Epistemic Specificity from a Communication-theoretic Perspective

Hans Kamp and Ágnes Bende-Farkas
IMS, Stuttgart University and RIL–HAS Budapest

Abstract

This paper offers a DRT-based analysis of epistemic specificity. Following Farkas (1996), we distinguish between scopal, partitive and epistemic specificity. After arguing in Section 1 that the three main variants of specificity are irreducible to each other, the paper then focuses on epistemic specificity.

In the analysis of epistemically specific indefinites we distinguish between specific use and specific interpretation. Specific use is defined as a relation between (the semantic representation of) a linguistic form and (the representation of) the speaker's mental state: In the speaker's state the sentence containing the relevant indefinite corresponds to a singular proposition. Specific interpretation is in a sense a derivative notion: It characterises the representation constructed by the hearer just in case he construes an indefinite as having been used specifically by the speaker, and builds his own representation accordingly.

The representation language we employ is a descendant of the original DRT-language presented in Kamp and Reyle (1993). This framework is tailor-made for the representations of attitudes of cognitive agents (for a recent discussion see Kamp (2013)); in the analysis reported in the present paper it enables us to distinguish between (i) the representation of an utterance that is derived via (standard) linguistic analysis, and (ii)– (iii) the representations that the individual discourse participants have or construct for this utterance. The key concept of the analysis is the notion of an anchored entity representation: Anchored entity representations are constituents of mental states that are causally linked, via their anchors, to the entities that they represent. In general, when a speaker uses a noun phrase to refer to an entity represented by one of her entity representations and thereby activates an entity representation in the mind of the hearer, the anchor of the hearer's representation, ER_H, will often be structurally different from that of the speaker's own entity representation ER_S, although normally the two representations will be coreferential. There will be a structural difference in particular when the speaker refers to the entity represented by ER_S through making a specific use of an indefinite noun phrase. If the hearer takes her to have used the indefinite specifically, he will construct an entity representation ER_H whose anchor links it to its referent *as* the entity represented by the speaker's representation ER_S. Anchors of this type are called 'vicarious anchors'.

The paper concludes with a brief discussion of the linguistic status of epistemic specificity. Data from English have been taken to suggest that specificity is an epiphenomenon, viz. that it need not be captured at the level of grammatical representation. But data from Romanian appear to suggest otherwise: The behaviour of Romanian indefinites marked with the Accusative preposition *pe* suggests that the specificity properties of these indefinites need to be marked at the level of (compositional) semantics, viz. at the level of those representations that are obtained directly from syntactic input. Our (tentative) conclusion is that the linguistic status of specific indefinites can be subject to cross-linguistic variation.

Preface

This paper is the result of reduction twice over. The 'grandmother' document was put together for a course on indefinites that we offered at the European Summer School in Logic, Language and Information in Helsinki in 2001. These notes cover the range of different uses of indefinites from specificity all the way to those that the literature often treats as cases of incorporation. The first

half of these notes, concerned just with specific indefinites, we subsequently revised and expanded, so much so that it has grown into a book length manuscript, and is thus is no longer suitable as a journal article.

The present paper covers part of the contents of the latter manuscript. It focuses on our own proposal for an analysis and formal treatment of what Farkas has termed 'epistemic specificity' (Farkas (1996)). The proposal builds on a DRT-based account of propositional attitudes and attitude reports, of which an unpublished manuscript has been circulated since the mid-nineties. A German translation of most of this document can be found in Kamp (2003). A somewhat condensed version, but with additional material on shared attitudes, which is relevant to the proposal in this paper, is part of Kamp et al. (2011). For an early, more informal version of most of the main ideas see Kamp (1990). In recent years the representation formalism we will be using has been referred to as 'MSDRT' (for 'Mental State DRT').

1 Introduction

Two characterizations of specific indefinites are salient in the literature:

(i) Specific indefinites are understood as being about some particular referent;

(ii) specific indefinites have wider scope than warranted by the syntactic configurations they occur in.

Following Farkas (1996) we refer to these properties as *epistemic* specificity and *scopal* specificity. Epistemic specificity concerns the way in which the use of an indefinite is related to the information state of the speaker who uses it. The notion has close ties to that of the *referential* use of definite descriptions introduced by Donnellan (1966) and of *speaker's reference* in the sense of Kripke (1977/1990), though these are connections which we will not pursue in this paper. The concept of scopal specificity goes back to the work of Fodor and Sag (1982). Fodor & Sag saw epistemic and scopal specificity as two sides of the same coin; but in the light of subsequent work that equation now seems problematic (e.g. Farkas (1996), von Heusinger (2002a), Schwarzschild (2002)). We discuss this relationship in some detail in the original lecture notes and in the book based on them (Bende-Farkas and Kamp (2001), Bende-Farkas and Kamp (forthcoming)) — in this article we focus on epistemic specificity.

The central aim of this paper is to offer an analysis of epistemic specificity as invoking a special relation between language used in communication and the cognitive states of participants in communication.

The view we will put forward is that epistemically specific uses of indefinite *NP*s involve the selection of particular individuals as intended 'referents', with the result that a sentence containing an *NP* that is being used as an epistemically specific one expresses a singular proposition, which is about the selected individual. The central role for the descriptive content of the *NP* is in the selection of its "referent". Whether and when this information should also be treated as part of the semantic representation of the sentence is a non-trivial question and for all we can tell at this point the answer may vary depending on the form of the indefinite *NP* and the grammatical constructions into which it enters as a constituent. We will address this question briefly in the final part of the paper.

The structure of the paper is as follows. In Section 1.1 we ask what qualifies a noun phrase as an 'indefinite' and state what we take to be the two most salient and most widely used criteria. In Section 1.2 we review three notions of specificity that are found in the literature — epistemic specificity, scopal specificity and partitive or presuppositional specificity — in order to make clearer what this paper is and what it is not about. Section 1.3 recalls the account of pronouns with indefinite antecedents that was central in early presentations of Discourse Representation Theory (especially (Kamp (1981)) and *mutatis mutandis* in other dynamic theories. The reason for this is that we are using a version of DRT as the framework within which we develop our analysis of specificity. We want to state clearly and up front that the use we are going to make of this DRT

extension is largely independent of the account of indefinites and 'donkey pronouns' that is still seen by many as DRT's main *raison d'être*. The treatment of specificity we will propose is *not* an adaptation of this earlier account of donkey anaphora, and the part that DRT plays in the present account is quite different from the one it plays in the earlier donkey pronoun analysis.

MSDRT, the version of DRT we will be using, was designed for the representation of attitudes and the semantics of attitude reports. We provide an outline of MSDRT in Section 2. This presentation, while significantly more extensive than in that earlier version, is still only an outline, but we hope that it will enable readers unfamiliar with the MSDRT formalism to get enough of a sense of how it works to be able to understand the substance of our proposal for the analysis of epistemic specificity, which is presented in Section 3. Readers who want to know more about the formal details of MSDRT are referred to other publications in which these details are spelled out.

Section 4 addresses some of the methodological question that are raised by the proposal of Section 3: Are the effects of epistemic specificity 'semantic' or 'pragmatic'? And what does this question exactly come to? We tentatively conclude that an answer to the first of these questions may have to vary from case to case — from language to language, and, within single languages, between indefinite noun phrases of different forms. Two cases, which suggest opposite answers to the question, are discussed explicitly: indefinites in English *there*-insertion constructions and Romanian indefinites marked by the particle *pe*.

1.1   What are indefinites and what are they for?

This paper is about a certain kind of use of indefinite noun phrases in different human languages. What we will have to say about indefinites presupposes a cross-linguistically applicable characterisation of the term 'indefinite noun phrase'. We are looking for a characterisation in semantic or pragmatic terms. (We couldn't rely on a form-based definition of the sort 'an indefinite *NP* is an *NP* beginning with an indefinite article', for that would simply shift the burden to the notion 'indefinite article'. Other form-based definitions would run into the same difficulty.)

Two criteria for indefiniteness stand out in the literature, one semantic and one pragmatic (Heim (1982)). The semantic criterion is stated in (1).

(1)   Indefinite *NP*s are those phrases which can be used to express existential quantification in simple clauses.

The 'simple clauses' referred to in (1) are those in which a typical transitive verb, such as *kill*, combines with a singular definite subject (e.g. a proper name) and an indefinite in direct object position. As (2) shows, each of the indefinites listed there gives rise to an existential proposition when inserted into the object position of such a clause.

(2)   John killed a bird/some bird/one bird/two birds/several birds.

There is also a pragmatic characterisation of indefiniteness, often referred to as the *novelty condition*.

(3)   Indefinite *NP*s serve to introduce new entities into the discourse.

The comparison of (1) and (3) provokes two basic observations. First, (1) and (3) define indefinites in terms of properties that relate to very different theoretical perspectives. (1) is a statement about the contributions of indefinites to logical forms that admit the standard quantifiers of classical logic and which therefore presupposes such a notion of logical form; (3) is not about the logical forms of sentences, but about discourse information and about the contributions that individual sentences and their constituents make to it. So it presupposes discourse structure, something that many of the logic-based treatments of sentence semantics ignore, or delegate to pragmatics.

Correlated with this difference is that between the logical roles that (1) and (3) attribute to indefinites. (1) presents indefinites as devices of quantification. (3) suggests that they are more like terms, which stand in some sort of reference relation to the entities which they are used to introduce.

The second observation is that while (1) and (3) are very different in the concepts they employ, it is not hard to see that they will tend to go together. On the one hand a statement to the effect that there is something which satisfies certain conditions is naturally construed as making an entity with just those properties available for further reference, and thus — in that sense — as introducing it into the discourse.

On the other hand, if a type of expression serves the purpose of introducing new entities, then at least from the point of the hearer, for whom the entity introduced by an expression of this type will typically be new, the force of the statement that is made by the sentence containing the expression will be that there exists an entity of the kind introduced by the expression which has the properties that the sentence assigns to it.

Plausible as these considerations may seem, we do not see them as a knockdown argument that expressions fitting one of the characterisations given by (1) and (3) must necessarily also fit the other. In fact, one of the points we will argue in this paper is that there is a sense in which the two can come apart. But there is also a sense in which they do not, and there is no question but that the two characterisations, conjoined or on their own, have proved a useful guide in cross-linguistic studies of indefiniteness.

1.2 Epistemic, Scopal and Presuppositional Specificity

In the opening paragraph of Section 1 we distiguished between epistemic and scopal specificity. Here we would like to say a few words about the considerations which led to these two notions. In addition, we briefly mention a third specificity concept, according to which an indefinite *NP* is specific if it is presuppositional. In this subsection we discuss these three specificity concepts, so that it will be clear what concept is targeted in the analysis we propose in Section 3.

Epistemic Specificity

The earliest discussions of epistemic specificity (Fodor (1970/1979), Ioup (1977)) were concerned with indefinite *NP*s that occur as constituents of the complements of attitude verbs such as *believe, know* or *want*. An example is (4):

(4)   Mary believes that her husband is seeing a real estate agent.

This sentence can be interpreted as attributing to Mary a belief about some particular person — viz. that Mary's husband is seeing this person. In the philosophical literature such attitude attributions are called *de re* (with respect to the constituent in question, here the *NP a real estate agent*) and it is widely assumed that the content of such an attribution is a singular proposition, in which a certain predicate is said to be true of some real individual.

Note that the *de re* interpretation of (4) may be purely existential from the perspective of the speaker: She may be of the opinion that there must exist some person of whom Mary believes that her husband is seeing her, but without herself having any knowledge who this person is (beyond the fact that she is a real estate agent). On the other hand, it is also possible for the speaker herself to have some particular person in mind, of whom she assumes that Mary believes that her husband is seeing her and whom she herself assumes to be a real estate agent and describes in these terms.

The *de re* interpretation of (4) is to be distinguished from its *de dicto* interpretation. On its *de dicto* reading (4) attributes an existential belief to Mary, to the effect that there is some real estate agent or other whom her husband is seeing, but there need not be a particular individual of whom she believes this. This *de dicto* interpretation of Mary's belief does not entail an existential interpretation of the indefinite from the perspective of the speaker. She may attribute a purely existential *attitude* to Mary, to the effect that there is an *N* with the property *P*, and she may do that even without herself assuming that *N*s exist. (For the example in (4) this possibility is rather far-fetched; we all know that there are real estate agents. But in general this is not so. I can ascribe to someone the belief that there is a rejuvenating fountain in Florida without making the assumption (or, for that matter, myself believing) that there are rejuvenating fountains anywhere in the world.)

There is an obvious connection between the *de re* and *de dicto* interpretations of (4) on the one hand and, on the other hand, two different ways in which we can interpret the statement (5), made by Mary herself.

(5)   My husband is seeing a real estate agent.

(5) can be understood as expressing Mary's belief that there is some real estate agent or other whom her husband is seeing, although she herself is in the dark about that person's identity. This interpretation correlates with the *de dicto* interpretation of (4). But (5) can also be taken as the expression of a belief of Mary's which attributes to some particular real estate agent the property that Mary's husband is seeing her; this interpretation resembles the *de re* interpretation of (4).

These two ways of taking the indefinite *a real estate agent* in (5) — as part of expressing a *de re* belief and as part of expressing a *de dicto* belief — are commonly referred to as *(epistemically) specific* and *(epistemically) non-specific*. When so used, the terms 'specific' and 'non-specific' correlate closely with the terms *de re* and *de dicto*: The belief that we take Mary to express when we interpret *a real estate agent* in (5) specifically is the belief that (4) attributes to her on the *de re* interpretation of *a real estate agent* and the belief that we take her to express by interpreting the indefinite as used non-specifically is the one that is attributed to Mary on the *de dicto* interpretation of (4). One of the tasks of this paper is to make this correlation between the different interpretations of (4) and (5) formally explicit.

Scopal Specificity

According to the logical tradition of which we spoke in Section 1.1 quantifying *NP*s are those whose translation into logical form requires the use of quantifiers. Such a quantificational translation necessarily assigns a scope to the quantifier or quantifiers used. It thereby also assigns, indirectly, scope to the translated *NP*. This scope assignment must accord with the semantic intuitions that speakers associate with the sentence in which the quantifying *NP* occurs. That is, the scope of the translating quantifier must agree with the scope constraints that speakers take to apply to the *NP* that it is used to translate.

Quantifying *NP*s — among which figure prominently those English *NP*s which begin with *every, no* or *most* — are notoriously subject to "island constraints" (cf. Ross (1967) on islands and *wh*-movement, and Rodman (1976) on island effects with quantifier movement). For instance, a quantifying *NP* has to obey the Complex *NP* constraint: if it occurs within a relative clause it cannot outscope the *NP* that governs this clause. Thus, in (6a) the quantifying *NP*s *every student of mine/no student of mine/most students of mine* cannot be understood as having scope over the *NP exam questions* (nor over the subject *NP three colleagues*).[1]

(6)   a.   Three colleagues contributed exam questions that *every* student of mine/*no* student of
                mine/*most* students of mine answered.
        b.   Three colleagues contributed exam questions that *a* student of mine answered.

In contrast, for (6b), in which the direct object position is filled by the indefinite *NP a student of mine*, a "wide scope" interpretation seems possible: "There is some student of mine such that three colleagues contributed questions which this student answered." (A narrow scope interpretation, according to which three colleagues contributed questions each of which was answered by at least one student of mine, is of course also possible in this case.) As a rule, when an indefinite like *a student of mine* allows for a non-narrow scope reading, the narrow scope reading is also possible, and the sentence is thus ambiguous.[2]

---

[1]That is — focusing on the *NP every student of mine* — the sentence cannot mean that there are three colleagues such that for every student of mine there were questions those colleagues contributed and that that student answered; even less can the sentence mean that for every student of mine there were three colleagues who contributed questions that were answered by this student (where not only the questions but also the colleagues can vary from student to student).

[2]One notable exception to this rule are positive polarity items (PPIs), which cannot occur in the (direct) scope of clausemate negation (see e.g. Szabolcsi (2004), and Farkas (2002)). That is to say, PPIs such as Hungarian *valaki, valami* 'someone', 'something' do not as a rule have the option to be interpreted in the scope of clausemate negation.

Over the years there has been a growing awareness that many island constraints are not as robust as was thought at first. (See for instance Fox and Sauerland (1996) and the research it inspired). This is true in particular for quantifying noun phrases beginning with *every*, which sometimes have wider scope readings than the original island constraints permit. Nevertheless, the contrast between (6a) and (6b) appears to be genuine enough and it points to an important difference between indefinites and "true" quantifying *NP*s like those beginning with *every* or *no*. It is this difference that finds articulation in the term *scopal specificity*. As we propose to use the term, an indefinite *NP* is given a *scopally specific* interpretation if it is interpreted as having a scope that is wider than the one indicated by its syntactic position and thus wider than would be possible for a true quantifying *NP* in that same position.

As we have described them, epistemic and scopal specificity appear to be notions of quite different kinds. But there is one obvious connection between them: if an epistemically specific use of an indefinite is one in which it is used to talk about some particular entity, then in the logical form of an utterance in which an indefinite is used in this way the existential quantifier corresponding to the indefinite must have maximal scope. However, scopal specificity is not just a matter of having maximal as opposed to local scope (pace Fodor and Sag (1982), who explicitly claimed that the indefinites could only have maximal or strictly local scope, i.e. they lack readings with so-called intermediate scope). It is a more complex phenomenon in that indefinites can get, besides a local scope which obeys the familiar island constraints for quantifiers and a maximal scope which includes all other material from the sentence, also 'intermediate scope' interpretations, which are neither local nor maximal. This important fact was first observed in Farkas (1981). An example of a sentence with an indefinite for which an intermediate interpretation is both possible and natural is the following variant from Chierchia (2002).

(7)   Every linguist has looked at all solutions that have been proposed in the literature for some linguistic problem.

Evidently, when an indefinite is interpreted as having intermediate scope, that use precludes it as a case of epistemic specificity. If it is true, as we suspect it is, that there are mechanisms which can lead to both intermediate scope and wide scope interpretations, then that suggests that there can be cases where an indefinite gets maximal scope even though it is not interpreted as an instance of epistemic specificity. If so, then epistemic specificity entails maximal scope but not conversely.

