, an utterance of Gx can be said to be licensed after a licensed utterance of $\exists xFx$, if one were to have the information that $\exists x(Fx \wedge Gx)$. Although this seems to provide a sensible qualification of the mere existential closure of Gx, it still appears to be too weak. Consider A's reply to B's question about a man A started to talk about:

(14) A: Yesterday, a man came into my office who inquired after the secretary's office. ...
(15) B: Was he wearing a purple jogging suit?
(16) A: If it was Arnold, he was, and if it was somebody else, he was not.

According to the weak entailment notion of licensing, A might motivate his reply by saying that several men came into his office yesterday, and that one, Arnold, was wearing a purple jogging suit, and that the others were not. But clearly it would be odd for him to do so. For, if he had started a story about Arnold, he should simply have said that he was indeed wearing he a purple jogging suit, and if he had started a story about any one of the others, he should have said that he was not.

It seems that neither a classical (universal or existential) nor a dynamic (strong or weak) notion of support offers us an appropriate notion of licensing open propositions. Still there is one more option available. One can be said to express an open proposition that Fx because of the information one has about a certain subject, not a determinate entity, that it is F. Before we can make this notion of licensing more precise we first have to present the notion of information which we borrow from systems of dynamic semantics.

On Information Aggregates

What, we think, we need in a definition of first order licensing is a notion of aboutness which does not necessarily involves singular propositions. We want to make some (formal) sense of the (colloquial) locution that people (ought to) know what they are talking about, without this requiring them to construe this information as about determinate individuals in the real world. For this purpose very useful devices have already been developed, viz., *DRT*'s discourse representation structures which serve to picture first order information, and dynamic semantic information states which can be taken to model the contents of these discourse representation structures.

Discourse representation structures code information concerning subjects referred to by means of 'discourse markers' (Kamp 1981; Kamp and Reyle 1993). In systems of dynamic semantics this kind of information is fleshed out using sets of variable assignments, viz., the sets of assignments of values to the variables or discourse markers which are possible given the information which agents have about them (Heim 1982; Zeevat 1989; Dekker 1996). We like to view the mentioned information structures as modeling the (first order) contents of the information states of epistemic agents, but also, more generally, as that of texts and utterances and whatever else can be seen to contain first order information. For this reason we give them a general label "information aggregates".

Information aggregates are defined relative to a set of possible worlds W, a domain of individuals D and a set of variables V. (For the sake of simplicity, we assume that all worlds have the same domain.) Information aggregates contain information about 'subjects' which are referred to using these variables. Although it is not strictly necessary, we have chosen to use one distinguished variable v to refer to the world. For a finite set of individual variables $X \subseteq V$, with $v \notin X$, there is a domain of possible variable assignments $W^{\{v\}} \times D^X$, which we abbreviate as $D_W^{X_v}$. Information aggregates then are defined as follows:

Definition 1 (Information Aggregates) *For finite* $X \subseteq V$ *and* $v \notin X$
$\Sigma^X = \mathcal{P}(D_W^{X_v})$ *is the set of information aggregates about* X

The information that a subject x has a property $\mathcal{P}$ is captured by an aggregate σ such that each assignment $g \in \sigma$ assigns an individual d to x and a world w to v such that d has the property $\mathcal{P}$ in w. Similarly, the information that subjects x and z stand in the relation $\mathcal{R}$ is captured by means of sets of assignments σ such that for each $g \in \sigma$, the pair $\langle g(x), g(z)\rangle$ stands in the relation $\mathcal{R}$ in $g(v)$. Etc. (Notice that we use 'x', not only to refer to the variable x, but also to the possible values of x, i.e., the subject labeled by x. We trust this will not give rise to confusion.)

In what follows we will use $D(\sigma)$ to refer to the domain $D(\sigma)$ of σ, which is X iff $\sigma \in \Sigma^X$. Notice that we take it to be convenient to distinguish the absurd aggregates (empty sets) in Σ^X and Σ^Y for any two different sets X and Y. One may assume that the various empty sets are labeled by their domains. Given this $D(\emptyset^X) = X$, if $\emptyset \in \Sigma^X$.

Information aggregates can be put to use in a system interpretation along the following lines, roughly. Definite and indefinite terms (descriptions, names, pronouns) can all be taken to set up or refer to subjects. Thus a sentence like (17) can be associated with an aggregate like (18) (with $X = \{x, z, u, w\}$):

(17) An actress gave it to him on a rainy Sunday.
(18) $\{g \in D_W^{X_v} \mid g(x)$ *is an actress in* $g(v)$ *&* $g(z)$ *is a rainy Sunday in* $g(v)$ *&* $g(x)$ *gave* $g(u)$ *to* $g(w)$ *on* $g(z)$ *in* $g(v)\}$

This aggregate precisely contains the information about four subjects x, u, w and z that x is an actress who gave u to w on a rainy Sunday z. It will be clear that, for the purpose of modeling natural language interpretation, we need some device to code the information that the subjects u and w are pronominal and that they are in need of being resolved by the context. A sophisticated use of variables and a separate resolution relation help to do the trick. Consider the following example:

(19) A bear has escaped.
(20) It is chasing him.

