The values of variables in dynamic semantics

Paul Dekker

To be is to be the value of a variable.
(Willard V. Quine)

1 Introduction

Although the development of *Discourse Representation Theory* (*DRT*, Kamp 1981; Kamp and Reyle 1993) in the early eighties has been prompted by a variety of empirical and theoretical considerations, those concerning the interpretation of anaphoric relationships in natural language are probably the most well-known. In the rise of the *DRT* paradigm, these relationships have been used to argue against compositional model-theoretic theories of natural language interpretation, and to motivate a theory of interpretation which encompasses a level of representation. In Groenendijk and Stokhof 1991, this motivation for a non-compositional level of representation has been criticized. Groenendijk and Stokhof present a non-representational, dynamic semantics for the language of predicate logic (*DPL*), characterized by a notion of meaning by means of which anaphoric relationships can be accounted for along the lines of traditional theory. Thus, they show that the interpretation of these relationships require a richer notion of meaning, rather than a different architecture of the theory of interpretation. As a matter of fact, something similar had already been implicitly shown with Heim's file change semantics (*FCS*, Heim 1982; Heim 1983). Although wrapped up in the representational metaphor of changing files and file cards, Heim presents the semantics of a language of logical forms which also deals with intersentential anaphoric relationships in non-representational terms. With this paper I want to combine the relative merits of *DPL* and *FCS*, and study a dynamic semantics for a language of predicate logic with a Heimian notion of meaning.

In the system of *DPL* sentences are interpreted as functions over a domain of information states. Since the primary target is the interpretation of anaphoric relationships, the type of information dealt with concerns the possible values of terms which at various stages in a discourse may serve as antecedents for anaphors later to come. This type of information is rendered as information about the values of variables. The idea is that a sentence which contains a potential antecedent term—typically, an indefinite noun phrase—shows up as an existential quantified formula $\exists x\phi$ in (predicate) logical form, and that the possible values of x under which ϕ is satisfied are passed on in the interpretation of subsequent discourse. A subsequent formula with a free occurrence of the same variable x can then be interpreted as coreferential with the preceding quantifier, and, thus, such a variable turns out to behave as the logical counterpart of a resolved anaphoric pronoun.

DPL and *FCS* both model the interpretation of anaphoric relationships using a notion of information about the values of variables (discourse markers, natural numbers, indices, ...), but *FCS* can be taken to show that there are two relevant aspects of this kind of information. First there is information about the variables the values of which are under discussion, and then there is information about their values. This first aspect of information, which *DPL* fails to deal with, is crucial to Heim's (un-)familiarity theory of (in-)definiteness. She judges it 'infelicitous' or 'inappropriate' to associate an indefinite noun phrase with a variable which already is in the domain of the information state which makes up the context of interpretation.

This paper serves to investigate what the adoption of Heim's felicity conditions amounts to semantically. It may be clear beforehand what (interpreted) sub-language is delineated when they are incorporated in a (dynamic) predicate logical language, but it is quite another question whether the delineated sub-language has a particularly interesting semantics. The answer given in this paper is that it has. The underlying information structure is shown to be a lattice, partially ordered by a so-called 'update'-relation. Further study reveals that the structure comes with a completely sensible, semantic, notion of a discourse referent. Furthermore, the structure helps us do away with a most unpleasant and debated feature of *DPL*, namely, its so-called 'down-date' property. Finally, it will appear that the dynamics of interpretation upon its Heimian reformulation can be located precisely in the conventions governing a proper use of variables. The corresponding inference system turns out to be pretty neat and well-behaved.

The paper is organized as follows. The next section sketches *DPL* and motivates the use of Heimian information states. In section 3 *EDPL* is presented, the Heimian system of interpretation for a language of predicate logic. Section 4 shows that the system is easily extended with a uniform account of adverbial and adnominal quantification.

2 On variables

2.1 Dynamic predicate logic

With *DPL*, Groenendijk and Stokhof propose a dynamic interpretation of the language of predicate logic which models, among other things, the compositional interpretation of intersentential anaphoric relationships in natural language. As in *DRT* and *FCS*, natural language noun phrases are associated with variables in *DPL*, and information states determine what values the variables can have given the conditions imposed on them in the course of a text. The *DPL* system of interpretation keeps track of the possible values of these variables in an on-line manner.

Consider the following two-sentence sequence which displays two intersentential anaphoric relationships:

(1) A cowgirl meets a boy. She impresses him.

Let us assume that, for the interpretation of the first sentence *A cowgirl meets a boy*, we have associated *a cowgirl* with a variable x and *a boy* with a variable y. Then the interpretation of that sentence will make us end up in a state which contains the information that the value of x is a cowgirl who meets a boy who is the value of y (if there is such a cowgirl, that is; otherwise the sentence is just false). This state is precisely the kind of state we need to interpret a continuation with *She* *impresses him*, that is, if *she* is associated with x again, and *him* with y. This second sentence then adds the information that the value of x impresses the value of y, and the state that results from interpreting the sequence of two sentences contains the information that the value of x is a cowgirl who meets and impresses a boy who is the value of y. Thus, the coreferential links between *a cowgirl* and *she*, and between *a boy* and *him* are accounted for semantically.

Information about the values of variables is cast in terms of (sets of) variable assignments in *DPL*. So, for concreteness' sake, let us suppose that, as a matter of fact, Alice is a cowgirl who meets a boy, Boris, and that Carol is also a cowgirl, who meets boy Don. In that case, and with the variables associated as indicated, the interpretation of the first sentence *A cowgirl meets a boy* yields an information state s which contains assignments g such that $g(x)$ is Alice and $g(y)$ is Boris, and assignments h such that $h(x)$ is Carol and $h(y)$ is Don. If we also suppose that Alice impresses Boris, but Carol doesn't impress Don, then the interpretation of the second sentence with respect in this state s will eliminate (the possible values of variables determined by) assignments h.

This more or less informal description of the *DPL* treatment of anaphora is worked out in full formal detail in a way which is presented shortly. Before, I have to point out that, in my presentation of it, the so-called functional formulation of *DPL* is adopted, instead of its (original) relational one. Although the relational one is 'cheaper' in the sense of being lower order, the functional one better matches the intuition that what is at issue is change of states of partial information.

The language of *DPL* is that of predicate logic, but in order to keep matters simple individual constants and function terms are disregarded. The semantics is defined with respect to an ordinary first order model $M = \langle D, F\rangle$ which consists of a non-empty set of individuals D and an interpretation function F which assigns sets of n-tuples of objects to n-ary relation expressions. (Reference to M is omitted whenever this does not lead to confusion.) The interpretation of a formula ϕ in a state s yields a new state, $s[\![\phi]\!]$, as specified by the following definition (here, $g[x/d]$ is that assignment $h = (g \setminus \{\langle x, g(x)\rangle\}) \cup \{\langle x, d\rangle\}$ which agrees with g on the values of all the variables except, possibly, x, and such that $h(x) = d$; $s[x]$ is the set of assignments $\{g[x/d] \mid g \in s \;\&\; d \in D\}$ which agree with an assignment in s on the values of all variables except, possibly, x):

DPL Semantics

- $s[\![Rx_1 \ldots x_n]\!] = \{g \in s \mid \langle g(x_1), \ldots, g(x_n)\rangle \in F(R)\}$
  $s[\![x = y]\!] = \{g \in s \mid g(x) = g(y)\}$
  $s[\![\neg\phi]\!] = \{g \in s \mid \{g\}[\![\phi]\!] = \emptyset\}$
  $s[\![\exists x \phi]\!] = s[x][\![\phi]\!]$
  $s[\![\phi \wedge \psi]\!] = s[\![\phi]\!][\![\psi]\!]$
- s supports ϕ with respect to M, $s \models_M \phi$, iff $\forall g \in s$: $\{g\}[\![\phi]\!]_M \neq \emptyset$
  ϕ is valid, $\models \phi$, iff $\forall M, s$: $s \models_M \phi$

The interpretation of an atomic formula in a state s involves the intersection of s with the set of assignments which verify the formula in a classical sense. The negation of ϕ preserves the assignments in s which falsify ϕ. Conjunction is interpreted as functional composition. A conjunction $\phi \wedge \psi$ is interpreted in s by, first, interpreting

ϕ in s, and, next, interpreting ψ in the state that results. A formula ϕ is supported by a state s iff all possible variable assignments $g \in s$ constitute a state $\{g\}$ with respect to which ϕ can be successfully interpreted.

The characteristic clause of the above definition is concerned with the interpretation of existentially quantified formulas. When we interpret a formula $\exists x\phi$ in a state s, we take into consideration all possible values for x and then interpret ϕ. The mediating state $s[x]$ contains the same information as s about the values of all variables except x, and about the value of x, $s[x]$ has no information whatsoever. The result of interpreting $\exists x\phi$ in a state s, then, is a state which contains precisely the information about the value of x which is expressed by ϕ.

The previous observations about the dynamic potential of *DPL*'s existential quantifier can be brought out using the following equivalences. The first one is classical, but the second one distinguishes *DPL* from static theories of interpretation:

Donkey equivalences (1)

$((\phi \wedge \psi) \wedge \chi) \Leftrightarrow (\phi \wedge (\psi \wedge \chi))$
$(\exists x\phi \wedge \psi) \Leftrightarrow \exists x(\phi \wedge \psi)$

It is typical of *DPL* that the second equivalence also holds if the variable x occurs free in ψ. These equivalences summarize, as it were, the *DPL* treatment of the following textbook example (which explains the equivalences' label):

(2) A farmer owns a donkey. He beats it.
$(\exists x(Fx \wedge \exists y(Dy \wedge Oxy)) \wedge Bxy) \Leftrightarrow$
$\exists x(Fx \wedge \exists y(Dy \wedge (Oxy \wedge Bxy)))$

This sequence of sentences turns out to be equivalent to the sentence *A farmer beats a donkey he owns*. So, although the two sentences *A farmer owns a donkey* and *He beats it* are assigned an interpretation of their own, the intersentential anaphoric relationships (between *a farmer* and *he*, and between *a donkey* and *it*) get established when the two are conjoined.