For reasons of space we cannot pursue the connections between epistemic and scopal specificity further in this paper. (But see the book-length manuscript Bende-Farkas and Kamp (forthcoming) currently in preparation.)

Presuppositional Specificity and Partitive Indefinites

There is a third notion of specificity that is found in the literature, that of partitive specificity. This notion was originally introduced in Enç (1991). See also Farkas (1996) for an elaboration on this and the other two notions of specificity, and Schwarzschild (2002), Portner and Yabushita (2002) and Breheny (2003) for proposals that derive specificity *tout court* from partitive specificity. As the term 'partitive specificity' implies, this notion is intimately tied to that of a *partitive indefinite NP*. Partitive *NP*s are those *NP*s whose semantic values are confined to some explicitly or implicitly given set X. Explicit, or overt, partitives are *NP*s in which X is given by some constituent within the *NP*, often as a *PP* adjunct headed by the preposition *of*, as in (8) below (though other syntactic realisations are possible as well).

(i)   a.   János nem látott valakit
               John not saw   someone-ACC
               'John didn't see someone', NOT 'John didn't see anyone
        b.   Mari nem hiszi,   hogy János látott valakit
               Mary not believes that John saw   someone-ACC
               'Mary doesn't believe that John saw someone', or 'Mary doesn't believe that John saw anyone'

(English indefinites beginning with *some* as opposed to *a* notoriously also have a tendency in this direction.)
    While acknowledging this point, we in turn draw attention to the fact that scope relative to negation has been known to be markedly different from scope relative to quantifiers or other operators.

(8)   One/some/three of the girls

Besides *overt* partitives, which contain an explicit set-denoting NP, there are also *covert* partitives in which there is no surface constituent denoting the restricting set X; for these the set must be recovered from the discourse context. An example is the *NP two women* in the second sentence of (9).

(9)   A small crowd had gathered in front of the church. Two women were crying.

A common feature of the interpretations of the indefinites in (8) and (9) is that their values are confined to a well-defined (usually comparatively small) set X.

Not all partitive *NP*s are indefinites. Besides indefinites like *one of the women* we find *NP*s like *the first of the women*, *the one of the women who managed to escape*, and so on. But for us here it is only the partitive indefinites that matter. Furthermore, we focus (merely for convenience's sake) on overt partitive indefinites.

One general constraint on partitive *NP*s is that the set X must offer room to more than one possible value for the *NP* as a whole. A clear indication of this is that the *NP* governed by the preposition *of* in overt partitive *NP*s (like those in (8)) must always be a plural. That is the case also when the partitive *NP* containing it is in the singular. Singular partitive *NP*s with singular *of-NP*s like *∗one of the woman* are unequivocally ungrammatical. So apparently partitive indefinites where X gives room for just one value are proscribed. The reason, presumably, is that when there is no room for variation and X thus fixes the value for the embedding *NP* uniquely, then what is required is a definite and not an indefinite description. (That is, instead of the ungrammatical *∗one of the woman* the *NP* that should be used is the simple *the (one) woman*.)[3]

The observation that room for variation is essential to the grammaticality of partitive NPs is consistent with a view of partitive indefinites according to which they do not differ substantially from non-partitive indefinites, at least not in any of the respects that are relevant to our concerns in this paper. Not only are partitive indefinites like other indefinites in that they fit the characterisations in (1) and (3); they also allow for both specific and non-specific uses, in both the epistemic and in the scopal sense of that distinction. That partitive indefinites can be used in a perfectly non-specific way is readily confirmed. An example is (10):[4]

(10)   The results indicate that for every exam one of the students must have got hold of the questions beforehand and passed them on to the others.

Sentence (10) could be uttered by someone whose only evidence is that for each exam the questions were stolen by one of the students (those who were due to take that exam), but without having any idea who in each case the culprit might have been. That is, (10) has an interpretation according to which *one of the students* is neither epistemically nor scopally specific. But of course it is equally possible for partitive indefinites to get interpretations that *are* scopally and epistemically specific. (In the case of (10) an epistemically specific interpretation is perhaps not very natural because of the epistemic modal *must*. But if we replace *must have* by *has*, such an interpretation is easily accessible.)

We conclude that the partitive–non-partitive distinction is orthogonal to the distinction between specificity and non-specificity, epistemic or otherwise.

---

[3]This doesn't by itself explain why *the one of the woman* is ungrammatical. We expect that this has to do with the fact that the preposition *of* is 'conceptually irreflexive': even though there are many different relations that *of* can be used to express — ownership, origin, part-whole and many more — , it is true of each such relation R that it is either impossible for R to hold between any entity x and itself; or in case there are x such that xRx, then this can only be because we are dealing with what is in some intuitive sense an 'exceptional contingency'. In the case of *the one of the woman* the relation between the referent of the embedding *NP* and that of the embedded *NP* would be that of identity; and that relation, being necessarily reflexive, would be incompatible with the general constraint that partitive *of* is subject to. For other constraints on partitive constructions see Jackendoff (1972) and Ladusaw (1982).)

[4]See also the original paper Enç (1991), where it is stated that partitivity is neither a necessary nor a sufficient condition for licensing the specificity marker (*a*) *certain*. (Cf. the discussion of example (59) on p. 19 in Enç (1991).)

Although partitivity is going to play no further part in what follows, we permit ourselves one further observation about the semantics and pragmatics of partitive indefinites. 'Partitive specificity' has received considerable attention in the literature. Notable in particular are proposals to derive scopal specificity from various construals of partitive specificity (representative references are the already mentioned Schwarzschild (2002), Portner and Yabushita (2002) or Breheny (2003), even if not all these authors refer explicitly to partitivity). The following comment on the general tenor of these proposals may help to make our own position clearer.

We have just seen that partitivity as such doesn't impose specificity constraints either in the sense of scopal or of epistemic specificity. So if there is such a thing as partitive specificity, that would have to be yet another kind of specificity. But can 'partitive specificity' be considered a species of specificity at all, on any natural interpretation of that term? On our understanding it cannot. In the final paragraphs of this subsection we explain why we think this.

One reason, we suspect, why there has been a temptation to regard partitivity as a species of specificity is that in many languages partitive indefinites share certain features with definites. For instance, in Turkish partitive indefinites in direct object position show the same Accusative morphology as definites, whereas non-partitive indefinites do not (see Enç (1991) and subsequent literature on Turkish Accusative marking). To the best of our knowledge a fully satisfactory explanation of this phenomenon is still outstanding. But the explanation may have to do with the fact we already observed that it is common for the embedded NPs of partitive NPs to be definites, which contribute an element of definiteness to the semantics and pragmatics of the partitive NPs that contain them.[5]

So definiteness enters into the semantics of partitives in some way. But that isn't saying very much. And in fact, proposals to the effect that there is a type of specificity — partitive specificity — which partitive indefinites exemplify claim more than just that its semantics involves definiteness somehow and somewhere). This is true in particular of those proposals that emphasize the presuppositional dimension of partitives. That the interpretation of partitive indefinites has presuppositional aspects to it seems true enough; in particular, someone who, like us, believes that definite noun phrases generally trigger 'identification presuppositions' (Heim (1991)) wouldn't contest this; for partitive indefinites (overtly or covertly) have definite NPs as embedded constituents and on the view just mentioned these import their presuppositions into the semantics of the partitive indefinite.

It is important to keep in mind, however, that these presuppositions only concern the determination of the *set within which* the semantic value of the partitive NP must be contained. And as we have seen, it is an essential feature of the semantics of indefinite partitives that this set does *not* determine the value of the NP completely, but leaves room for choice. So these presuppositions do precisely *not* play the role of identification presuppositions for the semantic value of the partitive indefinite as a whole.

On the other hand, the specific use of a partitive indefinite involves, like the specific uses of non-partitive indefinites, a felicity condition that some may want to see as a kind of presupposition too. The epistemically specific use of an indefinite, on our understanding of that term, presupposes that the speaker has identifying information for the entity that she uses the indefinite to talk about. But this is a presupposition in a different sense than the one that we take to be part of the dominant view of the role played by linguistic presuppositions in the use and interpretation of the expressions that trigger them. On that view linguistic presuppositions are always *interpreter*-oriented: it is the interpreter who should be able to resolve the presupposition. (Presuppositions are constraints on the speaker only in a mediated sense: the speaker must take care to choose her words in such a way that the interpreter can resolve the presuppositions they carry.)

If this is how presuppositions function, then the identifiability constraint that comes with the epistemically specific use of an indefinite is *not* a presupposition. It is an identifiability constraint that concerns the speaker; what the use of an indefinite (specific or otherwise) communicates to the interpreter is that *he* is *not* expected to be able to identify a referent; and that is true for partitive indefinites just as much as it is for non-partitive indefinites.

---

[5]As a matter of fact, the embedded NP of a partitive need not be a definite NP; it can also be an indefinite, as in *two of five students* (though only if that indefinite gets a specific interpretation). Such NPs appear to contradict what has become known as Jackendoff's Constraint (Jackendoff (1972) or Ladusaw (1982) — see also Footnote 4 on page 7), but, according to our judgment, they are acceptable.

One motivation for wanting to see partitive indefinites as similar to definites, and perhaps as more similar to them than to non-partitive indefinites, is that in a variety of languages partitive indefinites share morpho-syntactic properties with definites that they do not share with non-partitive indefinites. We are unsure about the explanation for this. It may be that such languages zero in on the circumstance that identification presuppositions are important for both definites and partitive indefinites, even if their effect is different: pinning down a unique referent in the case of a definite and imposing limits on the set within which the "referent" is located (but without reducing the set to a singleton) in the case of a partitive indefinite. But even if that is the explanation, this doesn't alter the fact that partitive indefinites are like other indefinites in being 'of indeterminate reference', and that they are — on our understanding of the term 'presupposition' — *not* presuppositional.

1.3   Indefinites in Dynamic Semantics

The formal background on which we rely in this paper is Discourse Representation Theory (DRT). More precisely, the background is provided by the extension MSDRT of DRT, sketched in Section 2, in which it is possible to represent complex mental states and analyse the meanings of simple and complex attitude reports as descriptions of mental states. MSDRT will then be used to articulate our analysis of epistemic specificity.

Something should be said about the role that DRT will play in all this, since that role is quite different from what can be found in other dynamic analyses of specificity (cf. Farkas (1996), Jayez and Tovena (2006), Brasoveanu and Farkas (2009) or Geurts (2010)). In fact, the use that we will be making of DRT here differs markedly from its deployment in the original DRT account of classical 'donkey' phenomena, exemplified in (11a).

(11)   a.   If Bill owns a donkey, then he beats it.
        b.   Bill owns a donkey. He beats it.

It should be stressed, however, that we are not simply discarding the old DRT treatment. We want to retain it for donkey sentences like (11a) (on their standard interpretation, according to which the indefinites of such sentences are not used specifically). It is just that this treatment will play no direct part in the analysis that we are going to propose for epistemic specificity. The analysis we propose here could even be combined with an E-type treatment of donkey pronouns (like *it* as it occurs in (11a) and (11b)).

To explain what we have in mind here we need to say a few things about the main differences between the DRT treatment of the examples (11a) and (11b) and the treatment they are given by E-type accounts, which have been seen by many as the principal competing approach toward donkey anaphora.[6] DRT and FCS[7] treat indefinite *NP*s like *a donkey* in (11a), (11b) as terms (and not as quantifying *NP*s, the standard view of indefinites from Frege up to and including Montague). The semantic contribution of indefinites is to introduce object identifiers into the discourse (as discourse referents in DRT, as file cards in FCS and as free variables according to some other formulations of the same basic idea), which can then be subsequently 'picked up' by anaphoric pronouns. The accounts are set up in such a way that they assign (11a) the truth conditions expressed by the predicate logic formula in (12a) and (11b) the truth conditions shown in (12b). Note that (12b) captures the *joint* truth conditions of the two sentences that together make up (11b) and that there is no straightforward way of reformulating (12b) equivalently as a conjunction of two logically independent closed formulas, the first of which captures the content of the first sentence of (11b) while the second formula captures the content of the second sentence.

(12)   a. $(\exists y)(donkey(y) \ \& \ own(b,y) \ \& \ beat(b,y))$

---

b. $own(b, d)$ & $beat(b, d)$

The E-type approach is quite different in spirit. It adheres to the traditional principle that indefinites are existentially quantifying noun phrases, and places the burden of the analysis on the treatment of the pronouns anaphoric to them. Donkey pronouns like *it* in (11a) and (11b) are semantically analysed via replacement by definite descriptions. For instance, on one version of the E-type approach the occurrences of *it* in (11a) and (11b) are to be replaced by the description *the donkey that Bill owns*. When this description is substituted for *it* in (11b) it induces an entailment or presupposition that the predicate $\lambda x.(donkey(x)$ & $own(b, x))$ has a unique satisfier. For (11b) this means that the sentence entails or presupposes that there is at most one donkey that Bill owns, while (11a) entails or presupposes that there is at most one donkey Bill owns in the event that there is at least one.

For (11a) this result is plainly undesirable and in more recent versions of the E-type approach this defect has been corrected by recasting the analysis within a situation-theoretic framework. (See Elbourne (2005) for details.) But the case of (11b) is different. Many speakers aver the intuition that a 'donkey discourse' like this is felicitous only if the speaker can uniquely identify a donkey that Bill owns and that she is talking about by using the *NP a donkey*. To the extent that this intuition is correct, it indicates that here it is the DRT/FCS analysis that is in need of correction. But the correction that is needed is, we now believe, quite simple: In order for (11b) to be felicitous the speaker must be making an epistemically specific use of the indefinite *a donkey*. (For an interpreter of (11b) the presence of *it* in the next sentence and its construal as anaphoric to *a donkey* is a clue that the indefinite must be given a specific interpretation.)

Let us summarise the discussion of the present subsection up to this point: Our use of DRT in this paper can be seen as a further elaboration of the principle that the indefinite *NP*s we have been looking at are to be analysed as terms. But the interpretation of those terms may vary: some get a specific, others a non-specific interpretation. Our notion of the specific use of indefinites also makes it possible to do justice to the intuition — which has been seen as favouring the E-type approach — that in discourses like (11b) the anaphoric link between pronoun and indefinite imposes some kind of uniqueness on the semantic value of the indefinite.

We conclude the present section by mentioning a point that we will take up again in Section 4. Once we commit ourselves to the position that the constraints on trans-sentential anaphora should be stated in terms of specificity we have to confront a further question: At what level of analysis is the mechanism that determines the choice between a specific and a non-specific interpretation supposed to operate? As a preliminary to the discussions of this issue in Section 4 we conclude our discussion of donkey anaphora with the two examples in (13).

(13) a. There is a donkey Bill owns. He beats it.
      b. There is a doctor in London. He is Welsh.
      (Sentence (13b) is a variant of an example from Evans (1980); it is discussed in Heim (1982).)

These examples show that the indefinite *NP*s of *there*-insertion constructions can be given a specific construal, just as this is possible for indefinites appearing in other syntactic configurations (such as e.g. *a donkey* in (11b)). This point is worth noting for the following reason. *There*-insertion constructions are highly selective with regard to the *NP*s which can appear in the position occupied by the indefinite *NP*s *a donkey* in (13a) and *a doctor* in (13b). For instance, definite singular *NP*s like proper names, pronouns or demonstratives are excluded from this position (Milsark (1977)). This indicates that the status of indefinites in *there*-insertion contexts is crucially different from that of referring terms. An indefinite in a *there*-insertion context behaves as an expression that introduces a variable that the *there is* constituent, which acts as a variable binder, needs as bindee.[8] But if this is what renders the combination of *there is* and indefinite grammatical, then the interpretation of the indefinites in (13) as specific, and thus as standing for some particular referent, must take place

---

[8]That *there*-insertion involves quantification was already hypothesised by Milsark in Milsark (1977). We present arguments and empirical evidence for this hypothesis in Bende-Farkas (2002) and Bende-Farkas and Kamp (2001).

at some other level of interpretation than the one at which *there is* and indefinite are represented as binder and bindee.[9]

2 Specificity and *de Re* Belief

In Section 1.1 we pointed out that a speaker can use the indefinite *NP a real estate agent* in (5) either non-specifically or specifically; her utterance may be backed merely by existential knowledge — that there is some person or other whom her husband is seeing and who works in the real estate business — or she may have some particular real estate agent in mind and believe that her husband is seeing that person.

(5)   My husband is seeing a real estate agent.

Uses of indefinites which are specific in this sense are extremely common. In fact, for many utterances of indefinites it is hard or impossible to imagine that they could have been used otherwise. One type of example of this we encountered in Section 1.3, where we noted that pronominal anaphora to an indefinite in a preceding sentence in sentence sequences like (11b) is felicitous only if the indefinite was used specifically.

In other cases it is the content of what is said and the circumstances in which it is said that suggest or imply that an indefinite must have been used specifically. An example is (14). Suppose Professor *A* says to her colleague *B* after a certain student approached *A* earlier that day with the question where she might be able to find *B*:

(14)   A student was looking for you this morning.

Then *A*'s use of *a student* will derive from her concept of the student who approached her, and the thought that induces her to make the statement will have been that that particular student was looking for *B*. Moreover, it is likely that *B* will, when *A* says (14) to him, interpret her words in this spirit: Presumably *A* is talking about some particular student who asked her where she could find me. In other words, *B* is likely to take *A* to have made a specific use of *a student*, and it is reasonable to assume that he will form, as the result of interpreting *A*'s utterance, a representation which reflects this.

It is these aspects of specificity — the link that exists in such cases between the use of the indefinite and an individuating concept in the mind of the speaker, and a second link that gets established in the mind of the interpreter through his assumption that such a link exists in the speaker's mind — which are the targets of the specificity analysis we develop in this and the next section.

In Section 1.2 we drew attention to the parallel between the specific–non-specific distinction and the *de re–de dicto* distinction: The difference between the specific and the non-specific use of *a real estate agent* in (5), we noted, resembles that between the *de re* and the *de dicto* interpretation of *a real estate agent* in the belief attribution (4).

(4)   Mary believes that her husband is seeing a real estate agent.