Neglecting the second (demonstrative) pronoun 'him", the contents of (19) and (20) can be modeled by the aggregates (21) and (22), respectively:

(21) $\{g \in D_W^{\{x\}_v} \mid g(x)$ *is a bear in* $g(v)$ *&* $g(x)$ *has escaped in* $g(v)\}$
(22) $\{g \in D_W^{\{u,w\}_v} \mid g(u)$ *is chasing him in* $g(v)\}$

These two aggregates can be merged by resolving u into x, thus yielding the following product of information of the little discourse (19)–(20):

(23) $\{g \in D_W^{\{x,w\}_v} \mid g(x)$ *is a bear in* $g(v)$ *&* $g(x)$ *has escaped in* $g(v)$ *&* $g(x)$ *is chasing* $g(w)$ *in* $g(v)\}$

Thus the update potential of a sentence can be seen to be the sum total of its meaning (conceived of as an information aggregate) and a certain merging with other aggregates under some resolution relation. Section 4 presents some more details on this merging operation; one may consult (Dekker 1997) for a more sophisticated treatment of anaphora along these lines.

Support and Substantiation

If we say that an information aggregate contains certain information about subjects referred to by $x_1, \ldots, x_n$ this means no more than that the aggregate has structured that information around n subjects. The fact that these subjects are labeled by the variables $x_1, \ldots, x_n$ is not of any great importance. An aggregate which contains the same information about subjects labeled $y_1, \ldots, y_n$ counts as an equivalent aggregate, like two discourse representation structures which are alphabetical variants of each other.

Given this, if we want to compare the information in any two aggregates, we first have to say which subjects in one are related to which subjects of the other. In our support relation we make this explicit. We will say that an information aggregate σ supports an aggregate τ under a linking relation r iff σ has more information about the subjects $x_1, \ldots, x_n$ in the domain of r, then τ has about the corresponding subjects in the range of r. (The domain $D(r)$ of such a relation r is $\{x \mid \exists y\colon \langle x, y\rangle \in r\}$, and the range $R(r)$ is $\{y \mid \exists x\colon \langle x, y\rangle \in r\}$.) The support relation can be defined as follows:

Definition 2 (Support) *σ supports τ wrt r, $\sigma \trianglelefteq_r \tau$, iff $r \subseteq D(\sigma) \times D(\tau)$ and $\forall g \in \sigma \exists h \in \tau\colon g \equiv_r h$ (where $g \equiv_r h$ iff $g(v) = h(v)$ & $\forall \langle x, y\rangle \in r\colon g(x) = h(y)$*

An aggregate σ supports τ wrt r iff r is a relation between the domains of σ and τ and every possible assignment in σ corresponds to one in τ which is identical modulo r. Thus, σ has more information about the subjects in the domain r than τ has about the subjects in the range of r.

The supports relation is quite a general one, because it may relate information aggregates relative to all different kinds of linking relations. Let us take a quick look at a couple of examples:

(24) if $r = \emptyset$, then $\sigma \trianglelefteq_r \tau$ iff $\{g(v) \mid g \in \sigma\} \subseteq \{h(v) \mid h \in \tau\}$, that is, iff σ has more information than τ in a classical sense
(25) if $r = \{\langle x, y\rangle\}$, then $\sigma \trianglelefteq_r \tau$ iff $\{\langle g(v), g(x)\rangle \mid g \in \sigma\} \subseteq \{\langle h(v), h(y)\rangle \mid h \in \tau\}$, that is, iff σ has more information about x than τ about y; so if y has some property $\mathcal{P}$ according to τ, then x has that property according to σ too
(26) if $r = \{\langle x, y\rangle, \langle z, w\rangle\}$, then $\sigma \trianglelefteq_r \tau$ iff σ has more information about x and z than τ has about y and w ($\{\langle g(v), g(x), g(z)\rangle \mid g \in \sigma\} \subseteq \{\langle h(v), h(y), h(w)\rangle \mid h \in \tau\}$), so if y and w stand in some relation $\mathcal{R}$ according to τ, then x and z stand in that relation according to σ too

etcetera. It is relatively easily seen that $\trianglelefteq$ is reflexive under the identity relation, and transitive modulo composition of links:

(27) $\sigma \trianglelefteq_r \sigma$, for $r \subseteq \{\langle x, x\rangle \mid x \in D(\sigma)\}$;
(28) if $\rho \trianglelefteq_r \sigma$ and $\sigma \trianglelefteq_s \tau$, then $\rho \trianglelefteq_t \tau$, for $t = s \circ r$.