Given the usual definitions of universal quantification and implication in terms of existential quantification, negation and conjunction, as a corollary of the above equivalence the following also hold (again, the first one is a classical equivalence and the second one typical for *DPL*):

Donkey equivalences (2)

$((\phi \wedge \psi) \rightarrow \chi) \Leftrightarrow (\phi \rightarrow (\psi \rightarrow \chi))$
$(\exists x\phi \rightarrow \psi) \Leftrightarrow \forall x(\phi \rightarrow \psi)$

The second donkey equivalences summarize the *DPL* treatment of the following museum piece donkey sentences:

(3) If a farmer owns a donkey he beats it.
$(\exists x(Fx \wedge \exists y(Dy \wedge Oxy)) \rightarrow Bxy) \Leftrightarrow$
$\forall x(Fx \rightarrow \forall y((Dy \wedge Oxy) \rightarrow Bxy))$

(4) Every farmer who owns a donkey beats it.
$\forall x((Fx \wedge \exists y(Dy \wedge Oxy)) \rightarrow Bxy) \Leftrightarrow$
$\forall x(Fx \rightarrow \forall y((Dy \wedge Oxy) \rightarrow Bxy))$

These sentences are assigned their so called 'strong' readings. By a compositional interpretation procedure, they are assigned the interpretation that every farmer beats every donkey he owns.

A particular property of *DPL* is that it doesn't have what one might call an 'update' semantics. Interpretation is defined as a function over information states which *modifies* information, but which doesn't properly *update* it. As has already been indicated, with the interpretation of an existentially quantified formula $\exists x\phi$ a state s is turned into state $s[x]$ which doesn't contain any information about the value of x, no matter what information s contains about the value of x. Put bluntly, an existential quantifier involves a 'downdate' of information. With an eye on the interpretation of natural language anaphoric relationships this is quite peculiar indeed. It means that if we interpret a potential antecedent term under its association with a variable x, then, by the same token, we make a preceding potential antecedent which happened to be associated with x unavailable for further anaphoric coreference. The remarkable thing about such a dumping of information about antecedent terms is that it is triggered by a mere indexing procedure, that, one might say, it is fully arbitrary.

With her file change semantics (*FCS*, Heim 1982; Heim 1983), which is also concerned with information about the values of variables, Heim employs an essentially richer notion of information than Groenendijk and Stokhof do. Heim presents a modern version of the so-called familiarity theory of definiteness of Christopherson and Hawkins, and gives a more or less formal elaboration of the idea that indefinite noun phrases set up some kind of 'discourse referents' which pronouns can refer back to. This idea, the origins of which can be traced back to the seminal Karttunen 1968a; Karttunen 1968b, is fleshed out by explicitly encoding information about the domains of variables about the values of which information states contain information. These domains are used to formulate felicity conditions on discourse. In Heim's system, it is judged 'infelicitous' or 'inappropriate' to associate an indefinite noun phrase with a variable which already is in the domain of the current information state, and, likewise, to associate a definite noun phrase (an anaphoric pronoun, for instance) with one which is not. Thus, Heim's felicity conditions demand a proper use of variables, and exclude unwanted forms of information downdate. (In fact, Heim's notion of felicity is the *FCS* equivalent of the *DRT* requirement that natural language noun phrases should be associated with 'new' variables.)

In this paper I want to show the advantages of incorporating Heim's felicity conditions in a *DPL*-style framework. Not only will this enable us to exclude violations of ordinary variable conventions which appear to have the unwanted side effects mentioned above, but also to gain better insight in the logic of updating information about the values of variables. As will be seen below, by observing Heim's felicity conditions, we delineate a fragment of our (predicate logical) language with pleasant (dynamic) semantic properties. So, in the remainder of this paper I develop and study a dynamic semantics for the language of predicate logic, *EDPL*, which differs from *DPL* in the following two respects. First, it is formulated using information states which contain information about the domains of variables under discussion. The notion of information involved is formulated, in particular, in terms of sets of assignments of values to *subsets* of the set of variables, that is, in terms of sets of *partial* variable assignments. Second, the interpretation function itself is used to ensure felicity. Reference to the value of a variable not yet introduced, and the introduction of a variable which has already been used, are both excluded on pain of undefinedness.

As will be clear, the resulting system *EDPL* can very well be conceived of as a version of *FCS*. There are some differences, although these mainly have to do with matters of presentation. Firstly, although I agree with Heim 1990, pp. 137–8 that *FCS* shows that a level of representation is not indispensable, things have been confused since Heim's file change semantics is defined in terms of the changes of the (semantic) satisfaction sets of *files* which, themselves, are presented in nothing other than a representational manner. (Besides, it also appears that the term 'file change semantics' itself must be taken to be more than merely a metaphor. Heim's criterion for the truth of an utterance with respect to a false file can only be sensibly understood with reference to (changes of) files, not (of) satisfaction sets of files, cf., Heim 1982, pp. 337–41).) Secondly, my focus of interest is different from Heim's. Heim's interest appears to be in a semantics which gets the linguistic facts 'right', and which fleshes out the (un-)familiarity theory of (in-)definiteness. For my part, I am interested in the implications which the adoption of Heim's felicity conditions has for the underlying information structures and for the logic of their updates. Thirdly, *EDPL*'s semantics is attached to the formal language of predicate logic. This, I think, provides for a better starting point for the study of the logic of updating information about the values of variables. Lastly, indefinite noun phrases are conceived of as existential quantifiers here, not as free variables. From a theoretical point of view the main insight gained here may be that as long as Heim's felicity conditions are observed, then there is no principled difference between an approach which treats indefinites as free variables, like Heim does, and one which treats them as existential quantifiers, like Groenendijk and Stokhof do.

2.2 Heimian information states

In this section I present and study the Heimian notion of information about the values of variables, which will be used, in section 3, in the definition of a dynamic semantics for a language of predicate logic. Following Groenendijk and Stokhof, I will define the interpretation of the language of predicate logic as a function on a domain of information states, which contain information about the values of variables. However, like Heim, I will use sets of *partial* variable assignments to encode information about the values of variables. For any subset of variables X, an information state about the values of X is a set of assignments of objects to the variables in X. An information state in *EDPL* thus determines the two aspects of information addressed above. In the first place, such a state s determines a domain of variables the values of which are at issue, viz., the joint domain of the assignments in s. In the second place, the state s contains information about the values of variables in that domain.

If D is the domain of individuals and $X \subseteq V$ a subset of the set of variables, then S^X, the set of information states about the values of X, and S, the set of all information states, are defined as follows:

Definition 1 (Information states)

$S^X = \mathcal{P}(D^X)$
$S = \bigcup_{X \subseteq V} S^X$

An information state about the values of a set of variables X is a set of assignments of individuals to the variables in X. Given such a state $s \in S^X$, I will refer to X as the domain of s, indicated by $D(s)$. Information states contain information about the values of the variables in their domain by restricting their valuations. An information state s is a set of valuations of the variables in the domain of s which are conceived possible in s, and, hence, excludes all other valuations to these variables. So, if x and y are in the domain of s, then s contains the information that the value of x is a man iff no i in s maps x on an individual that is not a man, and s contains the information that the value of x sees the value of y, iff for all i in s, $i(x)$ sees $i(y)$.

With respect to a fixed domain of variables, we can define the notions of a minimal and a maximal information state. For $X \subseteq V$, the minimal information state about the values of X is the set of all valuations of X, D^X, hereafter referred to as $\top^X$. A minimal information state has no information about the values of the variables in its domain, all valuations of the variables are (still) considered possible. A maximal information state about the values of X is $\{i\}$ for any $i \in D^X$. A maximal state has total information about the values of the variables in its domain: only one valuation of them is conceived possible. Furthermore, for any domain X, the absurd information state is $\emptyset$, referred to as $\perp^X$. An absurd information state excludes all assignments to the variables in its domain.

We can distinguish two borderline cases when we consider the set of information states about the values of a set X of variables. The set of variables X can be empty, and the domain of individuals D can be empty:

Observation 1

if $X = \emptyset$, then $S^X = \{\{\emptyset\}, \emptyset\}$
if $D = \emptyset$, then $D(s) = \emptyset$ or $s = \perp^{D(s)}$

There are only two states of information about the values of nil variables, the set containing the empty assignment, and the empty set. (In fact, this is the domain of truth values on its set-theoretical definition.) Interestingly, with respect to the empty domain the minimal information state and the maximal information state coincide. This reflects the fact that one can have no substantial information about the values of nil variables. Notice that, on the other extreme, the set S^V of information states about the values of all variables equals the set of *(D)PL* information states. So we see that, as a matter of fact, our Heimian information space encompasses the states of propositional logic and those of (dynamic) predicate logic.

Although this may go against ordinary usage, we can also consider states of information about the values of variables relative to an empty domain of individuals. Interestingly, if $D = \emptyset$, then we still have a true and a false information state about the values of nil variables (i.e., also in that case $S^\emptyset$ is $\{\{\emptyset\}, \emptyset\}$), but about the values of non-empty sets of variables we only have absurd information states. If $D = \emptyset$ and $\emptyset \neq X$, then $D^X = \emptyset$ and, consequently, $S^X = \mathcal{P}(D^X) = \{\emptyset\}$. Again, this is as it should be. One can only have absurd information about the possible values of a variable, if a variable cannot have a value. Even the seemingly tautologous information that the value of a variable is identical to itself is false then, simply because if there is no individual then there is no individual which is identical to itself. As we will see at the end of section 3.3, the possibility of an empty domain of individuals is not a mere curiosity which renders all existentially quantified statements false. In the meantime, however, I will assume non-empty domains of individuals, as is customary.