The extensive literature on the *de re–de dicto* distinction has led many — linguists as well as philosophers — to regard this distinction as an instance of genuine ambiguity. This view has been supported by the widespread use of formalisms that are able to express the distinction and that date back to the early days of formal semantics. For instance, if *believe* is analysed as a 2-place predicate *BEL* which relates individuals to propositions, and proposition-denoting terms are formed by attaching the intensional abstraction operator $^\wedge$ to formulas, then the two interpretations of (4) can be expressed as in (15).

(15)   a. $BEL(m, {}^\wedge(\exists y)(REA(y) \,\&\, SEE(Husb(m), y)))$
            (*de dicto* belief)

---

[9]In particular, our assumption about the nature of *there*-insertion is clearly incompatible with a straightforward explication of specificity of indefinites in *there*-insertion contexts along the lines of Fodor & Sag: obviously the indefinite cannot be an individual constant and a variable bound by *there is* at the same time.

b. $(\exists y)(REA(y) \;\&\; BEL(m, {}^{\wedge}SEE(Husb(m), y)))$
      (*de re* belief)

Quine stressed the very different commitments that these two forms make and drew attention to the problems connected with "quantifying into opaque contexts" (see e.g. Quine (1966a)), which is involved in (15b) but not in (15a). (15b) attributes to Mary a belief which is *de re* ('about the thing itself') with respect to the person that, according to Mary, her husband is seeing. (15b) represents this attribution in the form of a 3-place relation between (*i*) Mary, (*ii*) another individual (the real estate agent in question), and (*iii*) a property which Mary attributes to that individual. In contrast, (15a) simply attributes to Mary the existential proposition that there is some real estate agent her husband is seeing. Here a 2-place relation between believers and propositions suffices.

There has been much discussion over the question what conditions must be fulfilled in order that an agent can have a *de re* belief (or other attitude) with respect to some other person or thing. It has often been said that the agent must be "acquainted" with the entity at issue. Paradigmatic are cases where the subject is directly perceiving the entity, e.g. is looking at it under conditions where he can clearly see it. But it is difficult to know where to draw the line between cases in which the *de re* interpretation is legitimate and cases where it is not. Some philosophers have drawn the line extremely liberally, so that any information sufficient for unique individuation will be good enough (cf. Chisholm (1976)) Others have proposed to draw it more narrowly. But in any case, the debate over where the line should be drawn is not one that directly concerns the linguist. What is important for our present purposes is that at least in some situations agents are able to entertain thoughts that are *about* particular people, things, places, etc, and that there are linguistic forms which they can use in order to give expression to such thoughts.[10]

Our analysis of the *de re–de dicto* distinction is the similarly motivated as this more traditional analysis, if somewhat different in form. We assume that entity-related thoughts always involve a separate representation for the entity that the thought is about (i.e. the entity with respect to which the thought is *de re*), and moreover that this representation of the entity is *anchored*.[11] Such a theory requires a general formal framework for representing anchored and unanchored thoughts. The remainder of this section will be devoted to the presentation of that framework. However, and here our approach differs from the many discussions in the literature which deal only with the logic and semantics of attitude attributing sentences, we use the representations which this framework makes available not just for the semantic analysis of sentences which have the explicit form of attitude attributions, but also for the purpose of analysing the relationship between thoughts and the utterances to which they lead. It is one particular aspect of this relationship, viz. the relation between individuating entity concepts and (epistemically) specific uses of indefinites, that forms the central topic of this paper.[12]

The design of MSDRT was guided by two concerns that distinguish it from other formalisms for describing attitudes or attitude reports which were available at the time when it was first developed:

(*a*) The formalism should be able to represent not just single attitudes (or the contents of single attitude reports), but also complex mental states made up of several attitudes of distinct attitudinal "modes" (like, say, a belief and a desire) but whose contents may nevertheless be "referentially connected" in the sense that they target one and the same object. (And, similarly, it should allow for the representation of natural language reports of several attitudes that may be connected to each other in this sense, for instance through the use of anaphoric pronouns.) It will become clear

---

[10]Quine observed that belief sentences with 'exported existential quantifiers' unequivocally express *de re* attributions. For instance, *There is someone whom Mary believes her husband is seeing* only has the interpretation given in (15b) but not that in (15a). The focus in the present paper is on occurrences of indefinites in the complement clauses of attitude verbs — like that of *a real estate agent* in (4) — which allow for both a *de re* and a *de dicto* attribution.

[11]For more on anchors see Section 2.2. Further discussion can be found in Kamp (2003) and Kamp et al. (2011).

[12]We have been making use of this formalism in our own work since the second half of the nineties. (For one such application see e.g. Kamp (2006)). A detailed presentation of the formalism was written in 1996. So far only a (slightly abridged) German translation of this presentation has appeared in print (Kamp (2003)). A somewhat shorter version, but one that is like the longer versions in that it gives an explicit statement of the formalism's syntax and model-theoretic semantics, can be found in Kamp et al. (2011). For excellent discussions of a closely related formalism based on similar intuitions and with similar applications see the work of Maier, in particular (Maier (2015).

presently what we mean by "referentially connected".

(b) The formalism should provide us with the means to represent "singular thoughts" — thoughts which attribute a certain property to a certain object "directly", in the sense of, for instance, Kaplan (e.g. Almog et al. (1989)).

2.1 Internal Coreference

We start with concern (a), formulated at the end of the introduction to Section 2. Suppose a person who sees a gold coin in the middle of the road forms the desire to have it in her possession and the intention to go and pick it up, in order to satisfy this desire. In our formalism the thought complex consisting of these three attitudes — the belief, the desire and the intention — can be represented as in (16).

$$
(16) \quad \left\{
\begin{array}{l}
\left\langle BEL, \begin{array}{|c|} \hline x \; s_1 \; s_2 \\ \hline n \subseteq s_1 \quad n \subseteq s_2 \\ s_1 \colon gold\ coin(x) \\ s_2 \colon x\ be\ lying\ in\ front\ of\ i \\ \hline \end{array} \right\rangle \\[4em]
\left\langle DES, \begin{array}{|c|} \hline s_3 \\ \hline n \subseteq s_3 \\ s_3 \colon i\ have\ x \\ \hline \end{array} \right\rangle \\[3em]
\left\langle INT, \begin{array}{|c|} \hline t_4 \; e \\ \hline n < t_4 \quad e \subseteq t_4 \\ e \colon i\ pick\ up\ x \\ \hline \end{array} \right\rangle
\end{array}
\right\}
$$

Each of the components in (16) that represent these three above-mentioned attitudes consists of (i) a mode indicator and (ii) a DRS.[13] and the mode indicator indicates the kind of the represented attitude — whether it is a belief, desire, intention, doubt and so forth. The aspect of the mental state represented in (16) that the formalism is specially designed to capture is the "referential connection" between the three attitudes: the desire and the intention are aimed at the same object as the belief. Formally this is rendered by the reappearance of the discourse referent x, which was first introduced into the content representation of the belief, in the content representation of the desire and in that of the intention. And note well, the reusing of x in different attitude representations is designed to capture the internal, "psychological" nature of the coreferentiality links between the three attitudes: The subject whose state (16) describes *understands* her belief, desire and intention as concerned with the same object. It is this internal coreferentiality that plays an essential part in what drives her 'mental dynamics': Once the belief has been formed, the desire follows suit — *as* the desire to take possession of the very coin that the subject thinks she is seeing; and the belief and desire then jointly lead to the intention to go and pick up this coin.

These mental processes are independent of whether the perception that led to the belief (and hence to the desire and the intention) was grounded in reality or it was illusory — what the agent saw may have been a gold coin in actual fact, but it may also have been something else (the tinfoil top of a milk bottle perhaps, or a glittering piece of stone embedded in the pavement); or, worse even, the experience may have been a pseudo-perception, caused by some trick of the light. In this last case,

---

[13]We assume the basic notions of DRT — *Discourse Representation Structure* (*DRS*), *discourse referent, DRS-condition*, etc. — to be familiar. These can be found in several introductions to DRT, for instance Kamp et al. (2011), Geurts et al. (2015), Kamp and Reyle (1993) or van Eijck and Kamp (2011).

where there is no object that can be identified as the perceptual cause of the subject's perceptual experience, there is nothing in the world that x represents. But that doesn't affect how the subject will reason with the information represented in (16). For instance, it won't affect the transition from belief to desire and from belief and desire to intention. Nor does it affect the inferential and planning processes that will be set in motion by the combination of these three attitudes after all three have been put into place. (Not, at least, until the moment when the subject becomes aware of her mistake, e.g. when she goes to the middle of the road to pick up the coin and finds that there is nothing there.) All these mental processes are just as likely to happen when triggered by a radical misperception as they would be when triggered by a perception that is veridical.

Internal coreference of the kind that the recurrences of x in the desire and intention components of (16) are designed to account for is reminiscent of what is found in "standard" DRT, where discourse referents are often shared between the representations of successive sentences. But there is an important difference. In standard DRT the sharing of discourse referents is the device used to account for anaphoric relations between pronouns and indefinite antecedents (cf. Section 1.3). But whenever this happens, the second sentence is not assigned a semantics on its own, but is interpreted only in conjunction with the representation for the sentence that contains its antecedent. For a representational structure like (16) such a semantics, which does not assign any content to the desire and intention components by themselves, won't be good enough. We cannot rest content with assigning propositional content only to the merge of the three DRSs that serve as content representations for the three attitudes. For what kind of semantic content could be meaningfully attached to, say, the 'conjunction' of a belief and a desire? Obviously a semantics for such combinations of attitudes of different modes must assign a semantic content to each of them individually. Only such a semantics can be expected to provide us with a suitable basis for an account of the mental processes that are set in motion by such combinations of attitudes of different mode, with each attitude making its own particular contribution, by virtue of its content on the one hand and its mode on the other. (Two beliefs will tend to play different roles in mental processes when they differ in content, and a belief and a desire would not play the same role even if their content was the same.)

In fact, a semantics for complex representations like (16) must accomplish two things at once: On the one hand it must assign separate contents to the DRSs that serve as content representations for the different attitudes, while on the other it must do justice to the internal referential connections between them (those that are expressed through the sharing of discourse referents). The solution to this problem that we have proposed in earlier work (see Kamp (2003) and Kamp et al. (2011)) rests on the idea that whenever the content representations of two or more attitudes involve sharing of discourse referents, then one of the attitudes will always serve as presupposition for the others, in the sense that the latter will have a well-defined content only on the assumption that the former is true. (In cases where one of the attitudes is a belief while the others have different attitudinal modes, as in (16), it is the belief that will play this role of presupposition for the other attitudes. In (16) this asymmetry is indicated by the fact that the shared discourse referent x is "declared" in the DRS for the belief, in the sense that it belongs to the universe of that DRS). Such a conditional definition of the contents of the "referentially dependent" attitudes seems to be the best we can do. But for most purposes this is all we need.[14]

We will have a little more to say about the formal semantics for representations like (16) towards the end of this section (see Section 2.5). But first we need to address some other matters.

---

[14] The internal structure of mental states is most tangibly important in connection with various forms of reasoning in which propositional attitudes of different modes are combined to generate some further attitude or attitude complex, such as when beliefs and desires are combined in planning and the formation of intentions. We conjecture that the semantics alluded to provides the foundation required for such processes. More specifically it will enable us to distinguish between rational and non-rational instances of practical reasoning so long as the premises from which the reasoning proceeds include for each attitude with a referentially dependent content every other attitude on which that content depends. This is common and natural for reasoning processes that satisfy this constraint, and it can be argued that rational reasoning always satisfies it. If this is true then our semantics does provide a proper foundation for practical reasoning. We are aware that what we are saying in this footnote borders on the oracular. But there cannot be any question of going into greater detail about these matters here.

2.2 Singular Thought

The first of these matters alluded to at the end of Section 2.1 is directly connected with the second of the two concerns that the present formalism has been developed to address (see concern (b) on page 13). The formalism should be capable, we said, of representing certain attitudes as having singular propositional content. The device MSDRT has adopted to this end is that of an *anchored entity representation*. To explain what anchored entity representations are like and what they contribute to the mental state representations of which they are part, we begin by looking at one type of such representations, those whose anchors are based on visual perception. When an agent perceives a certain object — e.g. a bird she sees perched on the roof of the house opposite — then she will form a representation of this bird *as* the thing she is currently perceiving. We agree with those who take the contents of beliefs formed in this way to be singular propositions in which the perceiver attributes various properties to the thing or things perceived.[15] Such a propositional content is true in any world w (which can be the actual world w_0 or some other possible world) iff the object that is the actual cause of the perceptually grounded representation — the object, in other words, that is the perceptual cause in w_0 of the perception which produces the agent's perceptually anchored entity representation — has in w the property which the thought attributes to it.

In our formalism the representation of singular content is mediated by so-called *anchored entity representations*. Anchored entity representations are structures of the form shown in (17):

(17) $\langle[\text{ANCH}, x],\ K\rangle$

Here x is a discourse referent; "[ANCH, x]" serves as indicator that the representation as a whole plays the part of an anchored representation of the object that x stands for; the discourse referent x that immediately follows "ANCH" will stand for the entity represented by $\langle[\text{ANCH}, x],\ K\rangle$ in argument positions of other DRS conditions, and is known as the *distinguished discourse referent* of the entity representation; K is a DRS which describes the nature of the anchor. In the case we are discussing, where the anchor is one of visual perception, K will contain information to the effect that the agent is currently seeing the represented object, together with some information about where the object is spatially located in relation to her, as well as, normally, some information about what sort of thing it is. For instance — switching back to the case of (16) — the perception of the presumed gold coin in the middle of the road could give rise to an anchored representation of the form in (18).

(18) $\left\langle\ [ANCH, x]\ ,\ \begin{array}{|c|}\hline x\ s\ s'\ s'' \\ \hline \\ n \subseteq s \quad s{:}\,i\ see\ x \\ n \subseteq s' \quad s'{:}\,coin(x) \\ n \subseteq s'' \\ s''{:}\,1.5cm \leq diameter(x) \leq 2.5cm \\ \hline \end{array}\ \right\rangle$

The first two conditions of K in (18) state that the anchor is one of current perception.[16] The remaining two conditions give additional information about the represented object — that it is a coin and that its diameter is between 1.5 and 2.5 cm — and should be seen as information that is also constitutive of the agent's perceptual experience. (The choice of these conditions is obviously somewhat arbitrary; what conditions should be regarded as part of K and what conditions as parts of beliefs based on the agent's perception presumably depends on factors that are hard to analyse in any systematic way; but how such questions get resolved is of no direct importance for present purposes, so we let the matter rest.)

---

[15]Often such contents are multiply singular. For instance, in the case under consideration the agent presumably has a perceptually based representation for the roof as well as for the bird; the belief she forms on the strength of her perception is that the perceived bird is on the perceived roof. The content of this belief is doubly singular, with regard to the bird and with regard to the roof. There is in principle no upper bound to the number of anchored representations that can enter into the propositional content of a belief (or other mode of propositional thought).

[16]To be exact, it asserts that there exists a state s that holds at the current time n and that that state is to the effect that the agent (represented by the self-referring discourse referent i) is seeing the represented object.

The form of an anchored representation is to be seen as a witness to the agent's understanding of her representation as the reflection and result of some causal connection that links her to the represented object and that enables her to attribute properties "directly" to that object. There is a distinction to be drawn, however, between the agent having an entity representation of the form (17) and there actually *existing* an object whose perception produced the representation (in the manner the DRS K describes) and thereby qualifies as the object represented by the representation. As noted before, the agent who thinks she is seeing a gold coin may not only be wrong in the sense that what she sees isn't a gold coin, but also in the even more serious sense of there being no object of any sort that is causally responsible for her impression that she is seeing a coin. The agent will form a representation of the form (17) in such cases as well, but this representation radically misfires, in that there is no object of any kind in the actual world w_0 that is represented by it.

The DRSs of anchored representations, then, represent the agent's understanding of how her representations are causally connected to the entities they represent. When such a DRS is true to fact — i.e. if there is an entity that stands to the agent of the representation (and by implication to the representation itself) in the causal relation the DRS describes, then the representation does represent this entity, and in such a way that it can serve as a basis for singular thoughts. Specifically, a DRS containing the distinguished discourse referent of such an entity representation will represent a singular proposition, which is about the object represented by the entity representation. This is reflected by the possibility of using MSDRT to describe mental states like that of the agent from our gold coin example of Section 2.1 also as complexes in which entity representations are among the constituents. For instance, the relevant part of the mental state of our agent in the different gold coin scenarios that in Section 2.1 was described in the form shown in (16) can also be described as in (19).

(19)

$$
\left\{
\begin{array}{l}
\left\langle [ANCH, x],\;
\begin{array}{|c|}
\hline
x\ s_1\ s_2 \\
\hline
n \subseteq s_1 \quad n \subseteq s_2 \\
s_1\text{: } i\ see\ x \\
s_2\text{: } x\ be\ lying\ ifo\ i \\
\hline
\end{array}
\right\rangle \\[2em]
\left\langle BEL,\;
\begin{array}{|c|}
\hline
s_3 \\
\hline
s_3\text{:}gold\ coin(x) \\
\hline
\end{array}
\right\rangle \\[2em]
\left\langle DES,\;
\begin{array}{|c|}
\hline
s_4 \\
\hline
n \subseteq s_4 \\
s_4\text{: } i\ have\ x \\
\hline
\end{array}
\right\rangle \\[2em]
\left\langle INT,\;
\begin{array}{|c|}
\hline
t_5\ e \\
\hline
n < t_5\ e \subseteq t_5 \\
e\text{: } i\ pick\ up\ x \\
\hline
\end{array}
\right\rangle
\end{array}
\right\}
$$

The structure in (19) qualifies as the description of our agent's mental state in each of the three possible scenarios described in Section 2.1: perception of what is in fact a coin, misperception as a coin of something which is actually something else, and radical misperception, where the agent is under the impression that she is perceiving a gold coin, but there is no actual perception of any object. However, what content is represented by (19) depends on which scenario the agent is in. If her perception is veridical — when the object she perceives is in fact a coin — then that coin is the object represented by the entity representation of (19) and, with that, it is the value that is

assigned to the distinguished discourse referent x of that representation; and furthermore, by virtue of that assignment the content representations of the belief, desire and intention components of (19), which contain free occurrences of x, each determine a singular proposition that is about this value. The scenario in which the agent's perceptual experience is a case of radical misperception is entirely different in these respects. In this case the entity representation doesn't represent any object — it couldn't, since there isn't any object that has been perceived. And as a consequence, any propositional content representation that contains occurrences of x will fail to represent a well-defined proposition.