A most interesting type of support is one which holds between two aggregates σ and τ under a relation r having all τ's subjects in its range. It is special enough to give it a specific label.

Definition 3 (Substantiation) *σ substantiates τ (wrt l), $\sigma \mid\!\!<_{(l)} \tau$, iff $\exists r$: $\sigma \trianglelefteq_r \tau$ with $R(r) = D(\tau)$ (and $r = l$)*

(We want to refer, below, to the specific relation under which one aggregate substantiates the other, so that's why we have added a possible label 'l' for r.) If an aggregate σ substantiates τ, then for all subjects $y_1, \ldots, y_j$ in τ there are subjects $x_{i_1}, \ldots, x_{i_j}$ in σ such that σ contains more information about $x_{i_1}, \ldots, x_{i_j}$ then τ has about $y_1, \ldots, y_j$. Thus, σ can be said to contain at least as much information about at least as many subjects as τ. Substantiation is a reflexive and transitive relation, but it is not antisymmetric. This is as we want it to be. Two aggregates can be substantially equivalent without being identical, simply because they use different names for corresponding subjects.

Now we have defined our notions of information, support and substantiation, it is time we return to our original question: when can an utterance be said to be licensed by a certain information aggregate? We will say that an utterance is licensed by an aggregate σ iff it is substantiated, this relative to some linking relation l:

Definition 4 (Licensing) *σ licenses ϕ under l iff $\sigma \mid\!\!<_l [\![\phi]\!]$*

An agent with information aggregate σ is licensed to utter ϕ if the subjects in the interpretation of ϕ are rooted in subjects of σ. The subjects mentioned by an utterance of ϕ are required to correspond to subjects in σ which are at least dressed with properties in σ which are attributed to them by ϕ.

It will be clear that our notion of substantiation is both substantially weaker than the strong dynamic notion of support, and stronger than the weak dynamic notion of support. It is stronger than weak dynamic support because a substantiated utterance of "A man walks in the park" does not only commit us to the claim that there is a non-empty set of walking men in the park, it requires in addition that we have a particular subject in mind, whom we know as a man who is walking in the park. Substantiation is weaker then strong dynamic support, because a subsequent utterance of "He whistles" needs not commit us to the claim that every man who walks in the park whistles. All that it requires is that we have a subject in mind whom we know as a man who walks in the park and who whistles.

4 Linking Information Aggregates

Information aggregates in the actual world hang together in a world wide web of meaningful relationships. Aggregates may contain information about real individuals in virtue of factual (causal, perceptual) links with (individuals in) the actual world; they may also contain information about subjects of other aggregates in virtue of factual links of a conversational or intentional nature. In this section we want to try and apply the notion of an information aggregate in an analysis of cross-speaker anaphora. However, it is useful to take another look at *single* speaker anaphora first.

Single Speaker Anaphora

Consider again our little discourse (19)–(20), with associated information aggregates (21) and (22):

(19) A bear has escaped.

(21) $\{g \in D_W^{\{x\}_v} \mid g(x)$ *is a bear in* $g(v)$ *&* $g(x)$ *has escaped in* $g(v)\}$ $(= \sigma)$

(20) It is chasing him.

(22) $\{g \in D_W^{\{u\}_v} \mid g(u)$ *is chasing him in* $g(v)\}$ $(= \tau)$

An utterance of (19) by an agent a with information aggregate ρ is licensed if $\rho \ltimes_l \sigma$ under some relation l linking x to a subject of ρ. And he is licensed to utter (20) if $\rho \ltimes_{l'} \tau$ under some relation l' linking u to a subject of ρ.

In the case of an utterance of example (19) we may think of it as a pragmatic (or a dynamic semantic) fact about indefinites that they are used to refer to subjects which are not mentioned before, or to subjects the identity of which is deemed irrelevant. In example (20), however, a pronoun is used which needs to be resolved. An utterance of (20) thus invites us to look for a target to resolve the pronoun into and an obvious candidate is, of course, the subject which licenses the utterance of the first example. This candidate motivates the resolution which gave us the information product (23) above. For, if $\sigma \wedge_\varrho \tau$ is the product of the information of σ and τ under a resolution of u into x, and if l and l' link x and u to the same subject of ρ, then:

(29) $\rho \ltimes_l (\sigma \wedge_\varrho \tau)$ iff $\rho \ltimes_l \sigma$ and $\rho \ltimes_{l'} \tau$