Since our information states model two aspects of information about variables, two basic kinds of update of information can be distinguished. Update of information consists of either getting more information about the values of variables, by the elimination of partial variable assignments, or of extending the domain of partial variable assignments (or, of course, of a mixture of both). Before we give the *EDPL* definition of information update, it is useful to have at our disposal some notions relating assignments and states with different domains. With $D(i)$ I will indicate the domain of i ($D(i) = X$ iff $i \in D^X$) and I will say that j is an extension of i, $i \leq j$, iff $\forall x \in D(i)$: $x \in D(j)$ and $i(x) = j(x)$ (set-theoretically speaking, $i \leq j$ iff $i \subseteq j$). Using this notion of assignment extension, we can 'generalize' the notion of set membership in two directions. An assignment can have an extension in a state with a larger domain and a restriction in a state with a smaller domain:

Notation Convention 1 (Extension and restriction)

i has an extension in t, $i \lessdot t$, iff $D(i) \subseteq D(t)$ and $\exists j \in t$: $i \leq j$

i has a restriction in s, $i \gtrdot s$, iff $D(i) \supseteq D(s)$ and $\exists j \in s$: $i \geq j$

Assignment i has an extension in t if i is a restriction of some element of t. In this case, I will also say that i survives in t. It has a restriction in s if it is an extension of some element of s. Clearly, if $D(s)$ equals $D(i)$, then $i \gtrdot s$ iff $i \in s$ iff $i \lessdot s$.

The two notions of restriction and extension give rise to two generalizations of the subset relation, the substate relation and the update relation:

Definition 2 (Substate and update)

s is a substate of t, $s \sqsubseteq t$, iff $D(s) \subseteq D(t)$ and $\forall i \in s$: $i \lessdot t$

t is an update of s, $s \geq t$, iff $D(s) \subseteq D(t)$ and $\forall i \in t$: $i \gtrdot s$

The substate relation plays a part, mainly, in the definition of (*DRT*-like) notions of truth (or, rather, support) and entailment. If state s is a substate of state t, then t does not reject valuations of the variables in $D(s)$ which are conceived possible in s. So, if s is a substate of t, then t contains no more information than s has about the variables in the domain of s. State t may, however, have variables in its domain which s is silent about. The update relation holds between two states s and t iff t contains more information than s. In an update t of s only valuations of the variables in $D(s)$ are considered possible which are possible in s. Hence, t contains at least the information that s contains about the variables in $D(s)$. Moreover, t may have information about variables which s is silent about. So, the definition of $\geq$ precisely captures the notion of update informally described above. Notice that it is not the case that $s \sqsubseteq t$ iff $t \geq s$. However, we do find that $s \sqsubseteq t$ and $t \geq s$ iff $s \subseteq t$.

Both the update and the substate relation partially order the domain of information states:

Observation 2 $\langle S, \sqsubseteq \rangle$ and $\langle S, \geq \rangle$ are partial orders (i.e., $\sqsubseteq$ and $\geq$ are reflexive, transitive, and antisymmetric)

Of course, it might have been expected that a state is an update of itself (reflexivity), that an update of an update is an update (transitivity) and that if a state s is an update of an update of itself, then the update is identical to s itself (antisymmetry). The fact that the substate relation also induces a partial order may come as a surprise, for our entailment relation is stated in terms of the substate relation and dynamic entailment relations are *not* in general reflexive and transitive (cf., section 3.2).

The orders of S induced by $\sqsubseteq$ and $\geq$ have unique weakest and strongest elements. The weakest information state in $\sqsubseteq$ is $\top^V$, the state of no information about all variables, and the strongest state in $\sqsubseteq$ is $\bot^\emptyset$, the absurd state of information about no variables. The strongest state in $\geq$ is $\bot^V$, the absurd state of information about all variables, and the weakest information state in $\geq$ is $\top^\emptyset$, the state of no information about no variables. These observations can be pictured as follows (the ordering relations $\sqsubseteq$ and $\geq$ should be read as if rotated 90^o anticlockwise):

$$\begin{array}{cccc} & \top^V & \qquad & \bot^V \\ \sqsubseteq & \vdots & \qquad \geq & \vdots \\ & \bot^\emptyset & \qquad & \top^\emptyset \end{array}$$

Notice that $\sqsubseteq$ ($\geq$) ranges from the bottom (top) of propositional logic to the top (bottom) of predicate logic.

For the update structure $\langle S, \geq \rangle$ we can define a join and a meet, which we will call the 'product' and the 'common ground' of two information states, respectively:

Definition 3 (State product and common ground)

$s \wedge t = \{i \in D^{D(s) \cup D(t)} \mid i \gtrdot s \text{ and } i \gtrdot t\}$
$s \vee t = \{i \in D^{D(s) \cap D(t)} \mid i \lessdot s \text{ or } i \lessdot t\}$

The product of two states s and t contains the combined information contained in s and t. It contains information about the values of all the variables s or t contains information about, and it excludes valuations of these variables which are excluded by s or t. The common ground of s and t excludes possible values of variables excluded by both s and t. It is relatively easily seen that, for every two states s and t, the product $s \wedge t$ is the unique weakest state t' such that $s \geq t'$ and $t \geq t'$ and the ground $s \vee t$ is the unique strongest state s' such that $s' \geq s$ and $s' \geq t$. This implies that:

Observation 3 $\langle S, \geq \rangle$ is a lattice

2.3 Discourse referents

Before we turn to the semantics of *EDPL*, it may be clarifying at this point to join a discussion leveled by Landman against Heim's 'reconstruction' of Karttunen's notion of a discourse referent (cf., Karttunen 1976, Heim 1982, p. 249ff., Landman 1986, p. 110ff.). Heim starts of from Karttunen's stipulation that "(...) the appearance of an indefinite noun phrase establishes a 'discourse referent' just in case it justifies the occurrence of a co-referential pronoun or a definite noun phrase later in the text" (Karttunen 1976, p. 366), and she proposes to "(...) identify discourse referents simply with numbers: to establish a discourse referent means to carry a referential index in logical form" (p. 250). After having discussed anaphora and the lifespan of discourse referents, Heim draws the following conclusion (p. 262/3):

> (...) insight is gained by reducing the concept of a discourse referent and its lifespan to certain formal properties of logical form, rather than using it as an unanalyzed concept as Karttunen did. (...) Aside from explaining some facts that Karttunen could only describe, I hope to have dispelled some of the mystery that the notion of a discourse referent is likely to have presented, especially to philosophically-minded semanticists: If they are not referents, what are they? What is their relation to referents where there are any? Such questions are difficult if it is presumed that discourse referents, although not real referents, are individuals of some sort. But under the current reconstruction—or elimination, if you like—of the concept, they are no more puzzling than the question: What is the meaning of a variable?

Landman is partially satisfied with this conclusion (p. 110/111):

> I whole-heartedly agree with this last sentence, and I think it helps very much to sharpen the situation. What we *disagree* about, however, is the answer to the last question. For Heim the notion of variable meaning is unproblematic; being a philosophically-minded semanticist, for me the notion of a variable meaning is a mystery, and the question: What is the meaning of a variable? one of the main problems in semantics.

Landman is concerned with epistemically contingent identity conditions on discourse referents, and—like Heim, by the way—he is in search of objects in disguise: "Objects in disguise are typically *partial objects*; with respect to their identity conditions they are typically objects that can *grow into the same object*" (p. 112). Landman observes that variables or variable meanings are unsuited for modeling such objects in guises in a classical logic. The meaning of a variable x, Landman argues, must be taken to be a function from variable assignments to variable extensions, viz., the function that assigns to any assignment g the value $g(x)$ that g assigns to x. This means that "(...) for every variable it is immediately clear, which of the elements in the set of variable meanings is *its* meaning; given a variable meaning, it is immediately clear, of what variable it is the meaning. (...) in fact, *there is no reason not to equate the variable meaning with the variable itself* (or with its index)" (p. 111). But then (p. 112):

> If the variable meanings are just the variable themselves, then the basic entities at the level of Heim's discourse semantics are no less symbols than are discourse markers, and their identity conditions are equally trivial: either they are not the same file card, and then they are distinct, or *it is* the same file card. (...) If discourse representation theory wants to take these objects seriously, it cannot leave the analysis of identity statements to the classical logic of the metalanguage, but it has to give a non-classical semantics for them itself, a semantics which provides it with a domain of partial objects.

Landman concludes that the objects we talk about have to be modeled using a primitive domain of partial objects, 'pegs' as he calls them.

However, there is no real need to resort to such primitive pegs. The 'partial objects' Landman argues for are already implicitly present in our Heimian information states. These partial objects actually are the partial correlates of Landman's variable meanings, viz., partial functions from sets of variable assignments to variable extensions. In order to substantiate this claim, I will end this section with a formal specification of these partial objects, 'discourse referents', and show that they have the properties Landman argues partial objects have.

Heim and Landman raise the question what the meaning of a variable is. In the system of interpretation presented in the next section, the interpretation function is conceived of as a function updating information about the values of variables, and the interpretation of variables is, of course, taken to be information state dependent. So, if the standard, total, meaning of a variable can be identified with a function from (the set of all total) variable assignments to variable values, our partial counterpart of it can be identified, relative to an information state (a set of partial variable assignments), with a function from this subset of partial variable assignments to variable values. What I will call the 'denotation' $[x]_s$ of a variable x in a state s, then, is the restriction of the Landman-meaning of x to the assignments considered possible in s. These variable denotations can be conceived of as the discourse referents of a state s:

Definition 4 (Discourse referents)

$[x]_s$, the denotation of $x \in D(s)$ in s, is the function $\{\langle i, i(x)\rangle \mid i \in s\}$
$\mathcal{D}_s$, the set of discourse referents of s, is $\{[x]_s \mid x \in D(s)\}$
$\mathcal{D}$, the set of discourse referents, is $\bigcup_{s \in S} \mathcal{D}_s$

An object $\delta \in \mathcal{D}_s$ can be called a total object iff it is a constant function, so, if $[x]_s$ is a total object, then according to s the identity of the value of x is fixed. An object $\delta \in \mathcal{D}_s$ can be called impossible iff s is absurd, and, consequently, δ is the empty function.