This leaves the middle case, where there is an object perceived, but one that is not of the kind that the agent takes it to be. When the object is misperceived, and consequently misrepresented in the resulting entity representation, can the entity representation be nonetheless regarded as representing the object — is misrepresenting a form of representing? — or is no object represented in this case? Perhaps the answer should depend on how seriously the object is misperceived and misrepresented. Mistaking a piece of tinfoil for a coin seems to be a less drastic misperception than misperceiving a tree trunk as a child, or a child as a tree trunk. Arguably we have a legitimate representation of the perceived object in cases of the first kind but possibly not when the miscategorisation is as extreme as it is in cases of the second kind. This is another issue where it is hard to know where to draw the line. But for our purposes in this paper it doesn't matter how the 'middle cases' are handled, so we will let this matter rest.[17]

That the contents represented by mental state descriptions like (19) depend on the status of their entity representations — do they represent an object or don't they? — should be plausible enough. But remember that this is a difference that is inaccessible to the agent herself. So it won't affect what reasoning processes her mental state will impel her to engage in and the 'rationality' of those processes cannot depend on it. This means that if we want to account for the rationality of those processes in terms of the semantics of representations like (19), then we will have to rely on a different notion of content than the one we have just described, which does depend on whether entity representations succeed in representing.

The notion of content we are after in this connection is a version of what in the literature on mental representation and propositional attitudes has been discussed as *narrow content*.[18] Narrow content is content that can be ascribed to a mental state or mental representation in isolation from any possible causal link of the state or representation. For us the central question in this connection is: What is the narrow content of an entity representation? And here is the answer we propose to this question: The narrow content of an entity representation $\langle [\text{ANCH}, x], K \rangle$ is just the content represented by the DRS K. Or, more fully: The narrow content of a mental state description like (19) is the content of the representation that we obtain by replacing each of its entity representations $\langle [\text{ANCH}, x], K \rangle$ by a representation of the belief that its internal anchor is true — that is, $\langle [\text{ANCH}, x], K \rangle$ is replaced by $\langle BEL, K \rangle$. In the case of (19) this procedure leads to (20).

---

[17]The distinction between veridical and non-veridical entity representations can also be expressed in terms of the distinction between *internal* and *external anchors*. The *internal anchor* of an anchored representation is that part of it which describes the causal relationship that the agent takes herself to stand in to the entity that is represented by her entity representation. Internal anchors are part of the 'psychology' of the representer. The *external anchor* of an anchored representation is the entity to which the agent stands in the causal relationship described by the internal anchor — provided, of course, that there is such an entity. The distinction between internal and external anchors enables us to recapture that between veridical and non-veridical representations as follows: If an anchored entity representation has an external anchor, then that external anchor is the entity that the representation represents; in this case the representation is veridical. If there is no external anchor, then the anchored representation does not represent anything and is non-veridical.

The notions of the internal and the external anchor of an anchored entity representation are convenient tools in the discussion of many of the relevant issues in this paper and we will often make use of them.

[18]See for instance **?** as one among many other references.

$$(20) \left\{ \begin{array}{l} \left\langle BEL, \boxed{\begin{array}{c} x\ s_1\ s_2 \\ \hline n \subseteq s_1 \quad n \subseteq s_2 \\ s_1\text{: } i\ see\ x \\ s_2\text{: } x\ be\ lying\ ifo\ i \end{array}} \right\rangle \\[3em] \left\langle BEL, \boxed{\begin{array}{c} s_3 \\ \hline \\ s_3\text{:}gold\ coin(x) \end{array}} \right\rangle \\[3em] \left\langle DES, \boxed{\begin{array}{c} s_4 \\ \hline n \subseteq s_4 \\ s_4\text{: } i\ have\ x \end{array}} \right\rangle \\[3em] \left\langle INT, \boxed{\begin{array}{c} t_5\ e \\ \hline n < t_5\ e \subseteq t_5 \\ e\text{: } i\ pick\ up\ x \end{array}} \right\rangle \end{array} \right\}$$

Note that (20) is a representation of the same type as (16), in that some of its attitudinal contents are referentially dependent on others. (To be precise, the contents of the second belief, the desire and the intention are all dependent on the content of the first belief, via the discourse referent x, which is declared in the representation of this belief (i.e. x belongs to the universe of the content DRS of this belief) and free in the representations of the three other content representations.) The model-theoretic semantics for (20) poses the same challenge that we pointed to in connection with (16) but the challenge can be met in the same way. (See in particular footnote 14.) Another point to observe about narrow content representations like (20) is that the beliefs derived from entity representations are existential: each such belief is to the effect that *there is* some object satisfying the conditions that its content DRS predicates of the entity representation's distinguished discourse referent x. That doesn't prevent x from recurring in some other, referentially dependent content representations. But it is this combination precisely — of x playing an existential role in the belief representation in which it is declared and yet having free occurrences in the content representations of other components of the given mental state — that is responsible for the special challenge.[19][20]

2.3 Representing Attitude Reports

Attitude reports are often defined as sentences used to attribute propositional attitudes to 'agents' (i.e. creatures suited for such ascriptions). Thus, given the perspective we have developed in this section, attitude reports are sentences that can be used to describe certain parts of the mental states of agents. But when we look at the way people use language to describe mental states (of others as well as their own), it is plain that the use of a single sentence to attribute a single attitude is only one out of the many different options that our languages make available for describing mental contents. As often as not our concern is to describe larger portions of mental states, consisting of several attitudes; and in such cases we usually need several sentences to accomplish this. Such attitude

---

[19]A representation equivalent to (20) can be obtained by merging its two belief components into one. We will make use of this option when reduction becomes relevant once more in Section 4.

[20]The existence belief that is made explicit in (20) is of course available to the agent whether the anchored representation is veridical (i.e. whether it does represent an entity) or not. That is, the belief could have been made an explicit part of (19) as well, by formally adding the representation of the existence belief in (20) to it. In general, every mental state representation with anchored entity representations can be considered equivalent to one that is obtained by adding representations of existence beliefs for all or any of its anchored entity representations.

reports, in other words, do not take the form of single sentences but of stretches of discourse.

Suppose we want to account for the semantics of single- and multi-sentence attitude reports within the framework of DRT and that we assume that the fragments of mental states attributed by such reports are to be represented in the formalism exemplified in (19). Then we will have to integrate this formalism into a DRT-based representation language (a language whose formulas are DRSs) which is also capable of handling sentences and pieces of discourse that serve other purposes than attitude attribution. The device that has been adopted in MSDRT for this purpose is the predicate Att. Att has three argument slots. Its first slot is for terms denoting the agents of the attitudinal states that Att is used to ascribe to those agents. The second argument position is for the kinds of representations that are exemplified by (16) and (19). The third argument slot is reserved for specifications of external anchors — what is meant by that will become clear below.

As discussed in Section 2.1, our DRS formalism represents both verbal and non-verbal predications as descriptions of events or states. We also adopt this representation format for predications involving the predicate Att. As discussed in Section 2.1, our DRS formalism represents both verbal and non-verbal predications as descriptions of events or states. We also adopt this representation format for predications involving the predicate Att. Since Att-predications provide descriptions of mental states, they should be seen as descriptions of states (and not of events). So the schematic form of an Att-predication should be as in (21).[21]

(21) $s: Att(\xi, \mathcal{K}, \ldots)$, where
        (i) ξ is a discourse referent (representing the attributee)
        (ii) $\mathcal{K}$ is a structure of the kind exemplified by (16) and (19).

Note well that the state (represented by the discourse referent) s is a state which consists in the agent ξ being in a mental state of the kind described by $\mathcal{K}$ (together with some further information that is contributed by the third slot, explained below). s is thus not ξ's mental state itself, but the state of ξ being in a certain kind of mental state.

(21) doesn't show what goes into the third argument slot of Att. To see what this slot is for consider the DRS in (23) below, the representation of a complex attitude attribution made by an agent Bernard to an attributee Rachel. This attribution representation ascribes to Rachel a past mental state which contains an entity representation and three propositional attitude representations. The last three of these are like the components of (19). But there are nevertheless important differences between (19) and (23). (19) is the representation of a certain kind of mental state. In contrast, (23) can serve two purposes, neither of which is a mere mental state representation. The first purpose is to serve as the representation of a complex attitude attribution that one agent makes in thought to some other agent. For example, (23) could be taken as representation of what Bernard thinks at a given point in time about what Rachel thought and wanted at some earlier time. But structures like (23) can also serve as representations of attitude reports. For instance, (23) is a possible semantic representation for the two-sentence report (22).

(22)   Rachel thought a gold coin was lying in the middle of the road.
        She wanted to have it and intended to go and pick it up.

One of the general aims of DRT is to show in detail how semantic representations of linguistic input can be constructed according to systematic principles from syntactic structures for the input sentence or sentences. To carry out the actual construction of (23) from syntactic structures for the sentences of (22) transcends the aims of this paper. (For explicit formulations of DRS construction algorithms see Kamp and Reyle (1993) or the as yet unpublished Kamp (2017).)

But although we won't present an explicit construction of (23) from (22), we will nevertheless

---

[21]Strictly speaking the state discourse referent preceding the colon in (21) is also an argument of the predication expressed by (21). On this view Att is a 4-place rather than a 3-place predicate. But for many of those to whom we have presented this material talking about Att as 4-place, with the state it describes as its first argument and the agent and mental state description in second and third argument position, seems to go against the grain, just as it goes against the grain for many to talk about a predicate like *child* as a 2-place predicate, with a first argument slot for states and a second argument for persons. To comply with these sensibilities we will go on referring to Att as being 3-place.

have some things to say about the construction process, which we believe will be understandable also in the absence of a fully explicit presentation of the construction.

$$(23)$$

$$s_0\colon Att\Big(r, \Big\{ \Big\langle [ANCH, x], \begin{array}{|c|} \hline x \quad s_1 \quad s_2 \\ \hline n \subseteq s_1 \quad n \subseteq s_2 \\ s_1\colon i\ see\ x \\ s_2\colon x\ be\ lying\ ifo\ i \\ \hline \end{array} \Big\rangle, \Big\langle BEL, \begin{array}{|c|} \hline s_3 \\ \hline n \subseteq s_3 \\ s_3\colon gold \quad coin(x) \\ \hline \end{array} \Big\rangle, \Big\langle DES, \begin{array}{|c|} \hline s_4 \\ \hline n \subseteq s_4 \\ s_4\colon i\ have\ x \\ \hline \end{array} \Big\rangle, \Big\langle INT, \begin{array}{|c|} \hline t_5 \quad e \\ \hline n < t_5 \quad e \subseteq t_5 \\ e\colon i\ pick\ up\ x \\ \hline \end{array} \Big\rangle \Big\}, \{\langle x, x' \rangle\} \Big)$$

with outer universe $t_0 \quad s_0 \quad r \quad x'$ and conditions $t_0 < n \quad t_0 \subseteq s_0 \quad Rachel(r)$.

What matters at this point is the semantics for structures like (23) and the contributions that their various constituents make to it. We focus on those aspects of (23) of which we assume that they need explanation also for readers who have a basic knowledge of DRT, but are unfamiliar with MSDRT.

First, note that (23) contributes to Rachel a mental state — described by the *Att*-condition — at a time t_0 which precedes the attribution time represented by the outer occurrence of n. ((23) should be taken as the semantic representation of a certain utterance of (22). The occurrence of n outside the *Att*-condition represents the time of this utterance.) This temporal relation is expressed by the conditions "$t_0 < n$" and "$t_0 \subseteq s_0$", which together convey that the state s_0 — the state to the effect that Rachel was a mental state of the kind described by the *Att*-condition — held at a time that preceded the attribution time represented by the outside occurrence of n. By means of such temporal DRS conditions mental states ascribed by *Att*-conditions can be located in relation to any times that are available within the larger representation in which the *Att*-condition is embedded.

Second, the semantics of the *Att*-condition itself. What this condition says is that at the time t_0 of the state s_0 Rachel's mental state was one that contained (among many others no doubt) components of the kind described by the four constituents of the second argument of *Att*: (i) an entity representation and (ii) three propositional attitudes, a belief, a desire and an intention, that are all referentially connected to the entity representation. We consider these constituents in turn.

(i) The entity representation. According to the *Att*-condition the entity representation represents some object x' to which Rachel was related through visual perception at the time t_0 when she was in the mental state attributed to her. (That the visual perception took place at the time of the attributed mental state is expressed by the condition $n \subseteq s_1$ that is part of the internal anchor for the entity representation. This is what this last condition expresses, because the occurrences of n that are internal to the *Att*-condition of (23) identify the time at which the attributee Rachel

entertains the mental state she is said to be in. Such *Att*-internal occurrences of n represent the attributee's "psychological present" at the time of she was in the mental state attributed to her. Unlike the *Att*-external occurrences of n they do not represent the time of the attribution. For details see Kamp (2011).)[22]

Just as the *Att*–internal occurrences of n represent the attributee's psychological present, the discourse referent i represents her "psychological self" — that concept of self that is involved in self-attribution and that must be a constituent of any thought that the agent herself can express by an utterance that contains one or more occurrences of the pronoun *I*.

(23) not only claims that at t_0 Rachel was in a mental state that contained an entity representation with the properties that are displayed in the second argument slot of its *Att*-condition; it also claims, through the third argument of *Att*, that this entity representation does represent an actual object. This object is represented — "from the outside", so to speak — by the discourse referent x'. This is the import of the presence of the pair $\langle x, x' \rangle$ in the third argument slot of *Att*, where x is the distinguished discourse referent of the entity representation that is being attributed to Rachel. (Note that x' is 'declared' outside the *Att*-condition: it is declared in the universe of the main DRS of (23). That makes it into an external representative of the entity represented by Rachel's entity representation.) The upshot of the presence of $\langle x, x' \rangle$ in *Att*'s third argument slot and the declaration of x' in the main DRS universe is that (23) represents that interpretation of (22) according to which Bernard attributes to Rachel a belief, desire and intention that are *de re* but where the *res* need not be known to him. According to the standard DRS semantics the presence of x' in the main universe of (23) signifies that according to the interpretation of (22) that (23) represents, there is some *res* that Rachel's attitudes are about, but that is all that the *de re* interpretation of (22) has to tell us about the *res*.

(ii) Next we turn to the propositional attitudes that the *Att*-condition of (23) asserts to have been part of Rachel's mental state at t_0. Most of what can and needs to be said at this point about these constituents of Rachel's mental state has been said already. But there is one, general, point about mental state ascriptions expressed by *Att*-conditions that deserves attention. The representations that fill the second argument slot of *Att* are *partial* descriptions of the attributed mental states, and in fact they are partial in two respects, or at two levels if you like: First, the mental state that is being described may contain many more components than the entity and propositional attitude representations that are mentioned in the description. (In practice, all the descriptions occurring in *Att*-conditions are partial in this sense.) Secondly, the content representations of the propositional attitude representations that the description does mention explicitly will usually be partial characterizations of the contents of the represented attitudes. For instance, Rachel's belief represented in (23) as (24) may have a stronger content — i.e. impose stronger conditions on the object represented by x — than what is expressed by the DRS of (24).
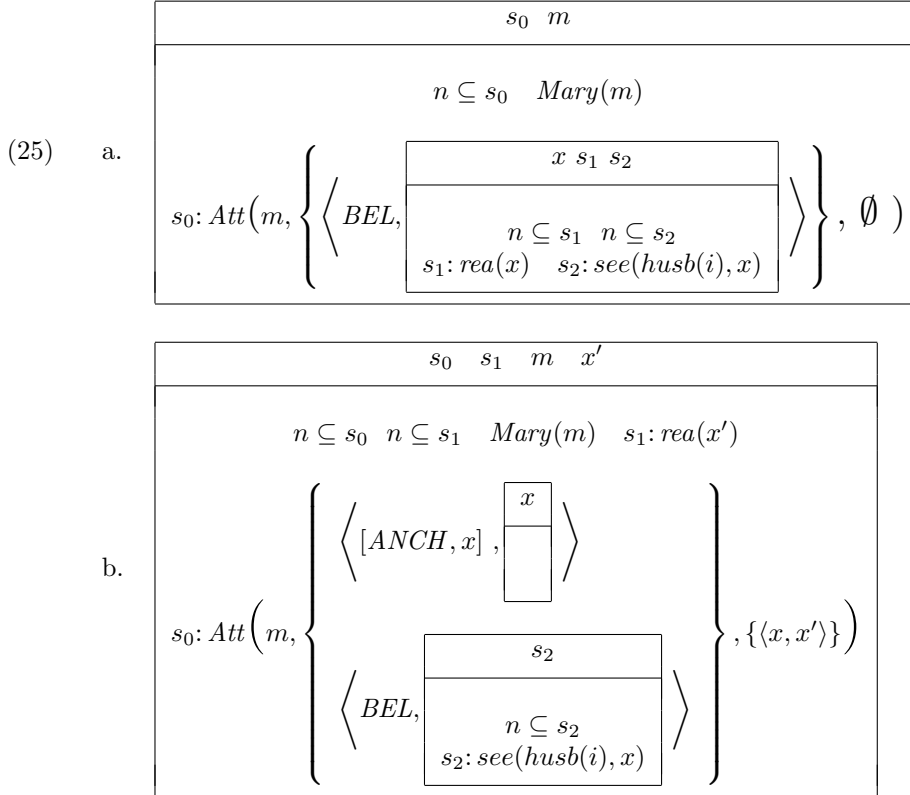
$$
(24) \qquad \left\langle BEL, \boxed{\begin{array}{c} s_3 \\ \hline \\ n \subseteq s_3 \\ s_3{:}gold \quad coin(x) \end{array}} \right\rangle
$$

We noted that (23), in its role of semantic representation for (22), is neutral with regard to the question whether the attributor Bernard has his own means of identifying the object that according to (23) is the joint target of Rachel's belief, desire and intention. We want to dwell a little longer on this point and return in this connection to what we said in Section 2 about possible interpretations of (4).

