Conceptual Covers in Dynamic Semantics

Maria Aloni*

aloni@hum.uva.nl

Abstract

In dynamic semantics three styles of quantification have been proposed that involve two different ways of interpreting free and quantified variables: (1) Variables as denoting single partial objects; (2) Variables as ranging over a number of alternative total objects. I will show that the first view leads to problems of underspecification and the second to problems of overspecification. I will propose a new style of dynamic quantification in which variables are interpreted in a way that avoids these problems: (3) Variables as ranging over a number of alternative definite objects (concepts). By relativizing quantification to ways of conceptualizing the domain, we avoid the cardinality problems which arise when we quantify over concepts rather than objects.

1 Quantification in Dynamic Semantics

In dynamic semantics, sentences describe transitions across a space of information states. Information states are generally defined as sets of possibilities (here world-assignment pairs) and meanings are state transitions. Updating with sentences may reduce the size of the states or may yield richer states. Atoms or negations narrow down the alternatives under consideration by eliminating the world-assignment pairs that do not satisfy the information contents of these sentences. Existentially quantified sentences instead add structure to the state by setting up new items as potential topics for further discourse: $\exists x\phi$ adds x and selects a number of possible values for it; the fact that in the output state(s) x is defined means that recurrences of x in later sentences can have the effect of anaphoric reference.

Information about variables is generally modeled in one of the following two ways:

1. Variables are interpreted as single *partial* objects.[1] The introduction of new items is defined in terms of *global extensions* that involve adding fresh variables and assigning them as possible values all elements of the universe of discourse. All of the values which variables can take are considered simultaneously.

2. Variables are taken to range over a number of *total* objects. The introduction of a new item is defined in terms of *individual extensions* that lead to the states resulting from adding a variable and assigning it a single element of the

*This paper has grown out of Aloni 1997.

[1] Partial objects are the structured entities that constitute the interpretations of variables in information states. In Dekker 1993, they are defined as functions that assign to each possibility in a state the value of the corresponding variable in that possibility. A partial object is called total if it is a constant function. In the picture below, the partial object corresponding to the interpretation of the variable x is represented by the vertical column below x.

universe as its value. The values variables can take are considered one by one as disjoint alternatives.[2]

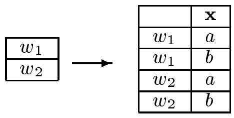

1. Global Extension

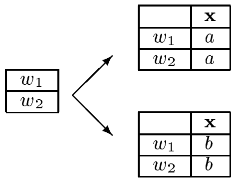

2. Individual Extension

Global extensions yield unique output states, whereas individual extensions produce as many different outputs as there are members of the universe. This involves splitting up the initial state into different alternatives: later sentences will be interpreted with respect to each of them in a parallel fashion.

In the literature three different interpretations have been proposed for the dynamic existential quantifier that involve one or the other way of interpreting free[3] or quantified variables:

Random Assignment (RA) is the standard interpretation. It is defined in terms of global extension; it involves assigning to fresh variables all individuals from the universe of discourse as possible values. In this way, quantified and free variables are interpreted uniformly as single indefinite partial objects, where further updates will tend to make these objects more definite and less partial.[4]

Slicing (SL) is defined in terms of individual extension; it involves splitting up the update procedure, so that the individuals that a variable can take as possible values are considered one by one, as disjunct alternatives, and not all at once. In this way, quantified and free variables are interpreted uniformly as ranging over a number of alternative total objects, where further updates will tend to eliminate certain alternatives.[5]

Moderate Slicing (MS) follows the slicing procedure as long as we are inside the syntactic scope of a quantifier, but lumps the remaining alternatives together once we are outside its scope. In this way, quantified variables range over a number of alternative total objects, whereas free variables are interpreted as single partial objects.[6]

These different styles of quantification lead to different results only in connection with notions that are sensitive to global properties of information states, i.e., notions that take a state as a whole and not pointwise with respect to the possibilities in it. This is not surprising: if we take states holistically it is obvious that it matters which possibilities are lumped together to form a state and which are kept separate. Examples of holistic notions are epistemic modals,[7] presupposition,[8] and the notion of support.[9]

Although the analysis of combinations of quantifiers and holistic notions motivated the use of (moderate) slicing instead of random assignment, I will argue that

[2] I use these pictures to represent shifts on information states. The tables correspond to information states. On the topmost horizontal row are displayed in bold characters the variables that are defined in the state. Each other horizontal row represents a world-assignment pair element of the state. More specifically the left column contains the world-coordinate and the right column the range of the assignment functions with each individual displayed right below the variable it gets assigned to. The universe is assumed to consist only of two individuals a and b.

[3] By a 'free' variable, I mean a variable x not occurring inside the syntactic scope of a quantifier $\mathcal{Q}x$. Typically, such 'free' occurrences may still be dynamically bound by a quantifier.

[4] Cf. Dekker 1993, Groenendijk and Stokhof 1991, Heim 1982, Heim 1983.

[5] Cf. van Eijck and Cepparello 1994.

[6] Cf. Beaver 1994, Dekker 1994, Groenendijk et al. 1996.

[7] Cf. Dekker 1993, Groenendijk et al. 1996, van Eijck and Cepparello 1994, Veltman 1997.

[8] Cf. Beaver 1994, Beaver 1995, Heim 1983.

[9] Cf. Dekker 1997, Groenendijk et al. 1996.

precisely in such contexts critical problems emerge for all three styles of quantification. Before turning to the illustration of these problems, let me introduce the relational dynamic semantics that supplies the general framework for the comparison of the three approaches.

Formal Framework

The core of the semantic framework that I will adopt is a relational version of the update semantics MDPL presented in Dekker 1993 with the addition of the presupposition operator introduced in Beaver 1995. The *language* $\mathcal{L}$ is a standard predicate logical language with the addition of two sentential operators, the modal operator $\Diamond$ and the presupposition operator ∂. Given $\mathcal{L}$, an *information space* $\mathcal{I}$ for $\mathcal{L}$ is a pair $\langle W, D\rangle$ where W, the set of possible worlds, is a non-empty set of interpretation functions for the non-logical constants in $\mathcal{L}$, and D, the domain of discourse, is a non-empty set of individuals.[10] Information states are sets of possibilities. They are defined as in Dekker 1993 and Heim 1982 as sets of world-assignment pairs in which all the assignment functions have the same domain.

Definition 1 [Information States] Let $\mathcal{I} = \langle D, W\rangle$ be an information space for $\mathcal{L}$. Let $\mathcal{V}$ be the set of individual variables in $\mathcal{L}$. The set $\Sigma_{\mathcal{I}}$ of *information states* based on $\mathcal{I}$ is defined as:

$$\Sigma_{\mathcal{I}} = \bigcup_{X \subseteq \mathcal{V}} \mathcal{P}(W \times D^X)$$

A possibility in an information state contains enough information for the interpretation of the basic expressions in $\mathcal{L}$.