We here touch upon an interesting point about our notion of substantiation which deserves further specification. Let us first give a general definition of the *information product* of a set Π of information aggregates relative to some linking function ϱ. A linking function ϱ assigns to each aggregate π in Π a finite linking relation $\varrho(\pi) \subseteq V \times D(\pi)$ and the product of Π under ϱ is an aggregate with domain $D(R(\varrho)) = \{x \in V \mid \exists y \exists \pi \in \Pi\colon \langle x, y\rangle \in \varrho(\pi)\}$. It is defined as follows:

Definition 5 (Information Product)

$$\bigwedge_\varrho \Pi = \{g \in D_W^{D(R(\varrho))_v} \mid \forall \pi \in \Pi\ \exists h \in \pi\colon g \equiv_{\varrho(\pi)} h\}$$

(Notice that it would make sense to use indexed information aggregates, because we might want to use the information of one information aggregate relative to different linking relations. We will neglect this point in the sequel.) The general point we want to make now is the following. Let $R(l) = D(R(\rho))$ and $\forall \pi \in \Pi$: $R(\varrho(\pi)) = D(\pi)$. Then:

(30) $\sigma \ltimes_l \bigwedge_\varrho \Pi$ iff $\forall \pi \in \Pi$: $\sigma \ltimes_{\pi(\varrho) \circ l} \pi$

Properly conceived, this fact amounts to what we would like to call compositionality of licensing. Define $[\![\phi \wedge_\varrho \psi]\!] = \bigwedge_\varrho \{[\![\phi]\!], [\![\psi]\!]\}$. We now find that somebody is licensed to utter ϕ and ψ under l and l', respectively, iff he is licensed to utter $\phi \wedge_\varrho \psi$ under i, where $\varrho([\![\phi]\!]) = l$, $\varrho([\![\psi]\!]) = l'$, and i is the identity relation on this product.

Let us now return to our examples (19) and (20). The product σ_b of the two aggregates σ and τ can be obtained using a linking function ϱ such that $\varrho(\sigma) = \{\langle x, x\rangle\}$ and $\varrho(\tau) = \{\langle x, u\rangle\}$. The pronoun "it" is then resolved because it is made coreferential with (the subject introduced by) the indefinite "a bear". Notice that the two linking relations involved in licensing the utterance of the two examples (19) and (20) thus motivate the anaphoric relation employed in the product.

Cross Speaker Anaphora

We have not yet considered the interpretation of the (assumed demonstrative) pronoun "him" in example (20). For an utterance of (20) to be licensed also this subject must be substantiated, and because we are concerned with a demonstrative here, it must be licensed by a subject presenting a demonstratively present individual. So, if someone interprets this specific utterance she will try and find out which individual the speaker may have had in mind when uttering the pronoun. Thus, the interpreter may construe the aggregate associated with example (20) also as being about that individual and relate it to a subject of her own, viz., the subject presenting that individual to the interpreter.

In a similar way, subjects introduced by means of indefinites noun phrases may be interpreted as being related to subjects of the interpreter herself. For although a speaker may use an indefinite noun phrase in order not to reveal the identity of the described source, still there should be a definite subject in the speaker's information aggregate licensing the indefinite, and the interpreter may find good reasons to think which real individual this subject is related to. But if that is the case, then the interpreter has a subject of her own which she takes the speaker to be speaking about, and about that subject she is licensed to add observations of her own.

Consider again the dialogue consisting of A's utterance of (1) and B's reply with (2):

(1) A man is walking in the park.
(2) He whistles.

If we come across a dialogue like this, it seems B knows (or thinks she knows) whom A is reporting about. Both A and B should have a subject in mind substantiating the contents of their respective utterances, and for the pronoun uttered by B to be resolved, we may take it to be the subject B thinks A reported about when he used the indefinite "a man". Thus, "he" is intended to be coreferential with "a man", and the information product of the two sentences can be obtained by resolving (the variable associated with) "he" into (the one associated with) the indefinite "a man". Thus, the two sentences together may provide the same information as if they had been uttered by one speaker, viz., that there is a man who is walking in the park and who whistles.

In the case considered above neither A nor B commits him- or herself to the product of information that there is a man who is walking in the park and who whistles. It suffices for A to have the information that a man is walking in the park, who constitutes a subject for him. For B it suffices that he has the information that the individual or subject B thinks A reported about whistles. Consider again the examples (5) and (6), uttered by R and N, respectively:

(5) A man is sleeping over there on a park bench.
(6) It is not a man, it is a woman, and she is not asleep, she is just sunbathing. Besides, it is not a parkbench.