The notion of a discourse referent presently developed out of our notion of information about the values of variables appears to be the type of thing Landman is looking for. In Landman's opinion, partial objects typically are things we assume in conversation, which we ascribe properties to and which can be followed through information growth. The most characteristic feature of these objects is that they need not, but can (turn out to) be identical to other partial objects, and that they can grow into less partial objects. The same can be said of our discourse referents. Discourse referents have precisely the required kind of partial identity conditions (running ahead of the definition presented in section 3.2, $s \models \phi$ is used to say that s supports ϕ):

Observation 4 $[x]_s = [y]_s \Leftrightarrow s \models x = y$, but $[x]_s \neq [y]_s \not\Rightarrow s \models x \neq y$

An identity statement $x = y$ is supported by a state s (or 'is true in s') *iff* the denoted discourse referents are identical. Clearly, also in case partial objects are at issue identity statements must have classical truth conditions. However, the fact that two discourse referents are not identical should not, and does not, imply that they *must* be different. This at most implies that they are not known to be identical, i.e., that they *might* be different. So, we see that, in fact, our discourse referents can be truly conceived of as *partial* objects.

Furthermore, although a particular discourse referent cannot be in the domain of two different information states, it can be followed under information growth, via the following 'approximation' relation:

Definition 5 (Approximation) If $\delta \in \mathcal{D}_s$, $\delta' \in \mathcal{D}_{s'}$, then δ approximates δ', $\delta \geq \delta'$, iff $D(s) \subseteq D(s')$ & $\forall j \in s'\ \exists i \in s : i \leq j$ & $\delta(i) = \delta'(j)$

If s' is an update of a state s, then for any discourse referent δ in $\mathcal{D}_s$ there is one, less or equally partial, object δ' in $\mathcal{D}_{s'}$ such that δ approximates δ'. Notice how intimately related approximation is with update and identity. The denotation of a variable x in s approximates the denotation of a variable y in s' iff s' is an update of s, and x and y denote the same discourse referent in s'. As a matter of fact, the universe of discourse referents is partially ordered by the approximation relation, which is as it should be. Of course, a discourse referent must be taken to approximate itself (reflexivity), and the growth of discourse referents into less partial objects must be taken to be a transitive relation. And what else could one expect if two discourse referents approximate each other, then that they are identical? In fact:

Observation 5 $\langle \mathcal{D}, \geq \rangle$ is a meet semilattice

One can always construct a meet of two partial objects. If they are partial objects which are known to be distinct, then their meet simply is the impossible object. There is not in general a join of two partial objects. However, if it is the case that two discourse referents, from two states s and t, are grown out of one partial object, then they are approximated by an object in the common ground of s and t, and then they are what one may call a 'shared' object. Thus we find, for instance, that for any $x \in D(s) \cap D(s')$, the denotation of x in s and s' is a shared object.

To wind up this discussion, let me recapitulate the results. I have given a definition of discourse referents, upon which they are shown to have properties characteristic of partial objects. Discourse referents have partial identity conditions, they can be followed under information growth (they can be dressed up with further properties, that is), and they can be shared by different information states. These discourse referents are constructed from the information states in which they live. So, we see that, with the Heimian notion of partial information about the values of variables at our disposal, we have no no need to assume, as Landman would have it, an additional primitive domain of partial objects. Irony has it that Landman criticizes Heim for her rejection of discourse referents, which actually are grown in her own information space.

3 EDPL

We are now ready to turn to the system of *EDPL*, a system in which the formulas of predicate logic are interpreted as functions updating Heimian information states, and which are enforced to obey Heim's felicity conditions. In the first subsection the *EDPL* interpretation function is presented, in the second the *EDPL* notions of support and entailment are discussed, and in the third subsection some of its logical properties are investigated.

3.1 Semantics of EDPL

The language of *EDPL* is that of predicate logic with identity and without individual constants and function symbols. A model for *EDPL* is an ordinary first order model $M = \langle D, F \rangle$ which consists of a set of individuals D and an interpretation function F which assigns sets of n-tuples of objects to n-ary relation expressions. (I will omit reference to M whenever this does not lead to confusion.) Interpretation is defined as a *partial* update function on the domain of information states. (Un-)definedness of interpretation is the *EDPL* equivalent of Heim's (in-)felicity (and of *DRT*'s fresh variable convention, for that matter). If a formula contains a free variable for which a state is undefined, then the interpretation of the formula is undefined for that state. Conversely, if a state already has a variable x in its domain, then the interpretation of a formula is undefined for it if it contains a quantifier binding x. Undefinedness persists in the following way. If ϕ is undefined for s, then so are $\neg\phi$ and $\phi \wedge \psi$. Furthermore, if ϕ is undefined for $s[x]$, then $\exists x\phi$ is undefined for s, and if ψ is undefined for $s[\![\phi]\!]$, then $\phi \wedge \psi$ is undefined for s.

Before we turn to the definition *EDPL* interpretation, it is convenient to introduce three more auxiliary notions: those of (state) restriction, subtraction and extension (here $i \leq_X j$ iff $D(j) = D(i) \cup X$ and $i \leq j$):

Notation Convention 2 (Restriction, subtraction and extension)
If $s \in S^X$, $t \in S^Y$ and $X \subseteq Y$, then

$$\begin{array}{lll} t^X & = \{i \in D^X \mid i < t\} & (\in S^X) \\ s - t & = s \setminus t^{D(s)} & (\in S^X) \\ s[x] & = \{j \mid \exists i \in s\colon i \leq_{\{x\}} j\} & (\in S^{X \cup \{x\}}) \end{array}$$

The restriction t^X of t to X consists of all those assignments to X which have an extension in t. The subtraction of state t from state s gives a state which contains all those assignments in s which have no extension in t. So, $s - t$ contains the information that s has about the variables in $D(s)$ supplemented with the information excluded by t. The notion of extension will replace the *DPL* notion of reinstantiation. The state $s[x]$ at most differs from s in that it is defined for x. About the values of variables in the domain of s the new state $s[x]$ contains precisely the same information as s, and, if x is not in the domain of s, $s[x]$ is completely impartial about the value of x. What is added, one might say, is the information that x has a value. (Clearly, it *is* possible that x is already in the domain of s (and then $s[x] = s$), but in the semantics of *EDPL* this possibility is ruled out on pain of undefinedness).

Let us now turn to the interpretation function itself. In the following definition I indicate the possibility of undefinedness in the source cases only:

Definition 6 (Semantics of EDPL)

$$\begin{array}{lll} s[\![Rx_1 \ldots x_n]\!] & \{i \in s \mid \langle i(x_1), \ldots, i(x_n)\rangle \in F(R)\} & \text{if } x_1, \ldots, x_n \in D(s) \\ s[\![x = y]\!] & \{i \in s \mid i(x) = i(y)\} & \text{if } x, y \in D(s) \\ s[\![\neg\phi]\!] & s - s[\![\phi]\!] & \\ s[\![\exists x\phi]\!] & s[x][\![\phi]\!] & \text{if } x \notin D(s) \\ s[\![\phi \wedge \psi]\!] & s[\![\phi]\!][\![\psi]\!] & \end{array}$$

where $s[\![Rx_1 \ldots x_n]\!], s[\![x = y]\!] \in S^{D(s)}$

The interpretation of an atomic formula $Rx_1 \ldots x_n$ in s preserves the assignments in s that map the variables $x_1, \ldots, x_n$ onto individuals $z_1, \ldots, z_n$ that stand in the relation R, in that order. The interpretation of $\neg\phi$ in s yields the state which consists of those assignments in s which do not survive the update with ϕ. So, it excludes from s those valuations of the free variables of ϕ under which ϕ would be verified. As in *DPL*, conjunction amounts to functional composition. New is the clause dealing with existential quantification. An existential quantifier $\exists x$ introduces arbitrary valuations of x. By the side condition, the variable x is required to be 'new', and the input information state is updated with the information that x has a value.

I now want to illustrate the *EDPL* interpretation function by means of some examples. Throughout, definedness is assumed. Since existential quantification and conjunction involve the composition of two operations on states, *DPL*'s characteristic donkey equivalences are retained:

Observation 6

$$(\exists x\phi \wedge \psi) \Leftrightarrow (\exists x(\phi \wedge \psi))$$
$$((\phi \wedge \psi) \wedge \chi) \Leftrightarrow (\phi \wedge (\psi \wedge \chi))$$

Like *DPL*, *EDPL* accounts for the fact that indefinite noun phrases (existential quantifiers) in one sentence may bind pronouns (free variables) in successive sentences. Again, this is achieved in a compositional way.

It may be elucidating to inspect an example in a little detail. Consider the interpretation of *A cowgirl meets a boy*, $\exists x(Cx \wedge \exists y(By \wedge Mxy))$, in a (non-absurd) state s. This gives us $s[\![\exists x(Cx \wedge \exists y(By \wedge Mxy))]\!] = s[x][\![Cx]\!][y][\![By]\!][\![Mxy]\!]$, which is the following set of assignments (where $X = D(s) \cup \{x, y\}$):

(5) $\{j \in D^X \mid \exists i \in s\colon i \leq_{\{x,y\}} j \ \&\ j(x) \in F(C) \ \&\ j(y) \in F(B) \ \&\ \langle j(x), j(y)\rangle \in F(M)\}$

This state, which will be referred to as state t for the time being, contains the information that the value of x is a cowgirl and the value of y is a boy and that the value of x meets the value of y. For any $i \in s$ and for any cowgirl z and boy z' met by z, there is an extension j of i in t such that $j(x) = z$ and $j(y) = z'$.