(4)   Mary believes that her husband is seeing a real estate agent.

---

[22]Different occurences of n within the same DRS have different import depending on whether they occur inside or outside *Att*-conditions. This dependence may be somewhat confusing on first encounter, but it is an essential part of the mechanisms used in DRT and MSDRT to represent detailed temporal information. (For instance, as (23) indicates, it is part of what enables MSDRT to represent that A has a thought (or makes a statement) at t_0 about a thought of B at some earlier time t_1, that is about what happened at an even earlier time t_2.)

We distinguished three different interpretations for (4), its *de dicto* interpretations and two *de re* interpretations, one in which the indefinite *a real estate agent* is taken as *de re* in relation to the attributee but nonspecific in relation to the attributor, and one in which the indefinite is taken as both *de re* and specific. Using DRSs of the form exemplified by (23) we can represent the first two of these interpretations as shown in (25a) – (25b).

(25)   a.

$$s_0\colon Att\Big(m, \Big\{ \Big\langle BEL, \boxed{\begin{array}{c} x\ s_1\ s_2 \\ \hline n \subseteq s_1 \quad n \subseteq s_2 \\ s_1\colon rea(x) \quad s_2\colon see(husb(i), x) \end{array}} \Big\rangle \Big\}, \emptyset\ \Big)$$

with outer box header $s_0\ m$ and conditions $n \subseteq s_0 \quad Mary(m)$

b.

$$s_0\colon Att\Big(m, \Big\{ \Big\langle [ANCH, x], \boxed{\begin{array}{c} x \\ \hline \\ \end{array}} \Big\rangle, \Big\langle BEL, \boxed{\begin{array}{c} s_2 \\ \hline n \subseteq s_2 \\ s_2\colon see(husb(i), x) \end{array}} \Big\rangle \Big\}, \{\langle x, x'\rangle\}\Big)$$

with outer box header $s_0\ \ s_1\ \ m\ \ x'$ and conditions $n \subseteq s_0 \quad n \subseteq s_1 \quad Mary(m) \quad s_1\colon rea(x')$
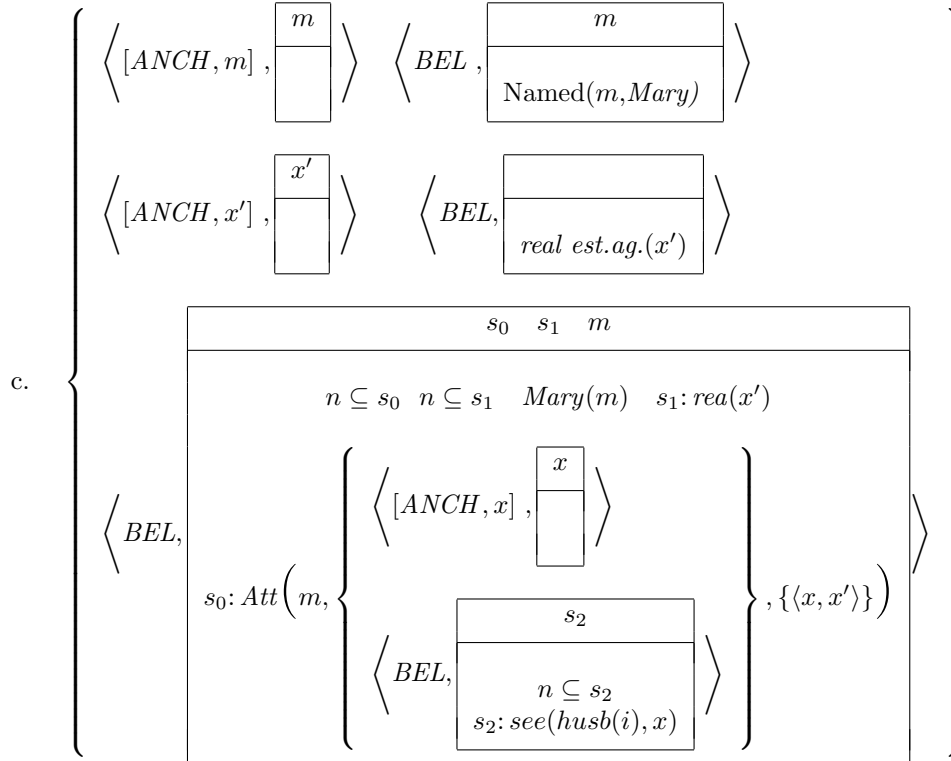
But the distinction between the second and the third interpretation of (4) cannot be captured at this representational level. (25b) is a semantic representation of Bernard's utterance of (4), and not a representation of Bernard's mental state. To represent the different thoughts that may have led Bernard to produce (4) we have to recast (25a) and (25b) as components of the mental state he was in at the time of his utterance. Presumably the thought that led him to make his utterance was a belief. On this assumption we can, using the format for mental state descriptions of (16), (19) and (20), and making the plausible assumption that Bernard has an anchored entity representation for Mary, represent the three mental states that might have led him to utter (4) as shown in (26a), (26b) and (26c). In case Bernard's mental state can be described as in (26a), his own interpretation of his utterance of (4) would be the one represented in (25a). If his mental state was either (26b) or (26c), then his interpretation of his own words would have been the one given by (25b).

(26) a.

$$\left\{ \left\langle [ANCH, m] , \boxed{\dfrac{m}{}} \right\rangle \quad \left\langle BEL , \begin{array}{|c|}\hline m \\ \hline \text{Named}(m, Mary) \\ \hline \end{array} \right\rangle \right.$$

$$\left. \left\langle BEL, \begin{array}{|c|}\hline s_0 \\ \hline n \subseteq s_0 \\ s_0 \colon Att\Big(m, \Big\{ \Big\langle BEL, \begin{array}{|c|}\hline x \ s_1 \ s_2 \\ \hline n \subseteq s_1 \quad n \subseteq s_2 \\ s_1 \colon rea(x) \quad s_2 \colon see(husb(i), x) \\ \hline \end{array} \Big\rangle \Big\}, \ \emptyset \ \Big) \\ \hline \end{array} \right\rangle \right\}$$

b.

$$\left\{ \left\langle [ANCH, m] , \boxed{\dfrac{m}{}} \right\rangle \quad \left\langle BEL , \begin{array}{|c|}\hline m \\ \hline \text{Named}(m, Mary) \\ \hline \end{array} \right\rangle \right.$$

$$\left. \left\langle BEL, \begin{array}{|c|}\hline s_0 \quad s_1 \quad m \quad x' \\ \hline n \subseteq s_0 \quad n \subseteq s_1 \quad Mary(m) \quad s_1 \colon rea(x') \\ s_0 \colon Att\Big(m, \Big\{ \Big\langle [ANCH, x] , \boxed{\dfrac{x}{}} \Big\rangle \quad \Big\langle BEL, \begin{array}{|c|}\hline s_2 \\ \hline n \subseteq s_2 \\ s_2 \colon see(husb(i), x) \\ \hline \end{array} \Big\rangle \Big\}, \ \{ \langle x, x' \rangle \} \Big) \\ \hline \end{array} \right\rangle \right\}$$

23

c.

$$\left\{ \begin{array}{l} \left\langle [ANCH,m]\,,\ \boxed{\begin{array}{c} m \\ \hline \\ \end{array}} \right\rangle \quad \left\langle BEL\,,\ \boxed{\begin{array}{c} m \\ \hline \text{Named}(m,Mary) \end{array}} \right\rangle \\[2em] \left\langle [ANCH,x']\,,\ \boxed{\begin{array}{c} x' \\ \hline \\ \end{array}} \right\rangle \quad \left\langle BEL,\ \boxed{\begin{array}{c} \\ \hline real\ est.ag.(x') \end{array}} \right\rangle \\[2em] \left\langle BEL,\ \boxed{\begin{array}{c} s_0 \quad s_1 \quad m \\ \hline n \subseteq s_0 \quad n \subseteq s_1 \quad Mary(m) \quad s_1\!:\!rea(x') \\ s_0\!:\!Att\Big(m,\ \left\{ \begin{array}{l} \left\langle [ANCH,x]\,,\ \boxed{\begin{array}{c} x \\ \hline \\ \end{array}} \right\rangle \\[1.5em] \left\langle BEL,\ \boxed{\begin{array}{c} s_2 \\ \hline n \subseteq s_2 \\ s_2\!:\!see(husb(i),x) \end{array}} \right\rangle \end{array} \right\}\,,\{\langle x,x'\rangle\}\Big) \end{array}} \right\rangle \end{array} \right\}$$

The representation level exemplified by (26a), (26b) and (26c) is not the end of all roads. Each of these mental state descriptions can serve in its turn as second argument in an *Att*-condition that attributes a mental state to some third agent. In this way it is possible to represent the thoughts that the recipient of (4) attributes to the speaker Bernard. For instance, if the recipient interpreted Bernard's utterance as the *de re* attribution represented in (25b), then he may either infer that Bernard's mental state is of the type represented in (26b) or that it is of the type represented in (26c). Alternatively, if the recipient interprets (4) as the *de dicto* attribution represented in (25a), then he will be more likely to conclude that Bernard's mental state is of the type shown in (26a). These various attributions by the recipient to Bernard can be represented explicitly by *Att*-conditions in which the second slot of *Att* contains a condition of the form $< BEL, K >$, where K is one of the representations in (26).

Along these lines our formalism is able to express a wide range of iterated attitude attributions in thought and language, of any embedding depth.

Much of what we have said about the interpretation of (4) also applies to the use and interpretation of (22). As noted, (23) is one of its possible utterance interpretations. But (23) leaves open whether or not Bernard, the speaker of (22), has himself an entity representation that identifies the object that Rachel has her belief, desire and intention about according to what (23) attributes to her. In this case too our formalism is capable of distinguishing between these two different mental states of the speaker, one in which he himself can identify the object and one in which he cannot; and here too the formalism can represent the different assumptions that that the recipient of (22) can make about the thoughts of the speaker; and so on.

It is also worth observing that although the attitude report (22) extends over several clauses, it too allows for a *de dicto* interpretation of the indefinite *a gold coin in the middle of the road* and not only for a *de re*-interpretation. At the level of utterance interpretation this leads to a representation that stands to (16) as (23) stands to (19). For reasons of space we do not display the semantic representation for this *de dicto* interpretation. But we trust that by now it will be clear enough what such a representation should look like in the formalism we have outlined.[23]

---

[23]Iterated nesting of *Att*-predications is also the tool for defining various notions of common knowledge and shared

2.4   Anchors that are not Based on Current Perception

So far we have focused exclusively on perceptually anchored entity representations, representations whose anchors say of the represented object that it is being currently perceived. But these are not the only possible kind of anchor. A first, modest, departure from the case of an anchor based on current perception are anchors based on memories of perceptions. The departure seems very modest indeed when the remembered perception is very recent. As you turn your gaze from the coin in the middle of the road to the left or the right (to see if there is any traffic approaching) your perception of the coin is, strictly speaking, no longer current. But it would seem highly artificial to take this as a ground for claiming that the entity representation you still have of the coin is no longer anchored (and therefore no longer enables you to have singular thoughts about the coin). If current perception of an object entitles us to entertain singular thoughts about it, then surely a memory of such a perception that is as fresh as this ought to entitle us as well. When the perception lies farther back in time, and its memory becomes hazier and less detailed, the question whether the entity representation still qualifies as anchored may become harder to answer on intuitive grounds. We ourselves are inclined to accept even memories of distant perceptions as a basis for (internally and externally) anchored representations, but others may well differ on this point.

Extrapolation from the paradigm of current direct perception is also possible in another direction. When you see footsteps in the sand, that will produce in you an anchored representation of the footsteps — this much follows from what we have said. But what about the person who left them? Assuming that these are the footsteps of a single person, then that person was the productive cause of those footsteps. And when this relation is combined with the causal relation involved in the perception, we get a causal chain which connects you with that person. Does that entitle you to an anchored representation of that person, which enables you to entertain singular thoughts about her? And if the answer to this first question is supposed to be yes, what about that person's mother, or her mother's mother? Such questions — questions about which causal connections suffice for the possibility of entertaining singular thoughts about the entities at the other end of the connection — have been raised in the philosophical literature ever since the days when this issue became important, often in the context of the distinction between *de dicto* and *de re*: Under what conditions are *de re* attributions legitimate? As far as we can tell, there is still no general agreement what the answer should be.[24]

On the whole our own inclination is towards a fairly liberal position, and that in two different respects. First, any causal relation that leads back from an object d for which an agent has an anchored representation to some other object d' that is uniquely identified by standing in this relation to d, entitles the agent to an anchored representation of d'. Second, whether an agent forms an anchored or an unanchored representation of such an object will depend in part on the agent's own intentions: She can form an anchored representation of the object. But she may also form an unanchored representation, in which the causal relation is part of the descriptive conditions that identify the object as their unique satisfier. For instance, the observer of the footsteps in the sand may, in addition to an anchored representation of the footsteps, form an anchored representation of the person who left them, but she may also form a representation of that person *as* the one that uniquely satisfies the condition of having left the given footsteps, for which she does have an anchored representation.

---

belief (Lewis (1969), Schiffer (1972) among many others). We will have no need of such notions in the remainder of this paper. So we leave it at the general observation that the inferences that interpreters make about the mental states of those whose utterances they interpret form the bottom layer of the hierarchically structured, recursively defined concepts that are central to much of the literature on the Common Ground. For more on forms of common knowledge, shared assumptions and common ground see Kamp (2016).

[24]Any sharp dividing line between causal relations that justify anchored entity representations and those that do not will be hard to find or defend. For a start, few people seem to be able to have clear intuitions that provide even a rough guidance to where the line should be drawn. Moreover — though obviously this is connected to the first difficulty — any proposal will be under serious threat from slippery slope effects: Starting from causal relations on the positive side of the line (the side of those relations which do support anchored representations), we modify the scenario step by step, obtaining a series of causal relations the last of which had been previously agreed to lie on the negative side. Our own position, outlined in the next paragraph, is motivated in part by our awareness that some kind of response to this difficulty is needed.

There is a class of relations which we have not mentioned so far, but of which we want to suggest that they can support anchored representations as well. These relations are important in connection with the interpretations that recipients of communication compute for the utterances that come their way, and with the propositional contents determined by those interpretations. And, more specifically, they are crucial to the account of epistemically specific indefinites that we will present in the next section.

The relations in question are causal connections that are grounded in verbal communication. They link the one who interprets what he is being told by someone else (or who reads a text or is listening to the radio, etc.) to the person talking to him (or to the author of the text or the one who has produced what is coming out of the radio) and thereby to anchored representations that that person may be assumed to have had for the entities referred to by the words she used or uses. For an example of this consider an utterance involving a proper name, as in (27).

(27)   Last night I got a call from my old friend Andrea Alfieri.

Suppose that the recipient is unfamiliar with the name Andrea Alfieri. In this case he is likely to assume that there is some particular person with the name the speaker has used, and he will also be likely to infer that the speaker must have an anchored representation for that bearer of the name. Under those conditions, we maintain, the interpreter is in a position to form an anchored representation of the person to whom the speaker has been referring, and we suggest that that is what typically happens. The anchor of this representation links it to whoever the external anchor of the entity representation of the speaker is, the anchor that was instrumental in her using the name to refer to that individual. In other words, the interpreter's representation represents its external anchor as the external anchor of the anchored representation that the interpreter assumes the speaker to have. Such anchors, which ride piggyback on the anchored representations of others, are called *vicarious anchors*. We will present a proposal for the form of vicarious anchors below in Section 3.

When the interpreter's assumptions that are implicit in the adoption of a vicariously anchored representation are correct, then this entity representation does have an external anchor, and that anchor is the same as the external anchor of the entity representation that the speaker relied on when she used the name. If the interpreter's assumptions about the speaker are correct, then the semantic representation he will construct of the speaker's utterance (of which the vicariously anchored entity representation will be an essential constituent) will express the same singular proposition that the speaker herself associates with her words: a proposition about the person that the speaker's use of the name referred to.

The case where the name a speaker uses is new for the interlocutor is of course just one possibility, and it is not the most common one. More often the name a speaker uses refers to an individual that is already familiar to the interlocutor under that name. In such cases the interpreter will already have an entity representation for the referent, which he will access when interpreting the speaker's present use of the name, and as a rule this earlier representation will be an anchored representation. The interpreter may then still add further information to that representation which captures his assumption that it represents the entity to which the speaker has referred by way of the name, thus adding in essence a new internal anchor to an anchored representation that is already in place. (The structure of entity representations that result from such additions involves some further complications, but since these are not directly relevant to the goals of this paper, we will not dwell on these.)

It is this process, we suggest, in which utterance recipients form vicariously anchored entity representations for the bearers of novel names, that allows a name to spread from one member of a speech community to the next, while preserving its ability to serve as a rigid designator of its bearer. By making the notion of vicariously anchored representation explicit, the present theory allows us to formulate a version of the 'causal chain theory' of proper names due originally to Chastain and Kripke (Chastain (1975), Kripke (1977/1990)). Interpretations of other types of definite *NP*s, definite descriptions and complex demonstratives, can also lead to vicariously anchored entity representations. And the list is not limited to definite *NP*s. Indefinite *NP*s can give rise to such representations as well. This is, we claim, what happens when the recipient of an utterance

containing an indefinite *NP* assigns an epistemically specific interpretation to that *NP*. Or, more accurately, interpreting a given use of an indefinite *NP* as an instance of epistemic specificity (as we understand the term) *is tantamount to* adopting a vicariously anchored representation for the entity that the recipient assumes the speaker had in mind when she used the *NP* to talk about it.

2.5 Sketch of the Model Theory

Before we leave our sketch of the formalism we have to make good on our promise to say something about its model-theoretic semantics. As we noted earlier, the main difficulty with that semantics is to account for the individual contents of attitudes like those represented in (16) and (20) while doing justice to the internal referential connections that often hold between such attitudes. We already noted that the key to the solution for this problem that is presented in earlier work on MSDRT is that the contents of some members of a set of referentially connected attitudes can be treated as presuppositions for the content assignments to others. The content complexes that reflect such presuppositional dependencies go by the name of *Information State Based Attitude Structures* or *ISBASs*. ISBASs are defined in terms of the central notions of Dynamic Semantics, *information state* and *context change potential*.[25] Like propositions, information states and context change potentials ISBASs are *intensional* constructs. This means that the models for our MSDRT formalism must be *intensional models*.