Apparently, the second utterance relates to the individual which the first can be taken to be about. Probably, R and N have laid their eyes on a certain individual lying in the park, and R's utterance is licensed because he thinks the individual is a man who sleeps on a parkbench, and N's utterance is licensed because she takes it to be a sunbathing woman lying on something which is not a parkbench. That is, R's aggregate includes a subject presenting the individual as a man who sleeps on a parkbench, and N's aggregate includes a subject she takes R to be reporting about, and which presents the individual as a sunbathing woman lying on something which is not a parkbench. Clearly, N intends her utterance to be about the same individual as R's utterance, so the pronoun "it" can be seen to be coreferential with the indefinite "a man". But although the information product obtained under such a resolution of the pronoun is inconsistent, clearly, neither R nor N are committed to the validity of inconsistent information here.

The examples we have considered sofar seem to involve demonstratively present individuals, or, at least the assumption of demonstratively present individuals. But cross-speaker anaphora is not restricted to such cases. Agents are very well able to make more sophisticated guesses as to whom a speaker is referring to. Consider again the following dialogue:

(11) H: A magistrate from Gotham village has confessed battering young girls.
N: They say he suspected them of sorcery. Do you know if more magistrates confessed?
H: I don't know.
N: Do you know who he is?
H: No idea, he preferred to remain anonymous.

Apparently, H and N agree to be talking about one and the same magistrate, perhaps one they read about in the newspapers. Assuming that the information they gathered from the stories they read concerns a definite magistrate, they each may have subject substantiating their utterances; furthermore, assuming the other one's subject to relate to the same individual, they may take their utterances to be coreferential. Thus, the dialogue in (11) can be construed as being about one individual, expressing it is a Gotham village magistrate who confessed battering young girls, which he suspected of sorcery etc.

Notice that nothing in dialogue (11) prohibits H and N from not *really* speaking about the same thing. For one thing, the information of H and N may derive from different newspapers which reported about confessions of different magistrates. For another, the whole thing may have been made up by an overzealous reporter. However, although H and N may be deceived, or although they may be wrong in equating their subjects in the way they do, the sum of information (and the lack of information) conveyed in the dialogue (11) still is appropriately understood as being about one subject.

The Licensing Link

The cases of cross-speaker anaphora we have considered involve reference to the subjects which a speaker has in mind when making a certain utterance. Observe that this would not make much sense if we would not require a licensed utterance to be substantiated by a definite link. When a speaker sincerely introduces subjects, or refers back to them, he must have definite subjects of his own in mind on our account. Thus, there can be said to be a *factual*, intentional, link between the subjects of utterances and the subjects of information aggregates licensing them. Notice that if all that was required was that there be *some* such link, then it seems there would be no ground for locutions like 'the subjects talked about'.

We think that the requirement that definite links support utterances also helps in explaining a couple of discourse phenomena other than cross-speaker anaphora. Consider again an utterance of example (14), followed by B's question (15) and A's answer (31):

(14) A: Yesterday, a man came into my office who inquired after the secretary's office.
(15) B: Was he wearing a purple jogging suit?
(31) A: He was indeed.

Suppose two men visited A's office yesterday, who both inquired after the secretary's office, suppose one (Arnold) was, and one was not wearing a purple jogging suit, and suppose A is fully aware of this. It seems A's answer to B can be simply true then. But it seems hard to see why this would be the case, if A's utterance of (14) is not understood as being about a definite subject of A which he knows as a man wearing a purple jogging suit.

For the same reason A may start a conversation with (12), and correct a reaction by B if B misidentifies the subject A started talking about:

(12) A: The queen was visited by a member of parliament yesterday. ...
(13) B: (I heard.) He was dead drunk.
(32) A: No, no, he was not; not the one I mean. Anyway, ...

Clearly A might agree with B that the queen was visited by a drunk member of parliament, but he still may be right in pointing out that he wanted to talk about another MP, who was not drunk. A can simply deny he is talking about the drunk MP, or deny that the one he talks about was drunk. Notice that if licensing only required there to be *some* substantiating link, then B's utterance could be argued to force A to talk about the MP who was drunk, which is absurd, of course.

The requirement of a factual link may also help to rule out a reply like (16) to B's question (15) in the situation considered above:

(14) A: Yesterday, a man came into my office who inquired after the secretary's office.
(15) B: Was he wearing a purple jogging suit?
(16) A: (Yes, and no.) If it was Arnold, he was, and if it was somebody else, he was not.

If A were licensed to utter (14) without having a particular man in mind, then it seems (16) would be the most informative answer he could give to B's question. But, as we argued above, such a reply would be pretty odd in the sketched circumstances. On our explanation it is odd because A must have had a definite man in mind and, in the given circumstance, he could (and should) have given either a positive or a negative reply.

Finally consider the following continuation of dialogue (14)–(15). Suppose A replies:

(33) A: I don't know.

and suppose A continues the story he started to tell, and ends up concluding:

(34) A: So, after all it *was* the man with purple jogging suit.