Definition 2 Let α be a basic expression in $\mathcal{L}$ and $i = \langle w, g\rangle$ a possibility in $W \times D^X$ for some $X \subseteq \mathcal{V}$. *The denotation* of α in i is defined as:

i if α is a non-logical constant, then $i(\alpha) = w(\alpha)$;

ii if α is a variable in X, then $i(\alpha) = g(\alpha)$, undefined otherwise.

Survival is a relation between a possibility and an information state.

Definition 3 [Survival] Let $\sigma \in \Sigma_{\mathcal{I}}$ & $i = \langle w, g\rangle \in \sigma'$ for some $\sigma' \in \Sigma_{\mathcal{I}}$.

$$i \prec \sigma \quad \text{iff} \quad \exists\langle w', g'\rangle \in \sigma : w = w' \ \& \ g \subseteq g'$$

A world-assignment pair i survives in a state σ iff σ contains a possibility j such that j is the same as i except for the possible introduction of new variables. I can now define the main semantic clauses and the notion of support in a parallel fashion.

Definition 4 [Support] Let $\sigma \in \Sigma_{\mathcal{I}}$ and ϕ in $\mathcal{L}$.

$$\sigma \models \phi \quad \text{iff} \quad \exists\sigma' : \ \sigma[\phi]\sigma' \ \& \ \forall i \in \sigma : i \prec \sigma'$$

A state σ supports a sentence ϕ iff all possibilities in σ survive simultaneously in at least one of the states resulting from updating σ with ϕ, where updates are defined as follows:

Definition 5 [The Core of the Semantics]

$$\begin{array}{rll}
\sigma[Rt_1, ..., t_n]\sigma' & \text{iff} & \sigma' = \{i \in \sigma \mid < i(t_1), ..., i(t_n) > \in i(R)\}; \\
\sigma[\neg\phi]\sigma' & \text{iff} & \sigma' = \{i \in \sigma \mid \neg\exists\sigma'' : \sigma[\phi]\sigma'' \ \& \ i \prec \sigma''\}; \\
\sigma[\phi \wedge \psi]\sigma' & \text{iff} & \exists\sigma'' : \ \sigma[\phi]\sigma''[\psi]\sigma'; \\
\sigma[\Diamond\phi]\sigma' & \text{iff} & \sigma' = \{i \in \sigma \mid \exists\sigma'' \neq \emptyset : \sigma[\phi]\sigma''\}; \\
\sigma[\partial\phi]\sigma' & \text{iff} & \sigma \models \phi \ \& \ \sigma[\phi]\sigma'.
\end{array}$$

[10] Possible worlds can be identified with interpretation functions because I am assuming that all possible worlds share the same domain.

Updating a state σ with an atomic formula preserves those possibilities in σ which satisfy the formula in a classical sense. The negation of ϕ eliminates those i in σ which survive after updating σ with ϕ. Conjunction is relational composition.

Modal sentences $\Diamond\phi$ are interpreted in Veltman's style, as consistency tests. Updating with $\Diamond\phi$ involves checking whether ϕ is consistent with the information encoded in the input state σ. If the test succeeds, i.e., if at least one world-assignment pair in σ survives an update with ϕ, then the resulting state is σ itself, so nothing happens; if the test fails, the output state is the empty set, i.e. the absurd state (cf. Veltman 1997).

∂ is Beaver's presupposition operator. $\partial\phi$ should be read as "it is presupposed that ϕ" and is interpreted as an update that is defined on a state σ only if ϕ is already supported in σ. Notice that presuppositions are not simple tests — the output state may vary from the input state in that it can contain new discourse items (cf. Beaver 1995).

Consistency tests, presupposition and support are holistic notions because they relate to properties of the whole state, not of its individual elements.

Three different systems can be developed from this core semantics depending on which of the three above-mentioned interpretations of dynamic existential quantification we adopt. To define them we need introduce the auxiliary notions of assignment operations, global extensions and individual extensions.

Assignment operations extend possibilities by adding fresh[11] variables and assigning to them as values individuals from the domain.

Definition 6 [Assignment Operations] Let $\langle w, g\rangle \in W \times D^X$ for some $X \subseteq \mathcal{V}$, $x \in \mathcal{V} \setminus X$ and $d \in D$.

$$\langle w, g\rangle[x/d] = \langle w, g \cup \{\langle x, d\rangle\}\rangle$$

In terms of assignment operations we define both global and individual extensions. Let $\mathrm{dom}(\sigma)$ be the set of individual variables defined in σ.

Definition 7 [Extensions] Let $\sigma \in \Sigma_{\langle D,W\rangle}, x \in \mathcal{V}\backslash\mathrm{dom}(\sigma)$ and $d \in D$.

i $\sigma[x] = \{i[x/d] \mid d \in D \;\&\; i \in \sigma\}$ (*global*)

ii $\sigma[x/d] = \{i[x/d] \mid i \in \sigma\}$ (*individual*)

Global extensions add fresh variables and randomly assign *all* elements of the universe of discourse to them. Whereas individual extensions enlarge the domain of the state by assigning *single* elements of D to fresh variables. Finally, we can define Random Assignment, Slicing and Moderate Slicing.

Definition 8 [Three Styles of Quantification]

$$\begin{array}{rcl} \sigma[\exists x\phi]_{RA}\sigma' & \text{iff} & \sigma[x][\phi]\sigma'; \\ \sigma[\exists x\phi]_{SL}\sigma' & \text{iff} & \sigma[x/d][\phi]\sigma' \text{ for some } d \in D; \\ \sigma[\exists x\phi]_{MS}\sigma' & \text{iff} & \sigma' = \cup_{d\in D}\{\sigma'' \mid \sigma[x/d][\phi]\sigma''\}. \end{array}$$

[11] As in Dekker 1993 and Heim 1982, variables cannot be reset because resetting variables would involve losing information about their previous values. This 'downdate' effect would be problematic for the notions of negation and support, which being defined in terms of survival depend for their significance on the fact that no operations are considered that cause loss of information. There are other means though to avoid the 'downdate' problem which allow reuse of variables, for instance the assumption of referent systems as in Groenendijk et al. 1996. As an aside, notice that once we assume the style of quantification I propose, we can reformulate the semantics in such a way that downdates are no longer problematic (cf. van Eijck and Cepparello 1994). Finally observe that the novelty condition is a source of partiality. In addition to presuppositions and formulas containing free variables, quantified sentences are partial updates as well. Since partiality is irrelevant for the issues discussed in this paper, I will sometimes be less careful about it in what follows.

Universal quantification is defined in the standard way in terms of negation and existential quantification.