On the other hand the models we need must also satisfy a number of other requirements that have to do with the DRS language of which the MSDRT formalism we are using is an extension. This DRS language is in essence that of Kamp and Reyle (1993), Ch. 5. The models described there are extensional models whose ontology includes besides a general category of individuals also special categories of times, events and states. (For details see Kamp and Reyle (1993).) Intuitively speaking, these models track the history of some possible world w through time, specifying for each moment t of the time structure of w what is the case or what is happening at t. We already noted that the models for MSDRT cannot simply be extensional models, but that they must contain the extensional DRT models of Kamp and Reyle (1993) as substructures. One way to achieve this (the standard way) is to assume that each model M determines a set of possible worlds W_M and that M associates with each world w from W_M an extensional model M_w of the Kamp and Reyle (1993) type.[26]

Special provisions need to be made for the evaluation of *Att*-conditions (cf. (21)). This is done as follows. We distinguish a special category of individuals, 'cognitive agents', who are capable of entertaining propositional attitudes. For each world w of the world set W_M of a model M, the model M_w assigns for each time t of its time structure, to each agent a that exists and is conscious in w at t a mental state $MS_w(a, t)$. This mental state is assumed to take the form of an ISBAS.

To explain how the function MS is used in the evaluation of *Att*-conditions is not all that simple, and for a detailed account we must refer to other work. (See in particular Kamp (2003).) Here we can only indicate how things work by looking at a couple of exceptionally simple cases.

Let $s: Att(\xi, \mathcal{K}, EA)$ be some particular *Att*-condition. We saw that such conditions allow for truth-conditional evaluation only when the third argument EA of *Att* provides external anchors for each of the anchored entity representations in $\mathcal{K}$. (As argued in Section 2.2, *Att*-conditions in which the third argument does not cover all the anchored entity representations resist truth-conditional evaluation; it is then only their narrow content reductions, in which each entity representation is replaced by the corresponding belief in the truth of its internal anchor, for which evaluation is possible.) So let us assume that EA provides external anchors for all the entity representations in $\mathcal{K}$.

What is it for an assignment function f, which maps the free discourse referents of $s: Att(\xi, \mathcal{K}, EA)$, to entities from a given model M, to verify the condition $s: Att(\xi, \mathcal{K}, EA)$ in M in a world w_M of M at a time t of M_w? This question will only arise when $s: Att(\xi, \mathcal{K}, EA)$ occurs as a member of the condition set of some DRS K_0, so let us assume that $s: Att(\xi, \mathcal{K}, EA)$ belongs to the condition set of

---

[25] For an authoritative presentation of the core of Dynamic Semantics see Groenendijk et al. (1996).

[26] For intensional models of this kind see Kamp et al. (2011).

K_0. For simplicity let us further assume that K_0 is proper.[27] (This is convenient and it constitutes no real restriction.) Then any assignment function f for which the evaluation of $s\colon Att(\xi, \mathcal{K}, EA)$ will arise (as part of the evaluation of K_0) will assign values to all the discourse referents that have free occurrences in $s\colon Att(\xi, \mathcal{K}, EA)$. These include the discourse referents that represent external anchors in EA (i.e. the second members x' of the pairs $< x, x' >$ in EA). The formal significance of these pairs $< x, x' >$ is that the discourse referent x represents what from an external perspective is represented by x'. This means that the value which f assigns to x' must also be the value of x. Formally this is captured by a stipulation in the verification conditions for $s\colon Att(\xi, \mathcal{K}, EA)$, accroding to which f is extended to an assignment function f' which is defined also for the first members x of pairs $< x, x' >$ in EA, and, for each such pair, f' assigns to x what f assigns to x'.

In order that f verify $s\colon Att(\xi, \mathcal{K}, EA)$ in M in w at t, f' should verify $\mathcal{K}$ in relation to $MS_w(a, t)$. It is at this point that we have to confront the difficulty mentioned earlier: that of accounting for the content of each of the attitude representations $\langle MOD, K_{MOD} \rangle$ belonging to $\mathcal{K}$ while doing justice to the internal coreference relations that connect the different attitude representations in $\mathcal{K}$. And it is at this point also that we cannot go into the necessary detail, but must refer to the publications on MSDRT that we cited earlier, in which its semantics is treated in detail. All we can do here is indicate what verification of Att-conditions comes to for the very simple cases we alluded to above.

We start with a case where the coreference problem does not arise — one in which the second argument of Att contains just one attitude component, as in (25a) and (25b). Let $< MOD, K_{MOD} >$ be the single propositional attitude component of the second argument of Att. Since K_{MOD} does not depend referentially on any other attitude component or components of $\mathcal{K}$ — by assumption there aren't any others —, K_{MOD} will determine a proposition $p(K_{MOD}, M, w, t)$ relative to M, w and t. This proposition is the set of all worlds w' in W_M such that for some $g \supseteq f'$ g verifies K_{MOD} in M in w' at t. f will count as verifying $s\colon Att(\xi, \mathcal{K}, EA)$ in M in w' at t in this case if (simplifying somewhat) $MS_w(f(\xi), t)$ contains a component $\langle MOD, q \rangle$ (where q is an M-proposition, i.e. a subset of W_M) such that q entails $p(K_{MOD}, M, w, t)$ (i.e. $q \subseteq p(K_{MOD}, M, w, t)$).[28]

The more challenging case is that where $\mathcal{K}$ has components $< MOD, K_{MOD} >$ that are referentially dependent on other components. It is here that the information state-based structure of the ISBASs that are the values of the function MS becomes important. In general the second argument $\mathcal{K}$ of an Att-condition defines, for any M, w and t, an ISBAS $J(\mathcal{K}, M, w, t)$. f verifies $s\colon Att(\xi, \mathcal{K}, EA)$ iff this ISBAS $J(\mathcal{K}, M, w, t)$ stands in the correct entailment relation to the ISBAS $MS_w(f(\xi), t)$.

The definition of this entailment relation is quite involved. We do not give it here and the following hint will have to do in its stead. Suppose that $< MOD, K_{MOD} >$ is a component of $\mathcal{K}$ that depends on only one other component $< MOD', K'_{MOD'} >$ and that $< MOD', K'_{MOD'} >$ does not depend on any attitude component of $\mathcal{K}$. Then, as in the previous case, $K'_{MOD'}$ determines in M in w at t a proposition $p(K'_{MOD'}, M, w, t)$ and K_{MOD} determines a function $J(K_{MOD}, M, w, t)$ which can be applied to the proposition $p(K'_{MOD'}, M, w, t)$ and returns upon application a stronger proposition (viz. the one that is determined in M in w at t by the DRS that can be obtained by merging K_{MOD} with $K'_{MOD'}$). In order that f verifies $s\colon Att(\xi, \mathcal{K}, EA)$ in M in w at t, $MS_w(f(\xi), t)$ must contain a component $\langle MOD', q \rangle$, such that q entails $p(K'_{MOD'}, M, w, t)$, and a component $<MOD', J >$ such that $J(q)$ entails the result of applying $J(K_{MOD}, M, w, t)$ to $p(K'_{MOD'}, M, w, t)$.

According to the verification conditions outlined above anchored entity representations in $\mathcal{K}$ contribute to the overall content of $\mathcal{K}$ only their "bare" referents (i.e. the values assigned by f to the discourse referents x' that represent their external anchors); the descriptive content of their

---

[27]A proper DRS is one in which all discourse referents are bound.

[28]We only demand that the proposition determined by K_{MOD} be entailed by the propositional content of some belief that $f(\xi)$ has in w' at t according to M, rather than insisting that the two propositional contents be identical. This is because we take it to be a general feature of attitude attribution that an attribution counts as true as long as it provides a truthful description of one of the attributee's attitudes; the description need not be exhaustive, in that it should capture the complete content of the attitude it describes — recall the remark on the partiality of mental state descriptions provided by Att-predications under (ii) of Section 2.3. (There is one qualification to this generalisation: This kind of partiality holds generally for "positive" attitudinal modes such as belief, desire and intention, but not for "negative" and "neutral" modes like doubt or uncertainty. In the present paper such "non-positive" modes are not considered.)

internal anchors is ignored. But, as we argued towards the end of Section 2.2, an anchored entity representation entails the belief that its internal anchor (the DRS that occurs as its second member) is true. Via these beliefs the contents of internal anchors can be made to matter to the verification of attitude conditions too.

3 Specific Use and Specific Interpretation

After the presentation of our background formalism in Section 2 and the informal remarks we have made in Section 1 about our own account of epistemic specificity, the present section, in which at long last we present this analysis, will hold few surprises.

The account consists of two parts:

(i) For a speaker to use an indefinite NP specifically is to use that NP to talk about an entity for which she has an anchored entity representation.

(ii) To interpret an indefinite NP as specific is to take the speaker to have used it to talk about one of her anchored entity representations.

We start with (i). As a preliminary step towards the definition of what it is to use an indefinite NP specifically we first introduce the more general notion of a speaker using an NP to *verbally represent* an entity-type constituent of a thought she is expressing in words. "Verbal representing" is a concept which we do not think can be reduced to simpler or more familiar terms and we treat it as a primitive. It is meant to capture the following intuition. Often (if not always) when a person expresses a thought in words, she had the thought, in some form, prior to the act of choosing the words in which she expresses it; and often in such cases her prior representation of the thought will involve mental entity representations. In such cases the sentence or sentences produced will often contain noun phrases which "translate" those entity representations, in the sense that the roles the noun phrases play in the sentence or sentences match the contributions that the corresponding entity representations make to the mental representation of the thought.

When an anchored entity representation of a speaker A properly represents some entity **d**, then A's use of an NP α to verbally realize this entity representation entails that she is using α to refer to **d**. This latter relation between α and **d** we refer to as "REP". So "REP(α,**d**)" means that A uses α to realize an anchored entity representation of hers whose external anchor is **d**.

Using the relation REP we can state our analysis of the epistemically specific use of an indefinite α as the use of α to talk about the entity represented by some entity representation. Formally:

(28) **<u>Definition 1.</u>** Let $\langle [ANCH, x_A], K' \rangle$ be a constituent of the mental state of speaker A and let **d** be the external anchor of this entity representation; let $\langle MOD, K \rangle$ be another constituent of this mental state such that x_A occurs in K and let S be a sentential utterance which A produces to express the thought represented in $\langle MOD, K \rangle$ and let α be an indefinite NP occurring in S such that $REP(\alpha, x)$. Then we say that A *uses α specifically in S* (to refer to **d**).

The notion of the *interpretation* of an NP as *used specifically* by the speaker is derivative from the notion given in **Definition 1**. Informally, an interpreter B is said to *interpret*, or *take*, an indefinite NP α occurring in an utterance S uttered by a speaker A *as used specifically* by A iff
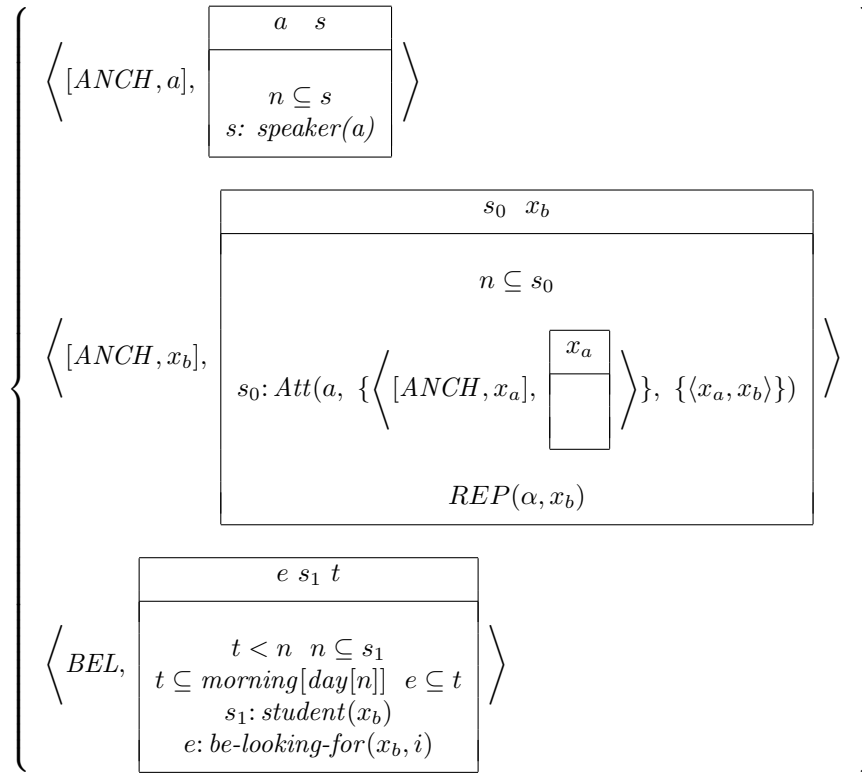
(i) B assumes that A used α to talk about the external anchor x'_A of an anchored entity representation $\langle [ANCH, x_a], K_a \rangle$ of hers; that is, B assumes that REP holds between α and the distinguished discourse referent x_a of this entity representation.

(ii) B represents the contribution that α makes to the content of A's utterance that he is interpreting in the form of an anchored entity representation $\langle [ANCH, x_b], K_b \rangle$.

(iii) The internal anchor K_B of this entity representation links its distinguished discourse referent x_b to the external anchor x'_a of A's entity representation, thereby making the external anchor of his

entity representation the same as the external anchor of the entity representation he attributes to A.

This is rather a mouthful, but an example should help. (29) is the representation which, according to the present proposal, results in the mind of B when he interprets the indefinite *a student* in A's utterance of (14) as having been used specifically by A.

(14)   A student was looking for you this morning.

(29)   (Recipient B's representation of the content of (14) on a specific
         interpretation of the indefinite *a student*)

$$
\left\{
\begin{array}{l}
\left\langle [ANCH, a], \boxed{\begin{array}{c} a \quad s \\ \hline n \subseteq s \\ s\!:\ speaker(a) \end{array}} \right\rangle \\[2em]
\left\langle [ANCH, x_b], \boxed{\begin{array}{c} s_0 \quad x_b \\ \hline n \subseteq s_0 \\ s_0\!:\ Att(a,\ \{\langle [ANCH, x_a], \boxed{\begin{array}{c} x_a \\ \hline \\ \end{array}} \rangle\},\ \{\langle x_a, x_b\rangle\}) \\ REP(\alpha, x_b) \end{array}} \right\rangle \\[2em]
\left\langle BEL, \boxed{\begin{array}{c} e\ s_1\ t \\ \hline t < n \quad n \subseteq s_1 \\ t \subseteq morning[day[n]] \quad e \subseteq t \\ s_1\!:\ student(x_b) \\ e\!:\ be\text{-}looking\text{-}for(x_b, i) \end{array}} \right\rangle
\end{array}
\right\}
$$

Comments:

    1. The last DRS in (29) is B's representation of the propositional content of A's utterance. It is assumed in (29) that this content representation has become part of a belief, i.e. that B has not only interpreted A's utterance but also accepted it as true and thus treats it as a belief. In general the impact of an utterance on the interpreter need not be his adoption of a new belief, not even when the utterance is an assertion, whose purpose is to convey new information to the audience; the utterance could just as well produce a doubt in the recipient, or a speculation whether what the speaker is asserting is true (cf. **?** on individual belief vs common belief, and possible variations in the contributions of utterances as regards participants' beliefs). However, for what remains of the paper it will be convenient to focus just on the case where the utterance is an assertion and the recipient accepts it, turning his semantic representation of the content of the utterance into the content representation of a new belief.

    2. The condition $morning[day[n]]$ stands for the morning of the day of the time n at which the representation is being entertained (for current purposes, where we are talking about face-to-face communication in which utterances get interpreted the moment they are made, this time can be identified with the utterance time).

3. The part of (29) that makes it an example of our formal analysis of epistemically specific interpretation is its second component — the anchored entity representation for the student that A referred to as *a student*. The anchor of this entity representation specifies its distinguished discourse referent x_b as the external anchor of some entity representation that B attributes to A — that entity representation of which he assumes that A has used α to talk about. Such internal anchors, which link their discourse referents to the external anchors of entity representations in the mind of someone else, we call *vicarious anchors*.[29]

It is perfectly compatible with B's having a vicariously anchored entity representation with external anchor **d** that B has no way of identifying it in any direct way. It is enough that B assumes there to be some **d** to which x_a is anchored. In this respect B's representation is similar to the one given in (25b) for sentence (4) — recall that (25b) represents the situation where the speaker of (4) takes the subject Mary to have an anchored representation for the discourse referent corresponding to the indefinite *a real estate agent*, but the speaker herself need not know who this person is. But nevertheless there is of course a crucial difference. In spite of the fact that B may not be able to say more about the entity represented by his anchored entity representation in (29), that entity representation will (if everything is as intended and expected) represent some particular entity, so that the content of the belief component of (29) is a singular proposition about this entity. The content of that belief of Bernard's of which (25b) can be understood to be a representation is not a proposition that is singular with respect to some real estate agent that Bernard believes Mary to have a *de re* belief about. That belief content is a proposition that involves existential quantification over real estate agents. The point of vicariously anchored entity representations is that they make it possible to have singular thoughts about things that one can identify only via the causal links between those entities and other agents with which one has interacted.

Generalising from this example we can define the notion of *interpreting* an indefinite NP *as used specifically* much in the way we already did informally above.[30]

(30)   **Definition 2.** Suppose that a speaker A utters a sentence S containing an indefinite NP α. An interpreter B of this utterance *interprets α as used specifically by A* iff

   (i) B takes A to have used α to refer to the external anchor **d** of some anchored entity representation $\left\langle [ANCH, x_a], \begin{array}{|c|} \hline x_a \\ \hline \\ \hline \end{array} \right\rangle$;

   (ii) B represents **d** in the form of an anchored entity representation whose vicarious internal anchor identifies the external anchor of its representation with the external anchor of the entity representation $\left\langle [ANCH, x_a], \begin{array}{|c|} \hline x_a \\ \hline \\ \hline \end{array} \right\rangle$ that B attributes to A.