Subsequent update of t with the sentence *She impresses him*, Ixy, gives us $t[\![Ixy]\!]$, which is the following set of assignments:

(6) $\{j \in D^X \mid \exists i \in s\colon i \leq_{\{x,y\}} j \ \&\ j(x) \in F(C) \ \&\ j(y) \in F(B) \ \&\ \langle j(x), j(y)\rangle \in F(M) \ \&\ \langle j(x), j(y)\rangle \in F(I)\}$

The state $t[\![Ixy]\!]$ contains the information that the value of x is a cowgirl, that the value of y is a boy and that the value of x meets and impresses the value of y.

The interpretation of $\neg\phi$ in a state s involves the subtraction of $s[\![\phi]\!]$ from s. This subtraction eliminates assignments in s that verify ϕ. Consider the following simple example:

(7) No man sees her.
$\neg\exists y(My \wedge Syx)$

The interpretation of this sentence in a state s yields the state $s[\![\neg\exists y(My \wedge Syx)]\!] = s - s[y][\![My]\!][\![Syx]\!]$. The state $s[y][\![My]\!][\![Syx]\!]$, which is subtracted from the original state s, is the following set of assignments (here $Y = D(s) \cup \{y\}$):

(8) $\{j \in D^Y \mid \exists i \in s\colon i \leq_{\{y\}} j \ \&\ j(y) \in F(M) \ \&\ \langle j(y), j(x)\rangle \in F(S)\}$

The subtraction of this state from s yields a state that consists of the assignments in s that do not survive in $s[y][\![My]\!][\![Syx]\!]$:

(9) $\{i \in s \mid \neg\exists j\colon i \leq_{\{y\}} j \ \&\ j(y) \in F(M) \ \&\ \langle j(y), j(x)\rangle \in F(S)\}$

In other words, the result of interpreting $\neg\exists y(My \wedge Syx)$ in s is a state that contains the information that the value of x is an individual such that no other individual can be found that is a man and sees her.

It may be noticed that the interpretation of a formula $\neg\phi$ is purely eliminative, which is to say that it does not yield extensions of the domain of variables under discussion. This corresponds to the fact that, usually, if a noun phrase stands in the scope of a negation it cannot serve as an antecedent for pronouns that occur outside the negation's scope. As a consequence, the law of double negation doesn't hold in *EDPL*. The interpretation of the double negation of ϕ in s preserves the assignments in s which verify ϕ, i.e., the ones that have an extension in the update of s with ϕ, but not the extensions themselves:

Observation 7 $s[\![\neg\neg\phi]\!] = s - (s - s[\![\phi]\!]) = \{i \in s \mid i < s[\![\phi]\!]\} = s[\![\phi]\!]^{D(s)}$

In other words, the double negation of ϕ conveys the same information as ϕ about the variables in the domain of an information state, but it cancels possible extensions of its domain. Adopting the terminology of Groenendijk and Stokhof, we may call $\neg\neg\phi$ the static closure of ϕ and write it as $\downarrow\phi$.

In *EDPL* universal quantification, and implication are defined, as usual, in terms of existential quantification, negation and conjunction, so, $\phi \rightarrow \psi = \neg(\phi \wedge \neg\psi)$ and $\forall x\phi = \neg\exists x\neg\phi$. Upon this definition, the second donkey equivalences are also valid in *EDPL*:

Observation 8

$$(\exists x\phi \rightarrow \psi) \Leftrightarrow (\forall x(\phi \rightarrow \psi))$$
$$((\phi \wedge \psi) \rightarrow \chi) \Leftrightarrow (\phi \rightarrow (\psi \rightarrow \chi))$$

It turns out that universally quantified formulas and implications have the following semantic effects:

Observation 9

$$\begin{array}{ll} s[\![\forall x\phi]\!] & = \{i \in s \mid \forall j \in s[x]\text{: if } i \leq j \text{ then } j < s[x][\![\phi]\!]\} \\ s[\![\phi \rightarrow \psi]\!] & = \{i \in s \mid \forall j \in s[\![\phi]\!]\text{: if } i \leq j \text{ then } j < s[\![\phi]\!][\![\psi]\!]\} \\ s[\![\forall x(\phi \rightarrow \psi)]\!] & = \{i \in s \mid \forall j \in s[x][\![\phi]\!]\text{: if } i \leq j \text{ then } j < s[x][\![\phi]\!][\![\psi]\!]\} \end{array}$$

So, for instance, the interpretation of $\forall x(\phi \rightarrow \psi)$ in a state s preserves the assignments i in s every extension of which in $s[x][\![\phi]\!]$ has an extension in $s[x][\![\phi]\!][\![\psi]\!]$. Let us briefly consider an example in which a free variable occurs:

(10) Every man sees her.
$\forall y(My \rightarrow Syx)$

The interpretation of this sentence in a state s yields the state $s[\![\forall y(My \rightarrow Syx)]\!]$, which is $s - (s[y][\![My]\!] - s[y][\![My]\!][\![Syx]\!])$. This state consists of the assignments i in s such that on every extension of i to y such that the value of y is a man, the value of y sees the value of x. The resulting state consists of those assignments i in s which assign an individual to x which every man sees.

The last example of this section is the universally quantified donkey sentence:

(11) Every man who owns a donkey beats it.
$\forall x((Mx \wedge \exists y(Dy \wedge Oxy)) \rightarrow Bxy)$

The interpretation of this sentence in a state s preserves the assignments in s every extension of which in $s[x][\![Fx]\!][y][\![Dy]\!][\![Oxy]\!]$ verifies Bxy. The resulting state contains an assignment i in s iff on every extension j of i to x and y, the value of x

beats the value of y if the value of x is a man and the value of y is a donkey which the value of x owns. Clearly, this is tantamount to the requirement that every man beats every donkey he owns.

To conclude this section, it may be worthwhile to point out that quantifiers in *EDPL* can as well be conceived of as atoms. (Cf., also, Vermeulen 1993 who makes a similar point.) That is to say, we might as well define $s[\![\exists x\phi]\!]$ as $s[\![\exists x \wedge \phi]\!]$, with the stipulation that $s[\![\exists x]\!] = s[x]$ if $x \notin D(s)$ (undefined otherwise). The results would be entirely the same. Interestingly, we would find by definition that $s[\![\forall x\phi]\!] = s[\![\neg(\exists x \wedge \neg\phi)]\!] = s[\![\exists x \rightarrow \phi]\!]$, that is, that $\forall x\phi$ says what ϕ says on the assumption that x has an arbitrary value.

Although conceiving of quantifiers as atoms does not in any way affect the semantic properties of the system, it does have some enlightening consequences. First, if quantifiers are atoms, then undefinedness originates from atoms only (atomic formulas and quantifiers). Second, if existentially quantified formulas are conjunctions which have a quantifying atom as their first conjunct, then the (special) donkey equivalences are instances of the (ordinary) associativity laws governing conjunction. Third, for the same reason, inference rules dealing with $\exists$ elimination and introduction would be subsumed under the rules dealing with $\wedge$ elimination and introduction. Interestingly, the quantifier needn't be dealt with in a deduction system for *EDPL*. Of course, this is not to say that an existential quantifier doesn't mean anything. By the declaration by means of the atom $\exists x$ that x has a value, it is asserted that something exists.

Although I do think these are good reasons for 'deconstructing' quantified formulas, I nevertheless refrain from doing so, and stick to ordinary usage in (most of) this paper. Only in section 3.3, where a sketch of *EDPL* deduction is given, I will adopt the alternative view.

3.2 Update, support and entailment

Before we turn to the definition of support and entailment in *EDPL*, it is useful to state some general observations about its interpretation function. In the first place, it is distributive:

Observation 10 $s[\![\phi]\!] = \bigcup_{i \in s} \{i\}[\![\phi]\!]$, if defined

This observation tells us that *EDPL* formulas do not express global properties of information states, but that they are only sensitive to the individual possibilities in an information state. In the second place, interpretation in *EDPL* always produces information update:

Observation 11 $s \geq s[\![\phi]\!]$, if defined

Although the assignments in the state that results from updating a state s with ϕ are not themselves elements of the original state s, as is the case in a real eliminative semantics, all these assignments are at least extensions of assignments in s. In other words, the information contained in an initial state of interpretation is properly preserved in the process of interpretation, and this fact positively distinguishes *EDPL* from *DPL*.formation about are sentenced undefined.

Technically speaking it is convenient that *EDPL* is an update semantics. Since interpretation in *EDPL* at most yields addition of information, the notion of the truth of a formula in an information state (or, rather: that of a formula being supported by an information state), can simply be defined in terms of the absence of proper update. So, more specifically, a formula can be said to be supported if it does not convey additional information about discourse referents introduced before:

Definition 7 (Support) s supports ϕ wrt. M, $s \models_M \phi$, iff $s \sqsubseteq s[\![\phi]\!]_M$

A formula ϕ is supported by a state s iff s is a substate of the update of s with ϕ, i.e., iff all assignments in s survive the update with ϕ. This is to say that s supports ϕ if s already contains the information conveyed by ϕ about the values of the variables in the domain of s. Obviously, we want to allow ϕ to be supported by a state s, also if in the update of s with ϕ new variables are introduced. If s contains the information that the value of x sees a donkey, then also a sentence *She sees an animal*, $\exists y(Ay \wedge Sxy)$, must be supported by s.