Since all of the operations defined are deterministic with the only exception of SL,[12] if either MS or RA are assumed as the interpretation of the existential quantifier, then the whole semantics is deterministic and can be stated in terms of partial functions.

I conclude this section by defining the notion of entailment.

Definition 9 [Dynamic Entailment]

$$\phi \models \psi \ \text{ iff } \ \forall \mathcal{I}, \forall \sigma, \sigma' \in \Sigma_{\mathcal{I}} : \sigma[\phi]\sigma' \Rightarrow \sigma' \models \psi$$

Entailment is defined in terms of support. ϕ entails ψ iff whenever a state is updated with ϕ, then all possible outputs are states that support ψ. We can now turn to the illustrations of the problems.

2 Underspecification and Overspecification

Of the two ways of interpreting variables that play a role in the three styles of dynamic quantification, the one that treats variables as single partial objects is too weak and leads to problems of *underspecification*. The other, which views them as place-holders for a number of alternative total objects, is too strong and leads to problems of *overspecification*.

Underspecification 1

If *quantified* variables are interpreted as *partial* objects, difficulties arise in connection with phenomena that involve quantification into the scope of holistic operators.

the suspect Treating variables in the syntactic scope of a quantifier as single underspecified objects has the unfortunate consequence of rendering the entailment $\exists x \Diamond \phi \models \forall y \Diamond \phi[x/y]$ valid (cf. Dekker 1993). So if we assume RA, sentences like the following:

(1) a. Someone might be the culprit.
    b. $\exists x \Diamond Px$

(2) a. Someone certainly is not the culprit.
    b. $\exists x \neg \Diamond Px$

will contradict each other. However, intuitively (1) and (2) express compatible pieces of information: you may hold the guilt of someone to be consistent with your information state and at the same time have evidence that someone else is innocent. The problem with RA is that the variable x, being introduced via global extension, will denote in both cases exactly the same single underspecified object, which either verifies the modal sentence $\Diamond Px$, or falsifies it.

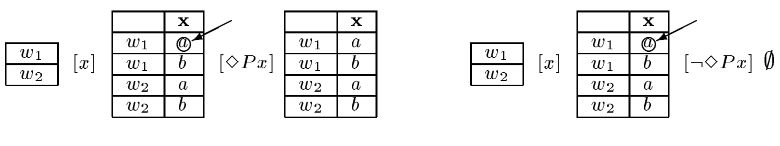

$w_1(P) = \{a\}$

[12] Conjunctions are non-deterministic only if one of the conjuncts is non-deterministic.

If at least one member of the universe has the property P in some world (say individual a in world w_1 as in the picture), (1) is accepted and (2) is rejected. If this is not the case the opposite holds. So (1) and (2) cannot be accepted at the same time. This undesirable result obtains because, to quote from Beaver 1994, in RA, quantified variables don't vary enough: the one value that a variable can take cannot be considered separately from the others, because all the possible values are lumped together. This type of underspecification is also the source of the problem discussed in the following paragraph.

the fat man This problem discussed in Heim 1983 concerns the projection of presuppositions from quantified contexts. Consider (3):

(3) a. A fat man was pushing his bicycle.
b. $\exists x[$ fat-man$(x) \wedge \partial(\exists y$ bike-of$(x,y))\wedge$ pushing$(x,y)]$

Intuitively, (3) projects the presupposition that the relevant fat man had a bicycle.[13] However, Heim 1983, which assigns a random interpretation to variables, predicts the universal presupposition *Every fat man has a bicycle* for sentence (3), which is intuitively too strong.[14] The ∂ clause is interpreted with respect to the state resulting from adding x and updating with fat-man(x). If x is introduced by Random Assignment, this local state may contain several alternative values of x for each surviving world, namely all fat men in that world (a and b in the picture below). If any of these values is not a bike owner, then the ∂ clause turns out undefined. That is, in each world all fat men (all possible values of x) must own a bike, otherwise the sentence is not accepted.

$$\begin{array}{|c|}\hline w_1\\\hline w_2\\\hline\end{array}\ [x]\circ[\text{fat-man}(x)]\ \begin{array}{|c|c|}\hline & \mathbf{x}\\\hline w_1 & a\\\hline w_1 & b\\\hline w_2 & a\\\hline w_2 & b\\\hline\end{array}\ [\partial(\exists y\ \text{bike-of}(x,y))]$$

Like Dekker's problem, Heim's problem results from the fact that in holistic updates all of the values that variables can take are considered all at once instead of one at a time.

Overspecification 1

If we use slicing, the two problems above do not occur.[15] However, the *total* interpretation of *free* variables that SL involves, leads to the loss of a number of attractive properties guaranteed by MS in connection with phenomena of identification in situations of partial information.

[13]In Kartunnen and Peters 1976, (3) is predicted to have the existential presupposition *Some fat man had a bicycle.* This prediction, as the authors admit, is clearly too weak, because intuitively, what should be projected in this case is the presupposition that the *same* fat man that verifies (3) had a bicycle, and *not* some other fat man. The problem arises because in K&P's system there is no obvious way to define scope and binding relations between the presupposition and the assertion, since these two components are represented by two mutually independent propositions. As is well known, in standard satisfaction theory (cf. Beaver 1994, Heim 1983), in which presuppositions are characterized as acceptance conditions, this problem does not occur. Since meanings are not split into two separate (assertive and presuppositional) components, but rather assertions and presuppositions are taken as different aspects of a single dynamic meaning, variables in the latter can be bound by quantifiers in the former and vice versa.

[14]In the same paper, Heim suggests remedying this inadequacy by stipulating the ready availability of an *ad hoc* accommodation mechanism in the evaluation of indefinite sentences.

[15]An alternative solution to underspecification 1 is obtained by defining presupposition (cf. Beaver 1992) and modality (cf. Beaver 1993) in a different way. However, by adopting (moderate) slicing (cf. Beaver 1994 and Groenendijk et al. 1996), we obtain the same results with minor surgery.

the culprit Consider the following examples discussed by Groenendijk et al. in (1996) that involve dynamically bound variables occurring in the scope of Veltman's epistemic operator:

(4) a. Someone did it. It might be you[16]. It might also not be you.
b. $\exists x Px \wedge \Diamond(x = \text{you}) \wedge \Diamond(x \neq \text{you})$

(5) a. Someone did it. It might be anyone.
b. $\exists x Px \wedge \forall y \Diamond(x = y)$

These are coherent pieces of discourse, but if variables range over alternative total objects, they are inconsistent. Take for example (5) which expresses an ultimate form of ignorance about the culprit's identity. If variables are place-holders for individuals, updating with (5) always yields the absurd state since it is impossible for one individual to be (possibly) identical to all the others (if $|D| > 1$). In MS, in which free variables are viewed as partial objects (4) and (5) are instead coherent, as should be the case.