On this scenario, if one were to inquire why A had initially replied that he didn't know whether the man was wearing a purple jogging suit a couple of answers may be perfectly acceptable. However, one answer we think is particularly marked: that A at the moment of uttering (33) did not yet know whom he was talking about, as if he had not yet made up his mind which story to tell.

The phenomena considered in this subsection all seem to be naturally accounted for on the assumption that subjects set up in a dialogue should be substantiated by specific subjects of the aggregates of the speakers. Although speakers may very well have only partial information about the real individuals concerned, and even partial information about their identity, they are not allowed to be confused about which of their own subjects they are referring to.

5 Conclusion

In this paper we have discussed the phenomenon of cross-speaker anaphora. We have first agreed with GSV that cross-speaker anaphora is a much more constrained phenomenon than single speaker anaphora, but we have argued that, upon closer inspection, it still appears to be governed by the very same principles. The principles governing the use of anaphora in single speaker monologue also govern the use of cross-speaker anaphora. In both cases speakers must be taken to express information about subjects of their own, and in both cases the subjects may be mentioned or introduced by other agents.

Our own analysis of (the licensing of) anaphora has made essential use of the idea that definite as well as indefinite noun phrases behave like free variables. Although definite noun phrases may come with uniqueness or familiarity assumptions, and indefinite noun phrases with novelty assumptions, they are all interpreted as a kind of functions. The meaning of a sentence has been conceived of as a function from the possible values which the terms in the sentence may have, to the proposition which the sentence express under these valuations. As we have seen, this conception has been fleshed out more generally in terms of sets of variable assignments, and we have used these 'information aggregates' to model the contents of both sentences and information states alike.

It is precisely our notion of an information aggregate which has allowed us to make sense of a notion of aboutness which does not necessarily invoke singular propositions or a too restrictive (too definite) interpretation of indefinite noun phrases. Yet, it does provide us with a definite enough notion of a subject which can be used to model speaker's reference, single and cross-speaker anaphoric reference, and linking more in general. Our first order extension of Grice's quality maxims in terms of substantiation may justify the assumption that one speaker refers to a certain subject which another one also may have information about.

Of course people may be deceived when they think they refer to an individual, or when they think they refer to the same individual as someone else and that's why we have chosen to model their information and the exchange of their information as involving (indefinite, maybe even fictional) subjects. However, we do think that these subjects come with something like a uniqueness assumption, viz., the assumption that there is a true answer to the question which individuals the subjects eventually stand in for. We also think that this assumption constitutes the main motivation for linking subjects. Upon the assumption that two subjects eventually relate to one and the same individual it makes sense to equate information about one with information about the other. In this way our story seems to fit in with that of the causal/intentional chain theory of reference, although we take the relevant assumptions to be basically pragmatic and defeasible.

Our basic concern in this paper has been with a proper notion of trading information. We have dealt with two aspects of information exchange, the licensing of utterances and the linking of subjects. Utterances are licensed, we claimed if they are substantiated, and any two subjects may be linked if they, eventually, relate to one and the same individual. The motivating idea is of course that if several epistemic agents have certain information concerning certain individuals, then accepting the information conveyed in discourses preserves veridicality of the information if the utterances are licensed and the resolutions appropriate. In the appendix we elaborate this point by showing that if all information is obtained from veridical perception, substantiated utterances and well-founded resolution then it is truthful.

Returning to the issue of anaphora, there is one main difference between single and cross-speaker anaphora. In the case of single speaker anaphora it is obvious of course that a speaker can refer to a subject of his own which he himself already has introduced to a discourse. In the case of cross-speaker anaphora this may be more difficult because the hearer must find reasons to match a speaker's subject with one of his own. This last type of linking may not in general be possible, and this fact may serve as an explanation for the fact that cross-speaker anaphora is more constrained than single speaker anaphora.

We have seen, in section 2, that cross-speaker anaphora is possible in more cases than those of the kind described sofar. GSV have mentioned multi-speaker monologue as a separate kind, but we think an analysis of multi-speaker monologue anaphora straightforwardly fits in with the analysis presented in this paper. What remains then are the cases of modally qualified cross-speaker anaphora. We think that these should fit in with the use of modally qualified single speaker anaphora.

Consider again the examples (1) and (9):

(1) A: A man is walking in the park.
(9) B: He must be wearing blue suede shoes then.

As we argued, B's utterance can be taken to refer explicitly to the information expressed in A's utterance, and state that if indeed a (or: the) man is walking in the park, then he is wearing blue suede shoes. This, strikingly resembles cases of modal subordination. Consider Roberts' (single speaker) sequence:

(35) A thief might break into the house.
(36) He would take the silver.