In order to illustrate the above definition, let us return to our earlier example *A cowgirl meets a boy*, $\exists x(Cx \wedge \exists y(By \wedge Mxy))$. This formula is supported by s iff $s \sqsubseteq t$, where t is the result of interpreting the formula in s. This turns out to be the case iff, in the model, there is a cowgirl who meets a boy. It may be noticed that a state of information either supports this formula or its negation. The formula contains no free variables, and, hence, it conveys no information about the values of variables in the domain of s. Things are different for the continuation *She impresses him*, Ixy. The mediating state t supports Ixy iff $t \sqsubseteq t[\![Ixy]\!]$, which would be the case if, as a matter of fact, all cowgirls impress every boy they meet. Apparently, this need not be the case, without such prohibiting a successful (non-absurd) update of t with Ixy. Without further assumptions, an interpreter may be inclined to accept both *She impresses him* and *She doesn't impress him* as a possible continuation of *A cowgirl meets a boy*.

As we can see here, if a formula ϕ is defined for a state s, it may be the case that s does not support either ϕ or its negation $\neg\phi$. The information that a state s has about certain variables $x, \ldots$ may be partial, and, thus, it may happen that s supports neither a formula $\phi(x)$ nor its negation $\neg\phi(x)$. On the other hand, if a formula is defined for it, then it is either true or false in any maximal state of information $\{i\}$. If we have that maximal state $\{i\} \models \phi$, then I will also say that i itself supports ϕ.

Using the observations about update and distributivity it can be shown that cases that support a formula are registered (i.e., extended) in the interpretation of the formula:

Observation 12 For all s, $i \in D^{D(s)}$: $i \in s$ & $\{i\} \models \phi$ iff $i < s[\![\phi]\!]$

An assignment i in s supports ϕ iff i survives in the update of s with ϕ. So, the update of a state s with ϕ contains (only) assignments that register, i.e., extend, the assignments in s that support ϕ. It also follows from observation 12 that a state supports a formula iff all of its elements do. This means that our global definition of support in terms of a state s and its update $s[\![\phi]\!]$ comes down to the same thing as *DPL*'s point-wise definition in terms of the 'updates' of singular assignments. The fact that a proper notion of support can be defined in this global manner is not just a matter of terminological or conceptual simplicity, for it appears that if epistemic modalities or presuppositions are introduced, a global notion of support *has to* be used. In that case the distributive definition which *DPL* is forced can be seen to give incorrect results (cf., Groenendijk et al. 1994; Beaver 1995).

Let us now turn to the *EDPL* definition of entailment. In *EDPL* a sequence of premises is said to entail a conclusion, if update with the premises is guaranteed to yield a state which supports the conclusion:

Definition 8 (Entailment in EDPL)

$\phi_1, \ldots, \phi_n$ entail ψ, $\phi_1, \ldots, \phi_n \models \psi$ iff $\forall M, s \in S^{\emptyset}$: $s[\![\phi_1]\!]_M \ldots [\![\phi_n]\!]_M \models_M \psi$

An inference is valid iff the update with the premises of a state of information about no variables always produces a state of information about the values of variables about which subsequent update with the conclusion adds no more information. The reader may notice that for $\phi_1, \ldots, \phi_n$ to entail ψ it is not required that ψ be true in the update of any state whatsoever with the premises $\phi_1, \ldots, \phi_n$. The reason is that we wouldn't want to call an inference invalid only because its interpretation is undefined for some information states. (For this would disqualify any sequence in which a variable occurs.) Still the notion of entailment as it is defined is general enough. If $\phi_1, \ldots, \phi_n$ entail ψ, then as long as the sequence $\phi_1, \ldots, \phi_n, \psi$ is defined, update with the premises always yields a state which supports the conclusion:

Observation 13 If $\phi_1, \ldots \phi_n \models \psi$, then $\forall s \in S$: $s[\![\phi_1]\!] \ldots [\![\phi_n]\!] \models \psi$ if defined

EDPL also licenses the deduction theorem:

Observation 14 $\phi_1, \ldots, \phi_n \models \psi$ iff $\phi_1, \ldots, \phi_{n-1} \models \phi_n \rightarrow \psi$

From the deduction theorem and the donkey equivalences it can be seen to follow that *EDPL* has a dynamic notion of entailment, like *DPL*. Free variables in the conclusion of an entailment may refer back to discourse referents introduced in the premises. For instance, $\exists x Fx$ entails Fx. This fact corresponds to the following line of elementary reasoning:

(12) If a man comes from Rhodes, he likes pineapple-juice. A man I met yesterday comes from Rhodes. So, he likes pineapple-juice.
$(\exists x(Mx \wedge Rx) \rightarrow Lx), \exists x(Mx \wedge Rx) \models Lx$

Digression We have seen that our notion of entailment corresponds to the strong implication of *EDPL* (which requires that *every* assignment extension in the update with the antecedent satisfies the consequent). But it has been argued, for instance by Schubert and Pelletier 1989, that the strong readings of donkey sentences are misguided, or, at least, are not the only readings these sentences have. A convincing example is the following one, labeled the 'dime implication':

(13) If I have a dime in my pocket, I'll put it in the parking meter.

On its most natural reading the dime implication says that if I have one or more dimes in my pocket, then I will throw one in the meter. However, if we interpret the dime implication, like the donkey implication, as one of strong implication, then the sentence would imply that I throw *all* the dimes I have in my pocket in the meter. With respect to this example, the strong reading is quite odd.

It is possible to define a notion of weak implication that assigns conditional sentences the weak truth conditions that Schubert and Pelletier argue for, and that preserves the internal dynamics of the implication. This is, in fact, Chierchia 1992's definition of implication. I use $\hookrightarrow$ to indicate the weak notion of implication:

$$\phi \hookrightarrow \psi = \downarrow\phi \rightarrow (\phi \wedge \psi)$$

The interpretation of a weak implication $\phi \hookrightarrow \psi$ in a state s gives us those $i \in s$ such that if i has an extension in $s[\![\phi]\!]$, then i has an extension in $s[\![\phi]\!][\![\psi]\!]$. If we interpret the dime implication employing this weak notion of implication, then the sentence is true in a state s if I throw a dime I have in my pocket in the meter, if I have dimes in my pocket at all.

However, it would be misguided to entirely reject the strong implication in favor of the weak one. Consider the following variant of the dime implication:

(14) If I have a dime in my pocket, I won't throw it in the meter.

Interpreting this sentence as a weak implication seems to give just as odd results as the strong interpretation of the original dime sentence. On the weak reading of the present sentence, its truth conditions only require that if I have one or more dimes in my pocket, I keep at least one dime but maybe throw all other dimes I have in the parking meter. On the other hand, on the strong reading it is required that I don't throw any dime I might have in my pocket in the meter, which is appropriate.

What is interesting about the alternative dime sentence is that it completely corresponds to the original dime sentence, that it can be uttered in precisely the same contexts as the original dime sentence, but that it, nevertheless, seems to require the strong interpretation of the conditional, whereas the original dime sentence appears to require the weak interpretation. In the fourth section of this paper, I will come back to the issue of strong and weak implication, and I will show there that the weak and strong readings of conditionals in fact fit in a more general scheme of (universal) adverbial quantification. For the remainder of this section, I will continue to assume the strong interpretation of the implication sign, since that is the one which corresponds to our notion of entailment. To conclude this section, I want to point at two welcome properties of the *EDPL* entailment relation. Dynamic entailment relations like that of Groenendijk and Stokhof 1991 and Veltman 1995 fail to obey certain classical structural rules. Both *DPL* and Veltman's update semantics are not idempotent, which is to say that in both systems the entailment relation is not reflexive. (For instance, in *DPL* $Fx \wedge \exists x Gx$ is not entailed by itself.) Furthermore, in the first system this relation isn't transitive either. (Although $(\exists x \neg Px \wedge \neg\neg\exists x Px)$ entails $\exists x Px$ and the last in its turn entails Px, the first by itself does not entail Px, but $\neg Px$.) And worse, it appears to be a quite complicated matter to spell out the *necessary* conditions on reflexivity and transitivity in *DPL*. In *EDPL* things are more tractable. Modulo definedness, reflexivity is preserved:

Observation 15 $\phi \models \phi$ iff defined

A formula entails itself if one can felicitly process it twice. This is a necessary and sufficient condition on reflexivity. Notice that this logical simplicity might be conceived of as a practical disadvantage, because one might want to have repeated access to the contents of a formula which has been assumed or derived in an inference, irrespective of the question whether its reiteration brings about undefinedness. But things aren't that bad. Any (defined) formula in an *EDPL* inference can be reiterated felicitously—with preservation of content—by an appropriate substitution of *bound* variables (cf., section 3.3). Notice that nothing similar to *this* fact holds in *DPL*, for an assumption Fx can be discharged, or annihilated, in *DPL* by a later assumption $\exists x Gx$.

Definedness also demarcates transitivity:

Observation 16 If $\phi \models \psi$ and $\psi \models \chi$, then $\phi \models \chi$ iff defined

We can conclude χ from ϕ if we can conclude it from a logical consequence ψ of ϕ, iff, after ϕ has been processed, a domain of variables is spoken about which is a proper domain for the interpretation of both ψ and χ. The definedness condition effectively excludes the (troubling) cases where the mediating formula ψ introduces variables which are free in χ.

3.3 Definedness and deduction

In this last subsection we will be concerned with the central role that definedness plays in *EDPL*. We have already seen that, on the one hand, definedness ensures update of information, and, on the other, that it is definedness which rules out examples which threaten classical inference patterns. In this section I will first present a syntactic characterization of definedness, and next argue that all of the dynamics of *EDPL* can be precisely located in its definedness conditions (viz., in Heim's felicity conditions). After that, the outlines of a deduction system for *EDPL* are sketched. From this sketch alone it can already be seen that, although *EDPL*'s Heimian variable conventions complicate matters, they substantially simplify things, too.