Underspecification 2

The use of moderate slicing avoids the problems noted above, but runs into several others connected with the notions of presupposition, support and coherence. The source of the difficulties here is the *partial* interpretation of *free* variables that MS involves.

the fat man again Heim's fat man problem arises not only for quantified variables, but for free variables as well. As an illustration, consider the following variation of (3), in which the occurrence of the variable x in the ∂ clause is dynamically bound by the existential quantifier:

(6) a. A fat man was sweating. He was pushing his bicycle.
b. $\exists x[\text{ fat-man}(x) \wedge \text{ sweat}(x)] \wedge \partial(\exists y \text{ bike-of}(x,y)) \wedge \text{ pushing}(x,y)$

If we assume a partial interpretation of free variables (RA and MS), then, for the same reasons as above, (6) projects the presupposition that every fat man who was sweating had a bike, which is intuitively too strong.

the wrong suspect Further difficulties for the partial view for free variables arise in connection with the notions of support and coherence. The notion of support (cf. def. 4) can be used to characterize when a speaker is licensed to utter a certain proposition. A speaker is licensed to utter ϕ iff her own information state supports ϕ. As a straightforward generalization we may say that a sentence is assertable iff there is a non-absurd state that supports it. In Groenendijk et al. 1996, texts satisfying this condition are called coherent texts; intuitively, such texts express mutually compatible pieces of information. Now, consider the following example (the pronoun in the second sentence should be read as co-referential to the indefinite in the first sentence):

(7) a. Someone might be the culprit. She is not the culprit.
b. $\exists x \Diamond Px \wedge \neg Px$

Intuitively (7) cannot be coherently asserted as a continuous monologue. The first and second sentence express incompatible pieces of information. You cannot hold the guilt of a person to be consistent with your information and at the same time

[16] In this example, the deictic pronoun *you* is assumed to be epistemically rigid.

have the information that the same person is innocent. But if we use MS (or RA) and treat free variables as denoting single partial objects, (7) surprisingly comes out coherent, i.e. there are states that support it. Take as input a state σ^* consisting of two possibilities that supports the information that either individual a (in w_2) or individual b (in w_1) is P. In MS (as in RA) that allows a partial interpretation of free variables, the first conjunct leads to a state with four possibilities in which both a and b are assigned as possible values to x for each world. Updating with the second conjunct keeps only those two possibilities that assign to x the individuals that are not P:

$$\begin{array}{|c|}\hline w_1 \\ \hline w_2 \\ \hline\end{array} \quad [\exists x \Diamond Px] \quad \begin{array}{|c|c|}\hline & \mathbf{x} \\ \hline w_1 & a \\ \hline w_1 & b \\ \hline w_2 & a \\ \hline w_2 & b \\ \hline\end{array} \quad [\neg Px] \quad \begin{array}{|c|c|}\hline & \mathbf{x} \\ \hline w_1 & a \\ \hline w_2 & b \\ \hline\end{array}$$

Even though the latter update eliminates possibilities, both possibilities in the initial state survive in the final state. So σ^* supports the sequence and hence the latter is coherent. It is impossible, however, for a state resulting from a successful update with the first sentence in (7) to support the second one. The fact that (7) still comes out coherent shows the 'non-compositionality' of the notion of support: we may have a state that supports a conjunction, whereas the same state updated with the first conjunct does not support the second one. Our notion of support predicts that a speaker who is licensed to assert $\phi_1 \wedge \phi_2$ as a whole, is not necessarily licensed to assert ϕ_2 after asserting ϕ_1 and this is counter-intuitive.[17]

To summarize, in both MS and RA, in which free variables are interpreted as single partial objects, texts like (6) are predicted to project too strong universal presuppositions, texts like (7) come out counter-intuitively coherent and, connected with this, we have a 'non-compositional' notion of support.

Overspecification 2

The *total* interpretation of *quantified* or free variables hides the conceptual presupposition that there exists a unique method of individuation across the boundaries of our epistemic possibilities. In Groenendijk et al. 1996, the total objects in an information state are taken to represent the ordinary individuals the agents are acquainted with. In particular, they are specified as objects of perception. Given our trust in our perceptual capacities, it is quite reasonable to assume that if an individual is standing in front of us, then the *same* individual will be standing in front of us in all our epistemic alternatives. So demonstrative identification as opposed to descriptive identification is suggested as the unique method of cross-identification and direct reference is specified as reference under such a perspective. The problem with this characterization is that it fails to account for phenomena of identity and identification in situations of partial or mistaken information, that are precisely the kind of phenomena that quantified epistemic logic should account for.

the man with the hood Suppose a man with a hood is standing in front of you and you haven't the faintest idea who he is. Groenendijk et al. have no obvious way of expressing this uncertainty. The following natural candidate, for instance, comes out inconsistent:

[17] Cf. Dekker 1997, in which this problem is solved by introducing a new notion of support. All underspecification problems can be solved in RA by adopting different analyses for the three holistic notions. However, if underspecification can be avoided by simply using another style of quantification, then by dropping RA we account for three groups of phenomena with a single move.

(8) $\neg\exists x \Box (x = \text{this})$

Since variables range over total objects, (8) is accepted in a state iff the semantic value of the demonstrative is *not* a total object. The problem is that, if demonstrative identification is taken as the unique individuation method, in all your epistemic alternatives, that same man is standing in front of you with a hood on his head, and so by definition you have identified him.

We could conclude that demonstrative identification was not the right notion and that we should investigate further to find a more adequate one. However, it is easy to see that this would not be the right way to go. Similar problematic cases can be constructed for any other possible characterization of the notion of direct reference. Direct reference cannot be characterized as reference to the objects individuated by the *one* favored mode of presentation, because there is not such a unique favored perspective. The following example supplies evidence for this point.

the soccer game Suppose you are attending a soccer game. All of the 22 players are in your perceptual field. You know their names, say a, b, c, ..., but you don't recognize any of them. Consider the following sentence:[18]

(9) a. Anyone might be anyone.
b. $\forall x \forall y \Diamond (x = y)$

It seems to me that (9) can be uttered in this situation. However, if we assume (moderate) slicing, (9) is inconsistent. The source of the difficulty is the uniqueness presupposition behind the total interpretation of variables. Intensional properties such as 'possibly being anyone' are not traits of individuals *simpliciter*, but depend on the perspective under which these individuals are looked at. Examples like (9) show that there is not one direct way of looking at the universe of discourse that characterizes the domain of quantification once and for all, rather different perspectives supply different sets of ultimately partial objects over which we can quantify.