It seems sensible to interpret (36) here as referring to the proposition in the scope of the 'might'-operator in (35). The second sentence can be used to claim that, if the possible state of affairs mentioned by example (35) turns out to be actual, that is, if in fact a thief breaks into the house, then, this thief will take the silver.

So, one might argue, both (9) and (36) refer back to a previously discussed proposition, the first to one expressed by another agent, the second to one qualified by a modal operator. In both cases the proposition may invoke a subject which can be referred back to under a suitable modal qualification (cf., 'must' and 'would', respectively). We think the two types of examples are therefore amenable to an essentially similar treatment, as, for instance, that of (Geurts 1995) and (Frank 1997). Unfortunately, space prohibits us to go into the required detail here.

We want to conclude with a more general observation about the (dynamic) semantics/pragmatics interface. In this paper we have dissected the meanings of sentences from the update of information they may provide under some resolution relation. We think that the acceptability of information expressed by utterances and the linking of subjects is best conceived of as a pragmatic phenomenon, since it is concerned with utterances rather than sentences and with information about the context of utterance and about the information of other speech participants. Still, this does not at all need to hamper a systematic study of the logic of information update and exchange, for we can just take the outcome of a resolution process for granted and work on the assumption that various bits of information do get accepted. Systems of dynamic or update semantics can thus be seen to be specific systems of interpretation in which such studies are carried out precisely on these assumptions.

Appendix

In this appendix we give a slightly more general statement of the notion of a first order information model employed in this paper, to enable a proof that licensed transmission of information is truth preserving.

First order information models are built upon first order *models* M which are triples $\langle W, D, F\rangle$ consisting of a set of possible worlds W, a domain of individuals D, and an interpretation function F, which assigns an interpretation F_w for a first order language to each world $w \in W$. An *information aggregate* concerning the values of a set of variables X is a set of assignments each assigning a world w to a special variable v and individuals from D to X. The set of information aggregates over M given V, $\mathcal{A}_{M,V}$, is $\bigcup_{v \notin X \subseteq V} \mathcal{P}(W^{\{v\}} \times D^X)$. Given a model $M = \langle W, D, F\rangle$ and a set of variables V, an *information model* $\mathcal{I}_{M,V}$ built upon M and V is a triple $\langle \mathcal{A}_{M,V}, R, S\rangle$, where $\mathcal{A}_{M,V}$ is the set of information aggregates over M given V, R is a representation function, specifying what is the information of an individual $d \in D$ in a world $w \in W$, and S is a linking relation, specifying for each world what are the sources, if any, of the subjects of information aggregates of individuals in that world.

For each $w \in W$, the *representation function* R_w is a partial function from D to $\mathcal{A}$. The information of an information container d in the domain $D(R_w)$ of R_w is characterized by $R_w(d)$. The set of (epistemic) subjects D^e_w in a world w is characterized as a set of indexed variables (really: pairs): $\{x_d \mid d \in D(R_w) \ \& \ x \in D(R_w(d))\}$. In order to obtain a uniform notation, we also index the set of individuals D in w as $D^o_w = \{d_w \mid d \in D\}$. The set of potential sources of information D^*_w in w then can be characterized as $D^e_w \cup D^o_w$. For each world w, the *source relation* S_w is a relation between D^*_w and D^e_w. The idea is that if a and b are in D_w and $\langle x_a, y_b \rangle \in S_w$, then the subject labeled x by a constitutes a source of the subject labeled y by b. Furthermore, if $\langle d_w, x_a \rangle \in S_w$, then d constitutes the (causal / perceptual) source of the subject x of a.

From now on we will consider one world $w \in W$ only, and let $R = R_w$ and $S = S_w$. We are interested in the relation between information aggregates and (individuals in) the world w. Let S^* be the transitive closure of S and $l_{w,a} = \{\langle d, x \rangle \mid \langle d_w, x_a \rangle \in S^*\}$ the linking relation specifying the ultimate sources of the subjects of agent a's aggregate. Truthfulness of a representation then can be defined as follows. R is *truthful* for a given S iff $\exists g \in \sigma$: $g(v) = w$ and $\forall \langle x, d \rangle \in l_{w,a}$: $g(x) = d$. It turns out that truthfulness is preserved when the source relation satisfies two properties which we label 'well-foundedness' and 'conservativity'.

A source relation S is *well-founded* iff (i) it is finite, (ii) the inverse of its transitive closure S^* is a function from D^e to D^o and (iii) it is irreflexive in the sense that for no $\langle x_a, y_b \rangle \in S^*$: $a = b$. Under a well-founded source relation, every subject has one unique ultimate source $d \in D$. A source relation S can be said to be conservative iff the information about an agent's subjects derives from their sources. In order to state this perspicuously, we need a couple of notation conventions.