Of course, *EDPL* definedness of a formula can be characterized syntactically. In order to find out whether a formula ϕ is defined, we have to go through its subformulas ψ and keep track of three things: the free variables $fv(\psi)$, the introduced variables $iv(\psi)$ and the quantified variables $qv(\psi)$ of ψ. Clearly, the set of free variables must always be a subset of the domain of an input state s, and the intersection of this domain with the set of quantified variables must be empty. By means of the following table we can figure out what the three sets are (here, $av(\phi)$ is the set of active variables in ϕ, the union of $fv(\phi)$ and $iv(\phi)$):

χ	$fv(\chi)$	$iv(\chi)$	$qv(\chi)$	conditions
$Rx_1 \dots x_n$	$\{x_1, \dots, x_n\}$	$\emptyset$	$\emptyset$	
$\neg\phi$	$fv(\phi)$	$\emptyset$	$qv(\phi)$	
$\exists x\phi$	$fv(\phi) \setminus \{x\}$	$\{x\} \cup iv(\phi)$	$\{x\} \cup qv(\phi)$	if $x \notin qv(\phi)$
	V	V	V	otherwise
$\phi \wedge \psi$	$fv(\phi) \cup (fv(\psi) \setminus iv(\phi))$	$iv(\phi) \cup iv(\psi)$	$qv(\phi) \cup qv(\psi)$	if $av(\phi) \subseteq \overline{qv(\psi)}$
	V	V	V	otherwise

Notice that if the side conditions are overruled, that is, if a subformula ψ of ϕ is necessarily undefined, then we equate the sets $fv(\psi)$, $iv(\psi)$ and $qv(\psi)$ with the full set V of variables. In this case V, of course, is not really the set of free, introduced and quantified variables of ψ. The reason for taking V here is that we use the sets $fv(\psi)$, $iv(\psi)$ and $qv(\psi)$ to determine whether ϕ is defined, and if all three sets of variables for a given subformula are equal to V, then the three sets of the whole formula ϕ are also equal to V, and, thus, ϕ, can be seen to be undefined, too:

Observation 17 ϕ is defined for s iff $fv(\phi) \subseteq D(s) \subseteq \overline{qv(\phi)}$

Clearly, if x occurs free in ϕ, and if a state s is not defined for x, then the interpretation of ϕ is undefined in s (cf., the requirement that $fv(\phi) \subseteq D(s)$). Undefinedness also occurs when s *is* defined for x, and x is quantified over in ϕ (whence the requirement that $D(s) \subseteq \overline{qv(\phi)}$). Moreover, we can see that if ϕ is necessarily undefined, then there can be no state s such that $fv(\phi) \subseteq D(s) \subseteq \overline{qv(\phi)}$. In that case, $fv(\phi) = V = qv(\phi)$, and clearly it is excluded that $V \subseteq \overline{V}$.

The specific kind of partiality of the *EDPL* interpretation function serves to set it apart from both *DPL* and a conceivable *static* system of interpretation. To the extent that *EDPL* really *is* dynamic, its dynamics resides in its definedness conditions because, as is shown in the appendix, under the assumption of definedness, *EDPL* operations can be defined in terms of the (commutative) product and ground operator. On the other hand, we also find that, under the assumption of definedness, *DPL* and *EDPL* are equivalent systems.

Observation 18 $\models_{dpl} \phi$ iff $\models_{edpl} \phi$ if defined

(This observation can be proved using the lemma:

Lemma If $fv(\phi) \subseteq D(i) \subseteq \overline{qv(\phi)}$ and $i \leq_V g$, then
1. if $j \in \{i\}[\![\phi]\!]_{M,edpl}$, then $\exists h$: $j \leq_V h$ and $h \in \{g\}[\![\phi]\!]_{M,dpl}$
2. if $h \in \{g\}[\![\phi]\!]_{M,dpl}$, then $\exists j$: $j \leq_V h$ and $j \in \{i\}[\![\phi]\!]_{M,edpl}$

This lemma in its turn is easily proved by induction on the complexity of ϕ.)

With the deduction theorem in mind, the above observation tells us that, under the assumption of definedness, the *DPL* and *EDPL* notions of dynamic entailment coincide. So, if we put things together, it appears that *EDPL* has located *DPL*'s dynamics in its definedness conditions.

To conclude this section on *EDPL* I will now sketch some inference rules for *EDPL*. They are part of a natural deduction system upon which I hope to report on another occasion. As usual, such a system specifies rules for building derivations, finite sequences of formulas $\phi_1, \ldots, \phi_n$ such that for $0 < j \leq n$, if ϕ_j is not an assumption then $\phi_1, \ldots \phi_{j-1} \models \phi_j$, this, however, under the additional assumption (which will be explained in a moment) that so-called 'linked' variables have identical values. Since *EDPL* derivations are conceived of as sequences of formulas, i.e., texts, they must conform to the *EDPL* variable conventions. In order to cope with these, more familiar inference rules have to be amended in two, related, ways. It has already been noticed that we cannot in general re-use a (subformula of a) formula which had been introduced earlier in the course of a derivation. This might render the derivation undefined. For this reason a 'proper copy' of such a (sub-)formula must be used, where a proper copy of a formula ϕ is obtained by substituting variables quantified over in ϕ by variables which are not in use at the relevant stage of the derivation. Thus, we can reiterate formulas without going undefined.

However, there is something more to this. If a proper copy ϕ' of a (sub-)formula ϕ is used, and if ϕ involves the introduction of a variable x, which is replaced by a variable y in ϕ', then we do want the information we derive about y to be of equal application to x, and the other way around. For this reason, we also have to keep track of the substitutions which have occurred in a derivation, and conceive of the variables which have been substituted for one another as being linked. All the time, such linked variables are assumed to have identical values.

A typical deduction rule, which displays both features, is the one for reiteration. The rule looks like this:

$$\begin{array}{rll} & \vdots & \\ m. & [g]\phi & \\ & \vdots & \\ n. & [f]\phi & [\text{reit.}, m] \end{array}$$

Here, f and g are injections which substitute variables in ϕ, and which obey the following requirements (which, by the way, can always be satisfied). The domain of f consists of the variables *quantified over* in ϕ and these are replaced by variables which are not active at step $n-1$ and which are not free in ϕ. By this substitution operation, $[f]\phi$ is guaranteed to be defined at step n, if ϕ was defined at step m (or, for that matter, if $[g]\phi$ was). The injection g may replace variables *free* in ϕ. The only condition is that, at step $n-1$, any free variable y in ϕ has been *linked* to $g(y)$, which means that the variables $g(y)$ and y in the course of the deduction in effect have been identified. Finally, it is recorded that every variable which has been *activated* by ϕ is explicitly linked to $f(x)$. Thus, an introduced variable x is in effect identified with $f(x)$.

Apart from the substitutions which are meant to ensure definedness and preservation of information, the conjunction and negation rules are pretty standard. So, for example, we use the following negation rule (one of the four dealing with $\neg$):

$$\begin{array}{rll} & \vdots & \\ m. & \neg\phi & \\ & \vdots & \\ m'. & [f]\phi & \\ & \vdots & \\ n. & \bot & [E\neg, m, m'] \end{array}$$

Here, f is another variables substituting injection, the domain of which is required to consist of the variables quantified over in ϕ. These observations then may suffice to give an impression of a deduction system for *EDPL*.

Probably the most striking thing about *EDPL* deduction is that it can do without rules for quantifiers. If we, as has already been discussed in section 3.2, conceive of quantifiers as atoms and of quantified formulas as conjunctions, then the quantifier rules are special cases of the conjunction rules. Moreover, assuming that the domains of individuals need not be non-empty, we can even do without introduction rules for quantifying atoms. In that case, any atom $\exists x$ in a derivation must be an assumption, or derive from an assumption. (For instance, from $\exists y$ one may deduce, by reiteration, $[f]\exists y = \exists x$ for any variable $x = f(y)$ not yet active. Clearly, if, by the declaration $\exists y$ that y has a value, the non-trivial claim is made that something exists, then we are also allowed to declare that x has a value, provided x is appropriately chosen.) It is remarkable indeed that, although the distinctive feature of *EDPL* is its dynamic existential quantifier, it needn't be especially dealt with in the deduction system. (In fact it is dealt with by the conditions on the other rules which guarantee a proper use of variables.) Of course, all this must not be taken to imply that the existential quantifier is semantically vacuous. Like an application of Heim's novelty condition on the use of x, a quantifying atom $\exists x$ has the semantic effect of extending the domain of an information state. By the same token it claims that something exists, so, semantically speaking, it has real existential import.

Although we can do without rules for quantifiers, it is interesting to see how they look like as *derived* rules of inference. To conclude this section I present the rules for the universal quantifier. Recall that, if $\exists x\phi$ is conceived of as a conjunction $\exists x \wedge \phi$, then $\forall x\phi$ is short for $\exists x \rightarrow \phi$. Using the rules for $\wedge$ and $\neg$, we then can determine the following derived rules for $\forall x$. The first differs mainly from the standard rule in the substitution of bound variables which may be required to preserve definedness:

$$\begin{array}{rll} & \vdots & \\ m. & \forall x\phi & \\ & \vdots & \\ n. & [f][z/x]\phi & [E\forall, m] \end{array}$$

provided z is active at step $n-1$, and f gives an appropriate substitution of the bound variables in ϕ. The (derived) rule for the introduction of the universal quantifier is the following:

$$\begin{array}{rll} & \vdots & \\ m. & \exists x & [\text{ass.}] \\ & \vdots & \\ n-1. & \phi & \\ n. & \forall x\phi & [I\forall, m] \end{array}$$

if defined (that is, provided that the free variables in $\forall x\phi$ are all active at step $m-1$). Notice that this introduction rule again mimics the standard one. In our dynamic setting, the applicability conditions of this rule are simpler, though. To conclude $\forall x\phi$ in a classical system, it must be required that we can derive $[z/x]\phi$ for a fully arbitrary z. But such a condition is automatically satisfied if we can derive ϕ on the mere assumption that x has a value. In *EDPL*, if one can prove ϕ on the assumption that x has a value, then ϕ will hold for all values of x (if any).