Synopsis The diagram below summarizes the contents of this section:

	RA	**SL**	**MS**
quant. var.	partial $\rightsquigarrow$ undersp 1	total $\rightsquigarrow$ oversp 2	total $\rightsquigarrow$ oversp 2
free var.	partial $\rightsquigarrow$ undersp 2	total $\rightsquigarrow$ oversp 1	partial $\rightsquigarrow$ undersp 2

3 Quantification under Conceptual Covers

In order to overcome the problems of over- and underspecification, I propose a new style of dynamic quantification that lies between random assignment and slicing, and which treats free and quantified variables in a uniform way. As in slicing, the interpretation proceeds on different parallel levels so that free and quantified variables range over alternative definite elements of some domain. In this way, variables vary enough to avoid the underspecification problems. On the other hand, the overspecification problems are solved by allowing not one but many ways of conceiving the individuals over which we quantify. Different sets of possibly non-rigid concepts that cover the whole universe and do not consider any individual more than once can constitute a suitable candidate for the domain of quantification.

Definite Subjects

In dynamic semantics, two levels of objects are assumed: the individual elements of the universe of discourse, and the partial entities that constitute the interpretations

[18] This is a modification of an example of Paul Dekker.

of variables in information states. The latter are introduced as items in conversation and can change, for instance by growing less partial, as the conversation proceeds. As we saw,[19] in Dekker 1993, these entities are called partial objects and are defined as functions that assign to each possibility in a state the value of the corresponding variable in that possibility. I extend Dekker's definition of partial objects and call a *subject* in an information state any mapping from the possibilities (world-assignment pairs) in the state to the individuals in the universe of discourse. Notice that in addition to explicitly introduced discourse items, potential items also count as subjects in a state.

Among the subjects, we can distinguish *rigid subjects* and *(in)definite subjects.* Rigid subjects are the constant functions among the subjects. Definite subjects are those that assign the same value to all possibilities that have the same world parameter. Definite subjects are contextually restricted (individual) concepts. They are *definite* in that they have a single value relative to a single world, but *partial* in that they may have different values relative to different worlds. Indefinite subjects are subjects that are not definite, i.e., those assigning different values to possibilities with the same factual content.

In RA and MS, the presence of indefinite subjects as interpretation of some discourse items in a state reflects the indeterminacy of the addressee's perspective. Note that from the speaker's point of view indefinite items are senseless.[20] Consider the following dialogue (Dekker):

K: Yesterday a man came into my office who inquired after the secretary's office.
J: Was he wearing a purple jogging suit?
K: If it was Arnold, he was, and if it was somebody else, he was not.

If we assume that K knows that Arnold *and* somebody else went to his office inquiring after the secretary's office, then K's reply is odd, because K should have made up his mind about whom he wanted to talk before starting to tell the story. But now imagine another scenario: assume that K, who is blind, but knows that Arnold always wears eccentric jogging suits, was wondering from the beginning whether it was Arnold who went to his office *or* somebody else. Then the dialogue becomes quite natural.

Speakers do not introduce indefinite subjects[21] (in the first scenario the dialogue is odd), but may introduce non-rigid subjects (in the second scenario the dialogue is natural). Speakers introduce definite subjects. Now, dynamic semantics models the addressee's updating procedure and addressees often lack information about which definite subjects speakers intend to refer to. Questions like *Who do you mean?* or *Who are you talking about?* represent states of ignorance of this kind. In RA and MS, this ignorance is modeled in the same way as ordinary ignorance about what is the case, that is by the presence in the information state of a number of world-assignment pairs in which the different individuals that the speaker might have in mind are represented by the different values that the relevant variable can take. A consequence of this strategy is the presence in states of indefinite subjects. In SL, instead, (lack of) information about the speaker's intentions is modeled on a higher level, namely by the presence of different alternative updates that run in a parallel fashion. Here the possible speaker's referents are modeled by the rigid subjects that constitute the interpretation of the relevant variable in the alternative parallel states. Going back to the dialogue above, in cases in which K is assumed to

[19] Cf. footnote 1.
[20] Cf. Dekker 1997 & van Rooy 1997.
[21] Of course there are also 'non-specific' indefinite NPs, but in that case no discourse referents accessible for inter-sentential anaphora are introduced. Note that in this framework the contrast between specific and non-specific indefinites is explained as a scope ambiguity.

be omniscient with respect to information about the world, K's reply is intuitively unacceptable, but it is not obvious how we can account for this intuition, if the item introduced by K is modeled by an indefinite subject (RA and MS). On the other hand, if we avoid indefinite subjects, but admit only rigid subjects (SL), then K's reply is never judged as acceptable, which is also incorrect.

I propose to let variables range over definite subjects. The interpretation of the existential quantifier will involve splitting up a state as in the slicing procedure. The definite subjects (possibly non-rigid ones) which the speaker might have in mind are considered one by one as disjoint alternatives.

	x
w_1	a
w_1	b
w_2	a
w_2	b

1. partial

	x
w_1	a
w_2	a

	x
w_1	b
w_2	b

2. total

	x
w_1	a
w_2	a

	x
w_1	a
w_2	b

	x
w_1	b
w_2	a

	x
w_1	b
w_2	b

3. definite

It seems that in this way we can avoid underspecification without falling into overspecification. Since variables are taken to range over alternative elements of some domain, we avoid Dekker's or Heim's problem. In addition, since they can vary over non-rigid subjects, we have a good hope of solving the overspecification problems as well. Definite subjects seem to be the "something in between" that we were looking for. However, quantification over concepts is quite an intricate affair. Difficulties arise almost immediately from the fact, evident from the picture above, that there are strictly more concepts in a state than individuals in the universe of discourse.

the winner It is easy to show that if we let quantifiers range over the set of all definite subjects, the semantics so obtained will validate the following scheme:

(10) $\forall x \Diamond \phi \rightarrow \Diamond \forall x \phi$

which is clearly undesirable.[22] Suppose a game has been played; (10) says that if it is known that there are some losers ($\neg \Diamond \forall x W x$), but we have no clue about who won, we have no way of expressing this ignorance (since $\neg \forall x \Diamond W x$). Another example showing the same point is the smallest flea case.

the smallest flea Consider the following two sentences:

(11) Any flea might be the smallest flea.

(12) The biggest flea might be the smallest flea.

[22] Quine, though discussing a different point, shows the implausibility of the equivalent scheme $\Box \exists x \phi \rightarrow \exists x \Box \phi$: "...in a game of a type admitting of no tie it is necessary that some one of the players will win, but there is no one player of whom it may be said to be necessary that he win." Quine 1953, p. 148.

If we quantify over all concepts, a generalized version of universal instantiation holds and we can derive (12) from (11). This means that in ordinary situations in which fleas differ in size, (11) is never accepted. There will always be an element in the quantificational domain that falsifies it, for instance *the biggest flea*. Thus ignorance about the smallest flea's identity is inexpressible in such situations.