Between them **Definition 1** and **Definition 2** contain the substance of our analysis of specificity.

---

[29] Note that B's vicarious anchor in (29) specifies no conditions for the internal anchor of the entity representation that B attributes to A. This reflects the fact that the content of sentence (14), for which (29) is B's semantic representation, provides no explicit information about the nature or content of the internal anchor of A's entity representation. Compare also the comments on the internal anchor of (25b) in Section 2.3.

[30] Sometimes an interpreter will represent an indefinite NP α by means of an anchored representation even though he doesn't assume the speaker to have made a specific use of α. These are cases where the interpreter takes himself to have identifying knowledge of the entity that satisfies the argument slot which α occupies in the given utterance. E. g. A says: *Last week Fred got himself a wife at last.* B, who knows that Fred has been courting Esmeralda for years and recently proposed to her, swearing that if she would turn him down he would remain single forever, sees himself in a position to provide an anchored representation for the NP *a wife*, and he can do this even if he doesn't take A's use of the indefinite to have been specific.

The difference between such cases and those covered by Definition 2, in which the recipient interprets the indefinite as having been used specifically by the speaker, should be clear from what we said about (29): Specific interpretation in the sense of Definition 2 yields a new entity representation with a vicarious anchor; the entity representation involved in the special case described in this footnote will typically not be a new one and as often as not its anchor will not be a vicarious one.

To conclude this section we repeat once more what we see as the main distinctive feature of our analysis. Unlike any other treatment of specificity that we are aware of, the analysis distinguishes explicitly between the perspective of the speaker and that of the hearer. Using an indefinite specifically and taking it to be used specifically are two different notions. The first relates the indefinite to the mind of the speaker, the second to the mind of the interpreter. Normally these relations will guarantee sameness of propositional content: when an indefinite that is part of an uttered sentence S stands in the two relations to, respectively, speaker and interpreter (and speaker and interpreter also observe the coding and decoding rules of the language in other respects), then the thought which the interpreter associates with S will express the same proposition as the thought which led the speaker to her utterance of S. But note that even in the best of cases the two thoughts will be different in form, for at the very least the representation of the speaker for the entity she is talking about will differ from the vicariously anchored representation that the interpreter introduces for it. The vicarious anchor of this latter representation refers to the speaker's representation and aligns itself with it. The speaker's own representation does not involve such a reference. And because of this difference in form the two thoughts can also differ in cognitive significance.

In the following section we focus exclusively on the perspective of the hearer. The question we will address has to do with the fine structure of semantic processing: What is the stage of the interpretation process at which taking an NP to have been used specifically makes its impact on the interpreter's representation?

4 Interpretation and Semantic Representation

In Section 2 we observed that anchored entity representations can be reduced to their purely internal content by ignoring their external anchors and turning their internally anchored representations into beliefs that entities of the described kinds exist.

In particular, the representation in (29) can be reduced to the one in (31), in which the new belief components — the belief that there is a speaker who made the represented utterance and the belief that there is a student whom this speaker is talking about by using the NP *a student* — have been amalgamated with the belief component of (29).

(31)   (Doxastic reduction of B's representation (29))
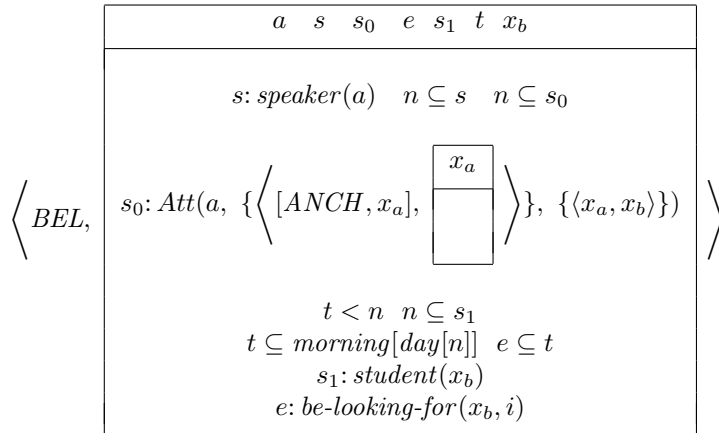
The content DRS of (31) invites comparison with the DRS given in (32) which can be obtained as a representation for (14) by applying a construction algorithm of the kind discussed in Kamp and Reyle (1993). Such a construction algorithm treats the indefinite *a student* as introducing a discourse referent x_b, together with the restrictor condition 's_1: student(x_b)'.[31]

---

[31] The discourse referents x_b and b have been chosen in order to make the relationship between (32) and (31) more transparent. In the construction of (32) the symbol b has been chosen to serve s the discourse referent introduced by the *NP you* of (14), and x_b to serve as the discourse referent introduced by *a student*. The choice of these symbols is arbitrary (so long as a new symbol is used when a discourse referent is introduced in the course of DRS construction).

$$(32)\quad \boxed{\begin{array}{c} x_b \quad s_1 \quad e \quad b \quad t \\ \hline \\ t < n \quad t \subseteq morning[day[n]] \quad e \subseteq t \\ s_1\colon student(x_b) \\ e\colon be\text{-}looking\text{-}for(x_b, b) \end{array}}$$

For ease of comparison let us consider the representation in (33) of a self-ascribing belief with the same propositional content as (32). (33) represents the belief that the addressee B of (14) can be expected to form when A says (14) to him. Its content representation corresponds to the content representation (32) in the sense that the argument position occupied by b in (32) is filled by the self-representing discourse referent i in (33).

$$(33)\quad \left\langle\; BEL\;,\; \boxed{\begin{array}{c} x_b \quad s_1 \quad e \quad t \\ \hline \\ t < n \quad t \subseteq morning[day[n]] \quad e \subseteq t \\ s_1\colon student(x_b) \\ e\colon be\text{-}looking\text{-}for(x_b, i) \end{array}}\;\right\rangle$$

When we compare the beliefs in (33) and (31), we see that the one in (31) contains a good deal of additional information, viz. that which it inherited from (29). So we can see (33) as yet a further reduction of (29), in which not only the referential presuppositions associated with the anchored representations have been eliminated — these had already been discarded in (31) — but the descriptions of those relations as well.

Our reason for bringing (32) and (33) into the discussion is this. What we take to be the classical view of the relationship between semantics and pragmatics is something along the following lines: Pragmatic mechanisms operate on content. So before these mechanisms can come into action, content (or at least some significant part of it) must already have been established. Since that is the task of semantics, semantics must have done its work, or some important part of its work, before pragmatics can kick in.

Translated into the terms of our DRT-based architecture this dependence of pragmatics on semantics takes the following form: First a DRS for a sentence is constructed (from a syntactic parse for it). In this way a content is assigned to the sentence. This content representation can then be used as input to the various pragmatic mechanisms that are needed to obtain the final interpretation of this sentence in its utterance context.

This general architectural constraint on the connections between semantic and pragmatic processing mechanisms is relevant for us because of an issue that has not yet been mentioned (though it was alluded to in the Introduction). It might be thought — and has been proposed explicitly, see Ebert et al. (2013a)) — that the effect of interpreting an indefinite as epistemically specific should be construed as a pragmatic 'add-on' to a specification of the semantics of the sentence containing it and not as part of the semantics of the sentence as such. When applied to *a student* in (14), and translated into the framework we are using, that proposal comes to this: First a semantic representation like (32) is constructed for the sentence, in which the indefinite *NP* is given an interpretation that is neutral with regard to the question whether the *NP* is taken as epistemically specific or non-specific. Then, at the level of pragmatics, this representation can be 'upgraded' by adding an anchored representation for the student that the interpreter takes the speaker to refer to (as well as one for the speaker herself; this is a harmless side effect of the way in which we have defined vicariously anchored representations). The result of this will be in essence the structure in (29), in which the anchored representations for speaker and student turn the propositional content of the belief into a doubly singular proposition.

Note also that the subscript $_b$ in x_b is only for mnemonic purposes. As far as the syntax and semantics of DRT and MSDRT are concerned, x_b is a single symbol, without relevant structure, just like the symbol b.

The entailment of (32) by the content representation of (31) is of a very simple type: one can get from this content representation to (32) simply by stripping it of some of the DRS-conditions in its condition set and of some of the discourse referents in its universe. Such deductions, where the conclusion can be obtained by eliminating parts from a single premise that is given in the form of a DRS, raise the question whether they could not be reversed: Start from the logically weaker conclusion and make your way back to the stronger premise by adding the discourse referents and conditions that are discarded when going from the premise to the conclusion.

There is also another aspect to this question. The passage from the content representation of (29) to the representation in (32) can be seen as the passage from the thought of a language user — the recipient of an utterance of (14) — to a user-neutral semantic representation of (14), of the kind that pretty much all established forms of formal semantics assign to this sentence (including DRT in the form in which it is developed in Kamp and Reyle (1993) or described in Geurts et al. (2015)). Or, making very much the same point, you can move from the mental state representation in (29) by doxastic reduction (which gets you from (29) to (31)) followed by stripping to get you from (31) to the belief representation in (33). And again we can ask the same sort of question: Are there cognitively plausible operations that may get us from the user-neutral (32) as starting representation to the mental state representation in (29)?

We will not present concrete answers to these questions. But we raise the questions nevertheless, because of their importance in connection with two divergent observations, which conclude this section and with that the paper as a whole.

The first of these has to do with indefinite *NP*s that occur in *there*-insertion sentences. Note that in the situation we described the effect that speaker *A* accomplishes by uttering (14) could have been achieved equally well by her using (34).

(34)   There was a student who was looking for you this morning.

In other words, the occurrence of the indefinite *a student* in (34) can be used specifically just as well as the occurrence in (14) and it can be just as easily taken to have been used specifically.

But there is an important difference between (14) and (34). As we suggested at the end of Section 1.3, the indefinite subject of a *there*-insertion sentence should be seen as contributing a variable that is bound by the existential operator expressed by *there be*. In DRT terms this means that a representation which results directly from this interface principle should have the form (32).

If that suggestion is correct, then the interpretation of an indefinite in a *there*-insertion sentence can only take place at a post-semantic interpretation level. For we cannot do justice all at once to the requirement that the discourse referent which is introduced by the indefinite is to be bound by the *there is*-operator and that it be bound at the same time by an internal anchor; that would put the representation truly in a double bind. However, what cannot be done all at once can be done in successive stages. Once the semantic representation for (34) has been constructed, it is then available for pragmatically driven operations that change the purely extensional binding by *there is* into binding by an anchored entity representation.

But what kinds of operations could be involved in such transformations, at what stage or level of processing they are to be applied and what triggers them? This question is in essence the same we just raised in connection with the inference from (29) to (33). As we have already said, we have no answers to these questions, but finding such answers will have to be a central part of our future work.

While specific interpretations of indefinites in *there*-insertion constructions can presumably enter into meaning representations only at some non-initial level of semantic processing, there are many other cases of specificity that suggest an opposite conclusion; here it seems plausible that specificity must make its impact already at the initial level, at which semantic sentence representations are computed directly from syntactic structures. These are the cases in which specificity is overtly marked. Overt specificity marking is found in language after language, and it seems quite common for a single language to have a range of different ways in which it can do this. It is a reasonable assumption, and arguably an inescapable one, that overt specificity marking devices are represented as part of the morphosyntactic structures from which initial semantic representations are computed,

and that that computation translates them directly into the semantic representations it outputs. Here we briefly discuss only one example of overt specificity marking, the Accusative *pe* marking of Romanian.[32]

An Example of explicit Specificity Marking: Romanian *pe*

On the one hand, Romanian *pe* must appear in front of definite *NP*s like demonstratives or proper names, as in (35a). In such cases it might be said to be pleonastic or a mere agreement marker: its presence only confirms what definites convey in any case, viz that reference is being made to some particular, uniquely identifiable individual. But *pe* can also occur in front of indefinites. Here it is optional, and when present it is not pleonastic.

(35)   a.   Ion a   angajat=o   **pe** Lucia.
            Ion PERF.3SG hired=ACC.3SG.FEM PE Lucia.
            "Ion has hired Lucia."
       b.   Ion a   angajat o   secretară.
            Ion PERF.3SG hired   a.FEM secretary.FEM
            "Ion has hired a secretary."
            (= he has hired someone in the capacity of secretary.)
       c.   Ion a   angajat=o   **pe** o   secretară.
            Ion PERF.3SG hired=ACC.3SG.FEM PE a.FEM secretary.FEM
            "Ion has hired someone who was a secretary."
            (= he has hired someone who was a secretary in some capacity or other)
       d.   Ion iubeşte un   agent   imobiliar.
            Ion loves   a.MASC agent.MASC real.estate-ADJ.MASC
            "Ion loves a real estate agent."
            (The speaker doesn't have any particular real estate agent in mind.)
       e.   Ion iubeşte **pe** un   agent   imobiliar.
            Ion loves   PE a.MASC agent.MASC real.estate-ADJ.MASC
            "Ion loves a real estate agent."
            (There is some particular real estate agent that Ion loves.)

Thus, (35b) means that Ion hired someone or other to be his secretary. (Maybe the person was a syntactician by profession, but with the job market being what it is becoming Ion's secretary was the best one could hope for.) In contrast, in (35c) the person hired must have been a secretary before, and it isn't implied that it was in order to be his secretary that Ion is said to have hired her — her new function could be that of Ion's cook or one of his gardeners. That is, *a angaja* 'to hire' can have a so-called resultative reading (cf. Moltmann (1997)), where the event of hiring causes the referent of the direct object noun phrase to have the property *secretary*. Only, in Romanian this reading is missing if the indefinite direct object is marked with *pe*. What this tells us is that the contribution of *pe*, or at least some part of its contribution, needs to be represented at some level of linguistic representation

   The two sentences of immediate concern to us are (35d)–(35e). These sentences correspond to the specific and the non-specific reading of Mary's statement in (5). (35d) is non-specific, while (35e), in which the indefinite is preceded by *pe*, only has the specific interpretation.

---

[32]Specificity markers tend to differ subtly in the constraints they impose on interpretation. Some examples are English *a certain, a/some particular, a specific, a/some given*, French *un certain, un … précis* (martin13), German *ein bestimmter, ein gewisser* — even for these languages and Romanian there are more options than the ones we have named; and these are just four of the world's thousands of languages. All in all, specificity marking is a rich field of semantic inquiry. It is an inquiry that is still in its early days. But a number of insightful observations about the semantics and pragmatics of some specificity markers have been made and some of these have been carefully and illuminatingly described. (See for instance Jayez and Tovena (2006) on English *a certain* and French *un certain* and Ebert et al. (2013a), Ebert et al. (2013b) on *bestimmt, gewiss*.)

   We are currently looking into the possibility of using MSDRT to formalise some of these results, with the ultimate aim of integrating them within a comprehensive account of the syntax-semantics-pragmatics interface.

We have argued that overt specificity markers must make their semantic contributions at the initial level of semantic representation. This argument, which applies in particular to the case of *pe*-indefinites, gains additional momentum from examples like (36), in which *pe*-indefinites form one constituent with the distributivity marker *câte*; this allows them to have narrow scope while yet retaining their specificity:

(36) Fiecare student cunoaşte PE câte un profesor de engleză.
Each(DET) student knows PE each(PRT) a.MASC professor.MASC of English
"Every student knows a certain English teacher"

(Example from Tigău (2013))

(36) is arguably an instance of dependent specificity in the sense of von Heusinger (2007). Although we are not presenting an actual analysis of sentences like (36) — we do that in Bende-Farkas and Kamp (forthcoming) — we do want to observe that adding the contribution made by *pe* at a non-initial level of semantic processing, after the quantification contributed by *fiecare student* has already been made part of the representation at the initial level, would be quite difficult and unnatural.[33]

There are two differences, we have suggested, between Romanian *pe*-indefinites (and, more generally, for any indefinites that are overtly marked for specificity) on the one hand and English indefinites beginning with *a* (or *some*) on the other: (*i*) What *pe*-indefinites make explicit, *a*-indefinites leave open; so there is a sense in which *a*-indefinites are ambiguous, whereas Romanian indefinites in direct object position — both those that are *pe* marked and those that aren't — are not. (*ii*) *Pe*- marked indefinites make their specificity contributions to initial semantic representations, whereas at least some *a*-indefinites that are given specific interpretations (those occurring in *there is*-constructions) can make their specificity contribution only at some later stage of semantic processing.

One way in which one might be tempted to rephrase this second difference is that the specificity contributions made by *pe*-indefinites are part of the semantics, whereas the specificity contributions of at least some occurrences of specifically interpreted *a*-indefinites must be processed at some stage after the initial representation has been computed, belong to pragmatics. When the difference is put in these terms, however, it should be kept in mind that the distinction between semantics and pragmatics that it presupposes is orthogonal to another one, according to which among the pragmatic

---

[33]One of the reviewers of an earlier version remarked that the sentence-level correlate of specificity is evidentiality, and wondered why specific indefinites are not acceptable in questions and commands. What we found is that Romanian *pe*-indefinites *are* perfectly appropriate in questions. (They are less appropriate in commands.)