Let us equate the world w with an indexed information aggregate $\{w\}_w$, and generalize $\equiv$ to the effect that $w \equiv_{l_{w,a}} g$ iff $w = g(v)$ and $\forall \langle d, x \rangle \in l_{w,a}$: $d = g(x)$. Furthermore, let $\sigma_a = R(a)$ if $a \in D(R)$ and $\sigma_a = \{w\}$ if $a = w$. For $a \in D(R)$, let Π_a be the set of aggregates $\{\tau_c \mid \exists \langle z_c, x_a \rangle \in S\}$ to which σ_a is linked, and ϱ_a the linking function for σ_a such that for $\tau_c \in \Pi_a$: $\varrho_a(\tau_c) = \{\langle z, x \rangle \mid \langle z_c, x_a \rangle \in S\}$. Now S is *conservative* wrt R iff $\forall a \in D(R)$: $\bigwedge_{\varrho_a} \Pi_a \subseteq \sigma_a$. Notice that substantiation is conservative.

We now can state the following result.

(37) If S is well-founded and conservative wrt R, then R is truthful given S for all a such that σ_a contains at least one subject

(The restriction to aggregates with at least one subject is needed because S only relates subjects, not aggregates perse.) Sketch of the proof. Suppose σ_a contains at least one subject, and let $\varrho_{w,a}$ be the ultimate linking function for Π_a: $\varrho_{w,a}(\tau_c) = l_{w,c}$ for $\tau_c \in \Pi_a$. Then we find that

$$\bigwedge_{\varrho_{w,a}} \Pi_a \trianglelefteq_{l_{w,a}} \sigma \text{ if } S \text{ is conservative and well-founded}$$

(Proof omitted.) For any $c \in D$ such that $\tau_c \in \Pi_a$ let $\Pi_{a\setminus c} = (\Pi_a \setminus \{\tau_c\}) \cup \Pi_c$ and $\varrho_{w,a\setminus c}$ the corresponding linking function. Then $\bigwedge_{\varrho_{w,a\setminus c}} \Pi_{w,a\setminus c} \trianglelefteq_l \bigwedge_{\varrho_{w,a}} \Pi_{w,a}$, with $l = Id(D(\Pi_{w,a}))$, so, by composition $\Pi_{w,a\setminus c} \trianglelefteq_{l_{w,a}} \sigma$. Since S is finite, S^* is finite, as is the set $\{\tau_c \mid \exists \langle x_a, y_c \rangle \in S^*\}$. So in a finite number of steps we may remove all $\tau_c \in \Pi_a$ such that $c \neq w$, ending up in a case where $\{w\}_w \trianglelefteq_{l_{w,a}} \sigma$. From this is follows that $\exists g \in \sigma$: $g(v) = w$ and $\forall \langle x, d \rangle \in l_{w,a}$: $g(x) = d$. End of sketch.

The present result can be taken to show that if all actual information in the worlds derives from (i) veridical perception, (ii) substantiated (conservative) exchange, and (iii) well-founded resolution, then all information is truthful.

Bibliography

Beaver, D.: 1995, *Presupposition and Assertion in Dynamic Semantics*, Ph.D. thesis, CCS, Edinburgh

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

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

Dekker, P.: 1997, On denoting descriptions, in P. Dekker, J. van der Does, and H. de Hoop (eds.), *Proceedings of the Groningen Workshop on Definites*, OTS, Utrecht, in preparation

Francez, N. and Berg, J.: 1994, *A multi-agent extension of DRT*, Technical Report LCL 94–9, Technion—Israel Institute of Technology

Frank, A.: 1997, *Context Dependence in Modal Constructions*, Ph.D. thesis, Universität Stuttgart, Stuttgart

Geurts, B.: 1995, *Presupposing*, Ph.D. thesis, Universität Osnabrück, Osnabrück

Groenendijk, J., Stokhof, M., and Veltman, F.: 1997, Coreference and modality in the context of multi-speaker dialogue, in H. Kamp and B. Partee (eds.), *Proceedings of the 1995 Workshop on Context Dependence*, IMS, Stuttgart and UFAL, Prague

Heim, I.: 1982, *The Semantics of Definite and Indefinite Noun Phrases*, Ph.D. thesis, University of Massachusetts, Amherst, Published in 1989 by Garland, New York.

Kamp, H.: 1981, A theory of truth and semantic representation, in J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Formal Methods in the Study of Language*, Mathematical Centre, Amsterdam, Reprinted in J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Truth, Interpretation and Information.* Foris, Dordrecht, 1984

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

Zeevat, H.: 1989, A compositional approach to discourse representation theory, *Linguistics and Philosophy* 12, 95–131