The examples above seem to show that quantifiers in natural language do not range over representations of individuals. If in sentences like (11) we were quantifying over representations, then we would have to accept the derivation of (12) from (11) as a trivial one. The fact, instead, that this conclusion strikes us as counter-intuitive means that natural language quantifiers do not work in this way. When we talk, we talk about individuals, not about representations of individuals, even in situations in which we lack information or are misinformed about them. To capture this feature of natural language quantifiers, we need a notion of *aboutness* which can work in situations of *partial information*. The traditional characterization of aboutness in terms of rigidity, implicit in (moderate) slicing, is inadequate in these cases. As we saw, in situations of partial information, we do not (because we cannot) quantify over total objects. However, to deny the claim that quantifiers range over individuals in a direct way, we need not assume that we quantify *over* representations - it is enough to say that we quantify over individuals, but *under* a representation. Natural language quantifiers range over individuals under a *perspective*. To give some content to this abstract claim, let's consider the following example[23] in which we see perspectives at work.

the butler Suppose a butler and a gardener are sitting in some room. One is called Alfred and the other Bill. We don't know who is who. In addition, assume that the butler committed a terrible crime. Consider now the following two discourses:

(13) The gardener didn't do it. So it is not true that anybody (in the room) might be the culprit.

(14) Alfred might be the culprit. Bill might be the culprit. So anybody (in the room) might be the culprit.

It seems to me that both (13) and (14) can be uttered in such a situation given the right circumstances. We can intuitively explain what is going on as follows: intensional properties such as perhaps being the culprit are not traits of individuals *simpliciter*, but depend on the perspective under which these individuals are conceived. In the two discourses, the universal quantifier though ranging over the same set, namely the set containing the two people in the room, identifies them from two different angles and for this reason no contradiction arises. In (13), individuals are looked at under the perspective of their professions; in (14) they are identified as bearers of some proper name. Under the latter identification mode, *the butler* may be Alfred or may be Bill. Yet if we assume the former perspective, we can think of *the butler* as standing for a single object contrasted with *the gardener*. Perspectives are determined by contextual factors. In these two specific cases, the relevant contextual information is supplied by the preceding sentences, which, by mentioning one concept or the other, suggest one or the other way of classifying the domain.

A natural way of representing a perspective over the universe of discourse is by means of a set of concepts. However, not all collections of concepts will do. The set of *all* concepts, for instance, is not a good candidate, as is evident from the winner and the smallest flea examples above. But there are many more inadequate conceptualizations.

[23] For more about examples of this kind cf. Gerbrandy 1999.

Take a situation similar to the one above. Again we have Alfred and Bill sitting in some room, we know that one of the two is the butler, the other is the gardener, but we don't know who is who. Suppose you are interested in determining whether for anybody in the room it is consistent with your information that an arbitrary property, say being bald, holds.

(15) Anybody might be bald.

Which of the following sentences constitutes a sufficient ground for a correct assertion of (15)?

(16) Alfred might be bald and Bill might be bald.

(17) The gardener might be bald and the butler might be bald.

(18) Alfred might be bald and the butler might be bald.

(19) Bill might be bald and the gardener might be bald.

In this particular situation, only the first two can ground (15). A derivation of (15) from either (18) or (19), would not be accepted as an example of correct reasoning. Even if explicitly suggested by the context, the sets consisting of *Alfred* and *the butler* or of *Bill* and *the gardener* are not good conceptualizations in this specific case.[24] The reason for this is that, intuitively, they do not provide a uniform perspective over the universe of discourse: they mix up different perspectives and they do not cover the domain of individuals in an exhaustive way. In the following section a way to formalize these intuitions is proposed.

Conceptual Covers

A conceptual cover is a set of individual concepts that satisfies two conditions: exhaustivity and disjointness.

Definition 10 Given an information space $\mathcal{I} = \langle W, D \rangle$, a conceptual cover CC over $\mathcal{I}$ is a set of functions $W \to D$ such that:

$$\forall w \in W \; \forall d \in D \; \exists! c \in CC : c(w) = d$$

In a conceptual cover, each individual d is 'seen' by at least one concept in each world (exhaustivity), but in no world is an individual counted more than once (disjointness).

Since conceptual covers are sets of concepts which exhaustively and exclusively cover the domain of individuals, conceptual covers and the domain of individuals have the same cardinality.

Fact 1 Let $\langle W, D \rangle$ be an information space. For any conceptual cover CC over $\langle W, D \rangle$, it holds that $|CC| = |D|$.

Irrespective of which perspective you assume, the number of individuals in the domain doesn't change. Given this result, we can say that conceptual covers are not simply sets of representations, but constitute a proper perspective over the universe of individuals. Different conceptual covers constitute different ways of perceiving one and the same domain.

Two typical examples of conceptual covers are the following (let $\mathcal{C}$ be the set of individual constants in $\mathcal{L}$):

[24] Their inadequacy doesn't follow from the fact that they use definite descriptions *and* proper names, but depends on the specific information supported in this case. In other situations, such sets can provide good conceptualizations.

1. $R = \{\lambda w\ d \mid d \in D\}$ (rigid cover)
2. $N = \{\lambda w\ w(a) \mid a \in \mathcal{C}\}$ (naming)

R is the set of constant concepts and N is the set of concepts that assign to every world the denotation of a certain individual constant in that world (assuming exhaustive and exclusive naming practices). These two covers can be taken to model the two identification modes that played a role in most of the examples above (cf. for instance (9)), namely identification by ostension and identification by naming.[25] However, these are just two among the many modes of individuation that we normally assume when we think or talk about objects in our everyday practices. Other families of individuation modes are for instance individuations by description as in example (13) or by recognition like in the cases in which we identify strangers by bringing to mind the visual image of their faces that we perceived at one time. Our theory has no problem providing enough conceptual covers to model this multitude of identification modes.[26]

I propose to let variables range over the elements of a contextually supplied conceptual cover. The existential and the universal quantifiers will behave as ordinary quantifiers, that is, even if, technically, they range over concepts, the effect obtained is that of quantification over genuine individuals. It is the insistence on this normal sense of quantifiers that motivates the two constraints on conceptual covers specified above, in particular the disjointness condition which serves to guarantee that the objects over which we quantify really specify determined individuals that can be said to be identical with themselves and distinct from one another. Non-disjoint sets of concepts do not characterize sets of genuine individuals in this sense. Consider again the situation described in the butler example above. Alfred and Bill are sitting in some room, we know that one of the two is the butler and the other is the gardener, but we don't know who is who. Take the set A consisting of the concepts *Alfred* and *the butler*. First of all, observe that A is not a conceptual cover. Given our assumptions, there will be some world w, in which someone is counted twice (namely the individual which is Alfred and the butler in w), and someone else is not 'seen' at all (namely the individual which is Bill and the gardener in w). So A is neither disjoint nor exhaustive. Now, given the situation, the two elements of A cannot be regarded as standing for two determined individuals. Since we wonder whether Alfred is the butler, *Alfred* and *the butler* might be one individual or two. On the other hand, inside exhaustive and disjoint sets of concepts, this kind of indeterminacy doesn't arise. Consider the set B consisting of *the butler* and *the gardener*, which, given our assumptions, is a conceptual cover. The concept *the butler*, which when taken in combination with *Alfred* gave rise to individuation problems, here, contrasted with *the gardener*, comprises a completely determined individual. Thus, only as an element of B and not as an element of A, *the butler* is capable of serving as value of some bound variable.