(i) a. Q. (L=)Ai văzut pe un profesor?
Q. (ACC.3SG.MASC=)PERF-2SG seen PE a.MASC professor.MASC
"Have you seen a certain professor?"

b. A₁. Nu l=am văzut, este plecat în concediu.
A₁. Not ACC.3SG.MASC=PERF-1SG seen, is left in holiday.
"I haven't seen him, he's on holiday."

c. A₂. Pe el nu l=am văzut, l=am văzut însă pe
A₂. PE he not ACC.3SG.MASC=PERF-1SG seen, ACC.3SG.MASC=PERF- 1SG seen though PE
profesorul lui Petre.
professor-DEFART.MASC.SG GEN Peter
"I haven't seen him, but I have seen Peter's professor.' '

The question in (ia) is about a particular individual. There is nothing marked about it; in particular the dialogue doesn't have a conspiracy connotation (à la "We know who we are talking about, but we're not telling the others"). In using a *pe*-marked indefinite the speaker introduces an entity into the discourse that is at the same time new *and* anchored (and identifiable). This is remarkable, since indefinites in questions and commands — Portner (2005), Portner and Zanuttini (2003) — are typically non-specific; if they can be given a specific interpretation at all, there is said to be an air of conspiracy about them. But, to repeat, the question in (ia) is perfectly appropriate as a query about a particular individual. The second answer, (ic), makes this even clearer: Had the question involved a non-specific indefinite, the addressee could have simply replied with the Romanian version of *I have seen Peter's professor*. But Peter's professor is clearly not the intended anchor for the question; if the addressee can recognise this intended anchor, then his answer has to be about *that* person and not about professors in general. It isn't clear to us at this point what implications this and other facts about the occurrence of *pe* in non-indicative speech acts may have for the question at issue here: at what level of semantic or pragmatic processing does *pe* make its contribution?

aspects of interpretation are those that relate, in some way or other, to the communicational situation in which the interpreted utterance occurs. In this second sense of "pragmatic" the representations of *pe*-indefinites *are* pragmatic, no less so than the contributions made by the specifically interpreted *a*-indefinites of English. According to our proposal for the analysis of epistemic specificity *all* specificity contributions are pragmatic in this sense, since all of them involve attributions by the interpreter to the mental state of the speaker, and thus go beyond the purely grammatical properties of the sentence or sentences uttered.

Romanian indefinite direct object NPs with *pe* differ not only from English indefinite *NP*s beginning with *a*, but also from Romanian indefinites that do not occur as argument *NP*s with Accusative case marking. The specificity behaviour of these latter Romanian *NP*s is much the same as that of *a*-indefinites in English. In light of what is proposed in this paper this entails that a Romanian sentence with a *pe*-marked indefinite direct object and a (non-marked) indefinite subject may have an initial semantic representation in which the direct object is already represented as a specific *NP*, while a specific interpretation of the subject *NP* will leave its imprint only at a subsequent stage of representation construction.[34]

We have argued that Romanian *pe*-indefinites make their specificity contributions to initial semantic representations, but that the specificity contributions of English *a*-indefinites in *there*-insertion sentences can only enter into the representation at some later stage. But what about specifically interpreted *a*-indefinites in other syntactic constructions, such as the occurrence of *a student* in (14)? At what level do such indefinites make their specificity contributions when they are specifically interpreted? Our discussion of (14) seems to point clearly to one conclusion: When an indefinite like that in (14) is given an epistemically specific interpretation, then this must happen after the construction of the initial semantic representation. For without this initial semantic representation, which affords the interpreter a first take on the content of the sentence that he is interpreting, he will have no basis for choosing between a specific and a non-specific interpretation. In other words, the choice between a specific and a non-specific interpretation of such indefinites will be possible only after an initial semantic representation has been put into place; consequently, when the interpreter chooses the specific interpretation option, the contribution that is made by this choice will become part of the semantic representation only after the construction of the initial representation.

But this conclusion is not inescapable. There may well be cases where the interpreter of a sentence containing an *a*-indefinite is in a position to infer just from the syntactic form of the sentence and the context in which it is used that the indefinite must have been used specifically. In such cases the interpreter would be in a position to integrate the specificity contribution that comes with this inference into the initial semantic representation he constructs for the sentence. If we are right in thinking that there are such cases, then the question at what point specificity information enters into the semantic representation is even more complex than is implied by our considerations so far: the point of entry for this information may vary even between different utterances of the very same indefinite-containing sentence.

5 Summary and pointers to what lies beyond

The aim of this paper has been to present a communication-based account of epistemic specificity. This account distinguishes between two aspects of epistemic specificity, the speaker-related aspect of *using an indefinite specifically* and the hearer-related aspect of *taking the speaker to have made a specific use of an indefinite*. The account makes use of MSDRT, a DRT-based formalism designed for the description of propositional attitudes and attitudinal states. Two features of this formalism

---

[34] The special specificity features of direct objects in Romanian instantiate a more general pattern, that of grammatical specificity marking being restricted to direct objects. In the literature this phenomenon is known as 'D(ifferential) O(bject) M(arking)', a term that reflects the morpho-syntactic restrictions to which it is subject. DOM is found in a considerable variety of languages, though it does not always serve to mark specificity (Aissen (2003), Enç (1991) or von Heusinger (2002b)). In those languages where differentially marked indefinites must be interpreted as specific and unmarked indefinite direct objects as non-specific, one finds the same complexity as in Romanian, with some cases of specificity marked morphologically, while others (e.g. subjects) cannot be marked in this way.

are essential to what we are saying about specificity: (i) Its ability to describe complex mental states (as sets consisting of entity representations and propositional attitudes); (ii) the analysis it offers of agent-internal and agent-external aspects of direct reference in terms of the distinction between the *internal* and *external anchors* of entity representations.

MSDRT has many applications in semantics and the philosophy of mind. Of special importance to linguistics is the use that can be made of it in the analysis of what happens in conversation. It is now widely acknowledged that much more goes on in verbal communication than is captured by a model according to which the speaker produces her utterance, the addressee constructs a representation of that utterance and that's it. In actual conversations the addressee will typically make inferences about the speaker's mental state on the basis of what the speaker is saying. (Taking the speaker to have made a specific use of some indefinite, in the way we have analysed that in this paper, is an instance of this, but it is only one out of many.) And the speaker, for her part, will often be able to draw conclusions about the state of mind of the addressee: for instance about whether or not he is agreeing with what she is saying, or whether he has correctly understood her. These inferences take the form of attributions that one conversation participant makes in thought about the mental state of the other. Furthermore, these mutual attributions allow for iteration. For instance A may attribute to B a thought which attributes a certain thought to *her*. In principle at least such attributions can be of unbounded complexity, with any number of nested attributions back and forth.

How *much* mutual attribution goes on in conversation is a matter of debate. But that some of this goes on in any normal conversation is beyond question. MSDRT is well equipped to describe such mutual attributions, so it can help us, in this way at least, to develop formally precise theories of verbal communication that address these inferential aspects of speaker-hearer interaction. The account we have proposed here of the specific interpretation of indefinites can be seen as a first small step on the road towards such a theory.

Developing such a theory is one of two long-range projects towards which this paper points. The other is the development of an account of semantic and pragmatic aspects of the interpretation of text. As the discussion in the paper suggests, both projects will probably need to postulate different levels of meaning representation and to articulate principles for moving from one representation level to the next. But beyond this there is little that we feel we can say with any confidence. How many levels of representation will be needed *in toto*; what information belongs to which levels; how interpretation proceeds from one level to the next; and how individual languages differ on any or all of these points — all these are questions that we are beginning to see how to ask. But finding the answers is a different matter and there is no way of telling how long that is going to take.

Authors' Addresses

Ágnes Bende-Farkas
agnesbf@gmail.com, agnesbf@nytud.hu
Research Institute for Linguistics
Hungarian Academy of Sciences
33. Benczúr St.
H-1068 Budapest, Hungary

Hans Kamp
hans@ims.uni-stuttgart.de
4. Koppentalstrasse
D-70192 Stuttgart, Germany

Between January 15 and May 15:
Department of Philosophy

University of Texas
1 University Station C 3500
Austin, Texas; 78712 USA

Acknowledgements

References

Aissen, Judith. 2003. Differential Object Marking: Iconicity vs Economy. *Natural Language and Linguistic Theory* 21:435–483.

Almog, Joseph, John Perry, and Howard Wettstein, ed. 1989. *Themes from kaplan*. Oxford University Press.

Bende-Farkas, Ágnes. 2002. Verb–Object Dependencies in Hungarian and English. A DRT-based Account. Doctoral Dissertation, Stuttgart University.

Bende-Farkas, Ágnes, and Hans Kamp. 2001. *Indefinites and Binding. From Specificity to Incorporation.* Lecture Notes at ESSLLI'01, Helsinki; available from `http://www.helsinki.fi/esslli/`.

Bende-Farkas, Ágnes, and Hans Kamp. forthcoming. *Specific Indefinites: Anchors and Functional Readings.* MS, IMS Stuttgart.

Brasoveanu, Adrian, and Donka F. Farkas. 2009. Exceptional scope as discourse reference to quantificational dependencies. In *Seventh international tbilisi symposium on logic, language and computation*, ed. Peter Bosch, David Gabelaia, and Jérôme Lang, Springer Lecture Notes in Computer Science, 165–179. Berlin–New York: Springer Verlag.

Breheny, Richard. 2003. Exceptional Scope Indefinites and Domain Restriction. In *Proceedings of Sinn und Bedeutung 7*, ed. Matthias Weisberger, 38–52.

Chastain, C. 1975. Reference and Context. In *Language, Mind and Knowledge*, volume 7 of *Minnesota Studies in the Philosophy of Science*. Minneapolis: U of Minnesota Press.

Chierchia, Gennaro. 2002. A Puzzle about Indefinites. In *Semantic interfaces: Reference, anaphora and aspect*, ed. Carlo Cecchetto, Gennaro Chierchia, and Maria Teresa Guasti, 51–89. Stanford: CSLI.

Chisholm, Roderick. 1976. Knowledge and Belief: "De Dicto" and "De Re". *Philosophical Studies* 29:1–20.

Cooper, Robin. 1979. The Interpretation of Pronouns. In *Selections from the 10th groningen roundtable*, ed. Frank Heny and Helmut S. Schnelle, number 10 in Syntax and Semantics, 61–92. New York: Academic Press.

Donnellan, Keith. 1966. Reference and Definite Descriptions. *Philosophical Review* 75:281–304.

Ebert, Christian, Cornelia Ebert, and Stefan Hinterwimmer. 2013a. The semantics and pragmatics of *bestimmt* and *gewiss*. In *Proceedings of NELS 40*, ed. Seda Kan. MIT Press.

Ebert, Christian, Cornelia Ebert, and Stefan Hinterwimmer. 2013b. The Interpretation of the German Specificity Markers *Bestimmt* and *Gewiss*. In Ebert and Hinterwimmer (2013), 31–74.

Ebert, Cornelia, and Stefan Hinterwimmer, ed. 2013. *Different kinds of specificity across languages*. Number 92 in Studies in Linguistics and Philosophy. Springer.

van Eijck, Jan, and Hans Kamp. 2011. Representing Discourse in Context. In *Handbook of logic and language*, ed. Johan van Benthem and Alice ter Meulen. Elsevier, 2 edition.

Elbourne, Paul. 2005. *Situations and individuals*. MIT Press.

Enç, Mürvet. 1991. The Semantics of Specificity. *Linguistic Inquiry* 22:1–25.

Evans, Gareth. 1977. Pronouns, Quantifiers and Relative Clauses (1). *The Canadian Journal of Philosophy* 7:467–536.

Evans, Gareth. 1980. Pronouns. *Linguistic Inquiry* 11:337–362.

Farkas, Donka. 1981. Quantifier Scope and Syntactic Islands. In *Proceedings of CLS 17*, 59–66. Chicago: Chicago University Press.

Farkas, Donka. 1996. Specificity and Scope. In *Langues et Grammaire 1*, ed. Lea Nash and George Tsoulas, 119–137. Paris.

Farkas, Donka F. 2002. Specificity Distinctions. *Journal of Semantics* 19:1–31.

Fodor, Janet Dean. 1970/1979. The linguistic description of opaque contexts. Doctoral Dissertation, MIT.

Fodor, Janet Dean, and Ivan Sag. 1982. Referential and Quantificational Indefinites. *Linguistics and Philosophy* 5:355–398.

Fox, Danny, and Uli Sauerland. 1996. Illusive Scope of Universal Quantifiers. In *Proceedings of NELS*, ed. K. Kusumoto, volume 26, 71–86. UMass.

Geurts, Bart. 2010. Specific Indefinites, Presupposition and Scope. In *Presuppositions in discourse*, ed. Rainer Bäuerle, Uwe Reyle, and Thomas Ede Zimmermann, number 21 in CRISPI, 125–128. Bingley UK: Emerald.

Geurts, Bart, David Ian Beaver, and Emar Maier. 2015. Discourse Representation Theory. Stanford Encyclopedia of Philosophy. `plato.stanford.edu`.

Groenendijk, Jeroen, Martin Stokhof, and Frank Veltman. 1996. Coreference and Modality. In *The handbook of contemporary semantic theory*, ed. Shalom Lappin, 179–213. Oxford: Blackwell, 1st edition.

Heim, Irene. 1982. The semantics of definite and indefinite noun phrases. Doctoral Dissertation, University of Massachusetts at Amherst, Amherst.

Heim, Irene. 1990. E-Type Pronouns and Donkey Anaphora. *Linguistics and Philosophy* 13:137–177.

Heim, Irene. 1991. Artikel und Definitheit. In *Semantik. ein internationales handbuch der zeitgenoessischer forschung*, ed. Arnim von Stechow and Dieter Wunderlich, 501–549. de Gruyter.

von Heusinger, Klaus. 2002a. Specificity and Definiteness in Sentence and Discourse Structure. *Journal of Semantics* 19:245–274.

von Heusinger, Klaus. 2002b. The Cross-Linguistic Implementations of Specificity. In *Meanings in contrast: The cambridge papers*, ed. K. Jaszczolt and Ken Turner, volume 2 of *Pragmatics and Beyond*, 405–421. Amsterdam and Philadelphia: John Benjamins.

von Heusinger, Klaus. 2007. Referentially Anchored Indefinites. In *Existence: Syntax and semantics*, ed. Ileana Comorovski and Klaus von Heusinger, number 84 in Studies in Linguistics and Philosophy, 273–292. Dordrecht: Springer.

Ioup, Georgette. 1977. Specificity and the Interpretation of Indefinites. *Linguistics and Philosophy* 1:233–245.

Jackendoff, Ray. 1972. *Semantic interpretation in generative grammar*. Cambridge, Mass.: MIT Press.

Jayez, J., and L.M. Tovena. 2006. Epistemic Determiners. *Journal of Semantics* 23:217–250.

Kamp, H. 2016. Entity Representations and Articulated Contexts. To appear in a collection in honor of David Kaplan.

Kamp, H. 2017. Semantics ii, ut. Unpublished notes for a graduate course offered at the University of Texas, Austin.

Kamp, H., J. van Genabith, and U. Reyle. 2011. Discourse Representation Theory: An Updated Survey. In *Handbook of philosophical logic*, ed. Dov Gabbay, volume XV, 125 – 394. Elsevier, 2nd edition.

Kamp, Hans. 1981. A Theory of Truth and Semantic Representation. In *Formal methods in the study of language*, ed. Jeroen Groenendijk. Amsterdam: Mathematisch Centrum. Republished in K. v. Heusinger and A. ter Meulen, eds. 2013. *The Dynamics of Meaning and Interpretation. Selected Papers of Hans Kamp.* Brill, Leiden. Pp. 329–370.

Kamp, Hans. 1990. Prolegomena to a Structural Theory of Belief and Other Attitudes. In *Propositional attitudes*, ed. C. Anthony Anderson and Joseph Owens, number 20 in CSLI Lecture Notes, 27–90. Stanford: CSLI.

Kamp, Hans. 2003. Einstellungszustände und Einstellungszuschreibungen in der Diskursräpresentationstheorie. In *Intentionalität zwischen subjektivität und weltbezug*, ed. Ulrike Haas-Spohn. Paderborn: Mentis.

Kamp, Hans. 2011. Representing *de Se* Thoughts and Their Reports. Unpublished.

Kamp, Hans. 2013. Introduction to Propositional Attitudes. In *The Dynamics of Meaning and Interpretation. Selected Papers of Hans Kamp*, ed. Klaus von Heusinger and Alice ter Meulen, 505–512. Leiden: Brill.

Kamp, Hans, and Uwe Reyle. 1993. *From discourse to logic*. Dordrecht: Kluwer.

Kripke, Saul. 1977/1990. Speaker's Reference and Semantic Reference. In *The philosophy of language*, ed. A.P. Martinich, 248–267. Oxford Uiversity Press.

Ladusaw, William A. 1982. Semantic Constraints on the English Partitive Construction. In *Proceedings of WCCFL*, ed. D. Flickinger, M. Macken, and N. Wiegand, volume 1. Stanford: California: Stanford Linguistics Association.

Lewis, David K. 1969. *Convention: A philosophical study*. Cambridge MA: Harvard University Press.

Maier, Emar. 2015. Parasitic Attitudes. *Linguistics and Philosophy* 38:205–236.

Milsark, Gary. 1977. Towards an Explanation of Certain Peculiarities of the Existential Construction in English. *Linguistic Analysis* 5:1–29.

Moltmann, Friederike. 1997. Intensional Verbs and Quantifiers. *Natural Language Semantics* 5:1–52.

Portner, Paul. 2005. The Semantics of Imperatives within a Theory of Clause Types. In *Proceedings of SALT 14*, ed. R.B. Young, 235–252. Cornell University, Ithaca: CLC.

Portner, Paul, and Katsushiko Yabushita. 2002. Specific Indefinites and the Information Structure Theory of Topics. *Journal of Semantics* 18:271–297.

Portner, Paul, and Raffaella Zanuttini. 2003. Exclamative Clauses: At the Syntax–Semantics Interface. *Language* 79:39–81.

Quine, W.v.O. 1966a. Quantifiers and Propositional Attitudes. In Quine (1966b). Originally appeared in *Journal of Philosophy* **56**, 1956.

Quine, W.v.O. 1966b. *The ways of paradox and other essays*. New York: Random House.

Rodman, Robert. 1976. Scope Phenomena, "Movement Transformations", and Relative Clauses. In *Montague grammar*, ed. Barbara Hall Partee, 165–176. New York: Academic Press.

Ross, John R Robert. 1967. Constraints on variables in syntax. Doctoral Dissertation, MIT.

Schiffer, Stephen R. 1972. *Meaning*. Oxford: Clarendon Press.

Schwarzschild, Roger. 2002. Singleton Indefinites. *Journal of Semantics* 19:289–314.

Szabolcsi, Anna. 2004. Positive Polarity — Negative Polarity. *Natural Language and Linguistic Theory* 22:409–452.

Tigău, Alina. 2013. Points of Similarity between Differential Object Marking in Romanian and German scrambling, with a Focus on Indefinite Direct Objects. Language, Individual & Society, Volume 7, Part 1. Available from http://www.scientific-publications.net.