To conclude, the elements of a conceptual cover represent the entities we quantify over, that we experience only via one or the other mode of presentation; yet it would be misleading to identify them with these modes of presentation.[27] The elements of a conceptualization are the individuals themselves just thought, conceived, identified in a particular way.

---

[25] Here I assume that the rigid cover represents demonstrative identification. However, since the notion of rigidity plays hardly any role once quantification is relativized to conceptualizations, I could have chosen otherwise and it would not have mattered.

[26] There are $(|D|!)^{|W|-1}$ conceptual covers over $\mathcal{I} = \langle D, W\rangle$.

[27] Or with Fregean senses characterized as ways of thinking the referent of some singular term.

Quantification under Conceptual Covers

To define quantification under conceptual covers, I need an operation that extends information states in the appropriate way.

Definition 11 [c-extension] Let $\sigma \in \Sigma_{\langle D,W\rangle}, x \in \mathcal{V}\backslash\mathrm{dom}(\sigma)$ & $c \in D^W$.

$$\sigma[x/c] = \{\langle w,g\rangle[x/c(w)] \mid \langle w,g\rangle \in \sigma\}$$

C-extensions lie between global and individual extensions. They introduce fresh variables and interpret them as certain definite subjects. Dynamic quantifiers are defined in terms of c-extensions; they range over elements of a contextually-given conceptual cover and only indirectly over the individuals in the universe. In this way, quantification is relativized to a particular way of conceptualizing the domain. Dynamic quantification is defined as follows:

Definition 12 [Quantification] $\sigma[\exists_Z x\phi]_\alpha\sigma'$ iff $\sigma[x/c][\phi]_\alpha\sigma'$ for some $c \in \alpha(Z)$.

Z is a CC-index, that is a free variable ranging over conceptual covers. The fact that each quantifier occurs with its own index allows different quantifiers to range over different sets of concepts. The parameter α represents the pragmatic context and is a function from CC-indices to conceptual covers. The interpretation function $[\bullet]$ maps formulae ϕ to their 'characters', i.e. to functions $[\phi]$ from contexts α to relations over information states $[\phi]_\alpha$.[28] By taking α-functions as parameters of the update functions, I avoid the important issue of how conceptual covers are pragmatically determined. The investigation of these selection processes and their constraints must be left for another occasion.

Quantification under conceptual covers formalizes the intuitive idea that quantifiers in natural language range *over* individuals *under* a perspective. It is easy to show that which perspective you choose only plays a role in certain circumstances, namely in situations of partial information for sentences that involve quantification into holistic operators or introduce some new discourse item. These are typically the constructions in our formal language that are used to represent linguistic phenomena involving some notion of aboutness, such as *de re* attitude attributions, knowing-who constructions and specific uses of indefinite NPs. When we talk *about* individuals in situations of *partial* information, we do it under a conceptualization.

4 Solutions

In this section, I will illustrate the role played by conceptual covers in the solution of the problems discussed earlier in this paper.

Underspecification Since variables in quantified contexts are taken to range over alternative definite objects, underspecification **1** does not occur. I will just consider Dekker's problem. Sentences (1) $\exists_Z x\Diamond Px$ and (2) $\exists_Z x\neg\Diamond Px$ do not contradict each other because different definite subjects are not absorbed into single indefinite ones (as in RA), but can be considered in isolation. For instance, let σ be a state and c_1, c_2 be two concepts in some conceptual cover CC such that only c_1 takes as values individuals that have the property P in some possibilities of σ. Such a state σ will support both (1) and (2) since $\sigma[x/c_1] \models \Diamond Px$ and $\sigma[x/c_2] \models \neg\Diamond Px$.

[28] In this way, the contextual dependency of indexed quantifiers is kept separate from that of anaphoric expressions. In the future, we may want to change this and account for the 'dynamics' of shifting conceptual covers by encoding in the states the information about the active conceptualizations and by incorporating in the meanings of expressions their potential to activate conceptualizations.

Underspecification **2** is also avoided, since only concepts may be introduced as new items. I just illustrate how the wrong suspect case is handled in the new system. Intuitively, example

(7) $\exists_Z x \Diamond Px \wedge \neg Px$

comes out incoherent because there are no possible concepts under any conceptualization that can satisfy the two conjuncts at the same time. Formally the incoherence of (7) follows from the fact that it is impossible for a state resulting from a successful update with $\exists x \Diamond \phi$ to support $\neg\phi$, in combination with the compositionality of support that follows from the following property of the new system (proof by induction on the complexity of ϕ):

Fact 2 [Non-Branching] $\forall \phi$, $\forall \alpha$, $\forall \sigma, \sigma' \in \Sigma_{\mathcal{I}}$:

$$\sigma[\phi]_\alpha \sigma' \implies \forall i \in \sigma : \forall j_1, j_2 \in \sigma' : i \prec j_1 \;\&\; i \prec j_2 \rightarrow j_1 = j_2$$

For any update in a non-branching system no two possibilities in the output state can extend one and the same possibility of the input state. Typical examples of systems in which the non-branching property does not hold are RA and MS, namely systems that allow a partial interpretation for free variables.

As an intuitive illustration of how the wrong suspect problem is solved, I will compare the interpretation procedures for (7) in MS, which allows branchings, and in the new system, in which the non-branching property holds. Let the input state be the state σ^* as above consisting of two possibilities that supports the information that either individual a (in w_2) or individual b (in w_1) is P. As we saw, in MS, that allows a partial interpretation for free variables, σ^* supports the sequence and hence the latter is coherent.

MS:

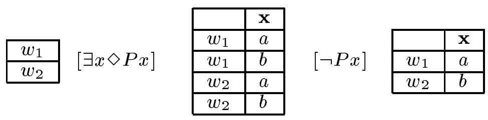

If on the other hand we adopt the style of quantification that I am proposing, we avoid this problem: the two initial possibilities do not survive together in any of the output states under any conceptual cover. In the picture below, we consider as an illustration the case in which $\alpha(Z) = R$.

NEW:

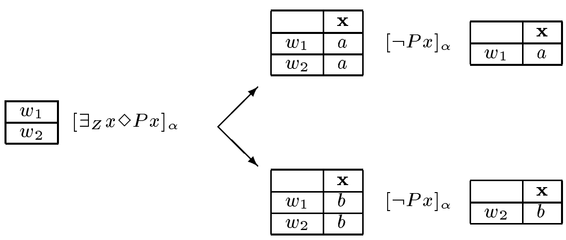

As a matter of fact, no state can be found that supports (7) under any conceptualization.

Overspecification Since variables are not taken to range over individuals *simpliciter*, but under a conceptualization, the overspecification problems are now solved. The inequalitarian attitude towards modes of individuating objects implicit in (moderate) slicing is overcome and different identification modes are given equal status. We can look at the individuals in the universe under different perspectives and, if the context justifies it, we can change perspective within the same discourse. Problems of identification can be represented as problems of mapping elements from different conceptualizations onto each other. Overspecification **1** cases are solved. Examples like

(5) $\exists_Z x Px \wedge \forall_Y y \Diamond(x = y)$

come out coherent, if interpreted in a context α that assignes different conceptual covers to Z and Y. E.g., if we let the existential quantifier introduce non-rigid subjects and the universal range over the rigid conceptualization ($\alpha(Y) = R$), a state like σ^* above supports (5).

Overspecification **2** is avoided in a similar way. Example

(8) $\neg\exists_Z x \Box(x = \text{this})$

can be accepted, if x is not taken to range over the cover representing demonstrative identification, R. A typical case is when Z is assigned naming. Thus we can express ignorance about the identity of some object of perception, and in addition, by shifting conceptualization we can account for any situation of partial identification in an enlightening way. Examples like *I wonder who Alfred is*, or *I wonder who the culprit is* are not problematic for this approach. The soccer game case is explained as well.

(9) $\forall_Z x \forall_Y y \Diamond(x = y)$

(9) is acceptable, if x and y are taken to range over different conceptualizations. In this specific situation, Z and Y are assigned as value respectively the cover by perception and the cover by names. Notice, however, that $\forall_Z x \forall_Z y \Diamond(x = y)$ is inconsistent.

Cardinality Since the set of all concepts is not among the conceptual covers, the winner problem and the smallest flea problem do not occur. The problematic scheme

(10) $\forall_Z x \Diamond\phi \rightarrow \Diamond\forall_Z x \phi$

is *not* valid and so sentences like *Anyone might be the winner* can be accepted in situations in which it is known that there are some losers. Furthermore, only a restricted version of universal instantiation holds:

(20) $\forall_Z x \phi \wedge \exists_Z y \Box(t = y) \rightarrow \phi[x/t]$

So, since the universal sentence *Any flea might be the smallest flea* can be accepted only under a conceptualization that does not contain *the biggest flea*,[29] the problematic implication to *The biggest flea might be the smallest flea* is blocked.

To summarize, if we assume quantification under conceptualizations, underspecification does not occur since only definite subjects may constitute interpretations of variables. At the same time, overspecification is also avoided since different occurrences of quantifiers may range over different sets of (possibly partial) concepts. Finally, by taking as domain of quantification only sets of concepts which exhaustively and exclusively cover the universe of individuals, we avoid the cardinality problems that normally arise when we quantify over concepts rather than objects.

[29] Unless we have a domain with a single flea.

5 Conclusion

The combination of dynamic quantification with holistic notions is a dim affair, because it adds to the obscurity of quantification into modal contexts[30] problems typical of dynamic environments. In this paper, I have tried to show that by bringing conceptual covers into the picture, we don't add obscurity to obscurity, but we shed some light on these difficult issues.

Acknowledgments

I am grateful to many people for conversations and discussions about this material, in particular David Beaver, Paul Dekker, Jelle Gerbrandy, Jeroen Groenendijk and Robert van Rooy. Anonymous reviewers helped to improve the final version.

References

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

Beaver, D.: 1994, When variables don't vary enough, in L. Santelmann and M. Harvey (eds.), *Proceedings of the Fourth Semantics and Linguistic Theory Conference*, DMLL Publications, Cornell University

Beaver, D.: 1995, *Presupposition and Assertion in Dynamic Semantics*, Ph.D. thesis, University of Edinburgh, To appear: Studies in Logic, Language and Information, CSLI, Stanford, CA.

Beaver, D. I.: 1992, The kinematics of presupposition, in P. Dekker and M. Stokhof (eds.), *Proceedings of the Eighth Amsterdam Colloquium*, ILLC, University of Amsterdam, Amsterdam

Beaver, D. I.: 1993, Kinematic Montague grammar, in H. Kamp (ed.), *Presupposition*, ILLC, University of Amsterdam, Amsterdam, Dyana-2 deliverable R2.2.A, Part II

Dekker, P.: 1993, *Transsentential Meditations. Ups and Downs in Dynamic Semantics*, Ph.D. thesis, ILLC/Department of Philosophy, University of Amsterdam, Amsterdam

Dekker, P.: 1994, Predicate logic with anaphora (seven inch version), in L. Santelmann and M. Harvey (eds.), *Proceedings of the Fourth Semantics and Linguistic Theory Conference*, p. 17, DMLL Publications, Cornell University

Dekker, P.: 1997, On first order information exchange, in A. Benz and G. Jäger (eds.), *Mundial '97*, München

Gerbrandy, J.: 1999, Identity in epistemic semantics, in *this volume*

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

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

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

Heim, I.: 1983, On the projection problem for presuppositions, in M. Barlow, D. Flickinger, and M. Wescoat (eds.), *Proceedings of WCCFL II*, Stanford, California

Kartunnen, L. and Peters, S.: 1976, Conventional implicatures in montague gram-

[30] Cf. the concluding paragraph of Quine 1956.

mar, in C. Oh and D. Dineen (eds.), *Syntax and Semantics 11: Presupposition*, Academic Press, New York

Quine, W. V.: 1953, *From a Logical Point of View*, Chapt. VII Reference and modality, Harvard University Press

Quine, W. V.: 1956, Quantifiers and propositional attitudes, *Journal of Philosophy* 53, 101–111, Reprinted in: W.V. Quine, 1966, *The Ways of Paradox and Other Essays*, Random House, New York

van Eijck, J. and Cepparello, G.: 1994, Dynamic modal predicate logic, in M. Kanazawa and C. Piñon (eds.), *Dynamics, Polarity and Quantification*, CSLI, Stanford

van Rooy, R.: 1997, *Attitudes and Changing Contexts*, Ph.D. thesis, University of Stuttgard

Veltman, F.: 1997, Defaults in update semantics, *Journal of Philosophical Logic*