4 Quantification

One of the motivations for the idea that indefinites are free variables derives from their seemingly 'chameleontic' behavior in various adverbially quantified contexts. In an adverbially quantified construction, the quantifying head can be taken to quantify, unselectively, over the possible values of all the indefinites in its restrictive clause, and, thus, it appears to bind them. Now, if a term can be bound by another operator, what else, one might think, can it be then a (free) variable?

In this section I want to show that upon their dynamic interpretation existential quantifiers already supply (adverbial) quantifiers with the required domains of quantification. In *EDPL* it is possible to give a perspicuous and uniform formulation of the various dynamic forms of (adverbial and adnominal) quantification which are found in the literature. The quantifiers are assigned so-called 'internally dynamic' interpretations, those that account for anaphoric relationships between antecedent noun phrases in the restriction of the quantifiers and anaphoric pronouns in their nuclear scope. I do not discuss 'externally dynamic' readings, in which antecedent noun phrases in the restriction or in the scope of a quantifier license anaphora beyond the scope of the quantifier.

4.1 Unselective quantification

Lewis 1975 argues that in many cases adverbs of quantification (like *always*, *sometimes*, *usually*) unselectively quantify over the values of 'free variables' (stemming from indefinite noun phrases) in their restrictive clause. The examples Lewis discusses are of the form *Sometimes/usually/always if* x *is a man, if* y *is a donkey, and if* x *owns* y, x *beats* y and, clearly, these are paraphrases in the logician's idiom of sentences like:

(15) Sometimes/usually/always if a man owns a donkey, he beats it.

Lewis points out that the quantifying adverbs *sometimes/usually/always* in fact quantify over the 'cases' that verify the restrictive clause. These cases are the admissible assignments of values to the variables that are free in the restriction, or, equivalently, the tuples of individuals that are possible values of these variables. So, the cases that verify the restriction x *is a man,* y *is a donkey and* x *owns* y are the maps from x to a farmer and from y to a donkey the farmer owns (i.e., the pairs consisting of a farmer and a donkey he owns). The adverbial quantifier quantifies over these cases, i.e., pairs of individuals in the examples above. If the head is *sometimes*, as in *Sometimes, if a farmer owns a donkey, he beats it*, the sentence says that some pairs consisting of a farmer and a donkey owned, are pairs of which the first element beats the second. If *always* is the head, the sentence tells us that every pair of a farmer and a donkey he owns stands in the *beat* relation. And with *usually*, we get that most pairs that consist of a farmer and a donkey owned, stand in the beat relation. (Clearly, if in a quantified construction $A(\phi)(\psi)$ the restriction ϕ contains three (or n) free variables or indefinites, the unselective adverbs quantify over triples (or n-tuples) of individuals.)

This part of the story about adverbs of quantification is formalized quite elegantly in *FCS*, *DRT*, and also in *DPL*. In these systems, unselective quantification about the values of any number of free variables is, uniformly, cast in terms of quantification about verifying assignments. *EDPL*, too, allows a straightforward interpretation of unselectively quantifying adverbs. For instance, the formula $Always(\phi)(\psi)$ tests, given an initial assignment i, whether all extensions of i that verify ϕ also satisfy ψ. Similarly, the formula $Sometimes(\phi)(\psi)$ tests whether some extensions that verify ϕ with respect to i satisfy ψ, and the formula $Never(\phi)(\psi)$ tests whether no extension of i that verifies ϕ satisfies ψ. I restrict myself to the adverbs (and determiners, cf., section 4.3) that satisfy the constraints of extension, quantity and conservativity.* For any such adverb of quantification A, with its usual set-theoretic interpretation $[A]$, the interpretation is defined as follows:

Definition 9 (Adverbs of quantification (symmetric))

$$s[\![A(\phi)(\psi)]\!] = \{i \in s \mid [A](\{j \mid i \leq j \;\&\; j \in s[\![\phi]\!]\})(\{j \mid j < s[\![\phi]\!][\![\psi]\!]\})\}$$

In a state s, for each assignment i in s, a symmetric adverbial quantifier A tests whether $[A]$ applies, first, to the set of extensions of i in the update of s with the restriction of the adverb, and, second, to the set of assignments that also verify the nuclear scope of the adverb. Given the logical restrictions on the adverbs stated above, quantification over extensions of variable assignments amounts to quantification over the values of the variables in these extensions, which, in effect, comes down to quantification over the sets of tuples of individuals introduced in the restriction of the adverb.

For example, if we interpret *If a farmer owns a donkey he always beats it* in a state s in *EDPL*, we get $s[\![Always(\exists x(Fx \wedge \exists y(Dy \wedge Oxy)))(Bxy)]\!]$. This is the set of assignments i in s such that on every extension of i to x and y, if the value of x is a farmer who owns a donkey which is the value of y, then the value of x beats the value of y. In other words, this formula tests whether all pairs of a farmer and a donkey he owns are pairs of which the first element beats the second element. A second example is *If a man gives her a present, she usually thanks him for it*, which is translated as $Usually(\exists y(My \wedge \exists z(Pz \wedge Gyzx)))(Txyz)$. The interpretation of this formula in a state s yields a state consisting of all those assignments i in s that assign an individual to x that renders thanks in most cases in which a man gives her a present.

There are some interesting correspondences between the sentential connectives of *EDPL* and adverbs of quantification (as before, $\downarrow\phi$ indicates the static closure $\neg\neg\phi$ of ϕ):

Observation 19

$$\begin{array}{lcl} Sometimes(\phi)(\psi) & \Leftrightarrow & \downarrow(\phi \wedge \psi) \\ Never(\phi)(\psi) & \Leftrightarrow & \neg(\phi \wedge \psi) \\ Always(\phi)(\psi) & \Leftrightarrow & (\phi \rightarrow \psi) \\ Not\_Always(\phi)(\psi) & \Leftrightarrow & \neg(\phi \rightarrow \psi) \end{array}$$

Conjunction (disregarding its external dynamics) and implication appear to fit in the more general scheme of adverbial quantification. A sentence *Sometimes, if a farmer owns a donkey he beats it* has the same truth conditions as the conjunction *A farmer owns a donkey. He beats it.* Furthermore, it turns out that the sentence *Always, if a farmer owns a donkey he beats it* is equivalent to the sentence *Every farmer beats every donkey he owns*, and the sentence *If a farmer owns a donkey he never beats it* to the sentence *No farmer who owns a donkey beats it*.

---

* Let Q be a quantifier that assigns any domain of individuals E a binary relation Q_E between subsets of E. Then Q satisfies extension iff always if $A, B \subseteq E \subseteq E'$: $Q_E(A)(B)$ iff $Q_{E'}(A)(B)$. Q satisfies quantity iff always if π is a bijection from E to E' and $A, B \subseteq E$: $Q_E(A)(B)$ iff $Q_{E'}(\{\pi(a) \mid a \in A\})(\{\pi(b) \mid b \in B\})$. Q is conservative iff always if $A, B \subseteq E$: $Q_E(A)(B)$ iff $Q_E(A)(A \cap B)$. In van Benthem 1986, Ch. 1,2 these constraints are used to single out the determiners that qualify as '(logical) quantifiers'. Adverbs and determiners that do not observe all three constraints can be treated within the *EDPL* framework, but they deserve a special treatment.

Here we come touch upon another point of interest. The two indefinites (existential quantifiers) in *Always, if a farmer owns a donkey he beats it* ($Always(\exists x(Fx \wedge \exists y(Dy \wedge Oxy)))(Bxy)$) have universal import. This example can be used to say something about all farmers and donkeys owned. However, this shift in quantificational force does not derive from the fact that the indefinites (quantifiers) are free variables in logical form, and that they are bound by the quantifying adverb. Rather, they have their ordinary dynamic semantic contribution. They introduce possible assignments of values to the variables they have been associated with, and these are quantified over. As a matter of fact, their 'chameleontic' nature does not only appear from their dynamic import. Also the sentence *If there is a man inside, John is home* means the same as *For every man this holds: if he is inside John is home*, and not the same as *There is a man such that if he is inside John is home*. It is a fact about the classical predicate logic that $(\exists x\phi \rightarrow \psi)$ is equivalent to $\forall x(\phi \rightarrow \psi)$ (provided x is *not* free in ψ, of course) and not to $\exists x(\phi \rightarrow \psi)$. But this, of course, is no reason to conceive of the existential quantifier from ordinary predicate logic as not being a quantifier after all.

4.2 Asymmetric quantification

Adverbs of quantification do not always unselectively quantify over the values of *all* variables which are introduced in their restriction. Several authors (Bäuerle 1985; Root 1986; Rooth 1987; Kadmon 1987; Kadmon 1990, see also Heim 1990; Chierchia 1992) have discussed examples in which adverbial quantifiers involve quantification over the values of a proper *subset* of the introduced variables. Following Rooth and Kadmon, I call this kind of quantification *asymmetric*. We find it in the following sentences:

(16) If a farmer owns a donkey, he is usually rich.
(17) If a DRUMMER lives in an apartment complex, it is usually half empty.
(18) If a drummer lives in an APARTMENT COMPLEX, it is usually half empty.

On its most natural reading, the adverb *usually* in the first example quantifies over farmers who own a donkey and not over farmer-donkey pairs. The sentence says that most farmers who own a donkey are rich. (If the adverb is taken to quantify unselectively, we get the different reading that for most pairs consisting of a farmer and a donkey he owns, it holds that the farmer is rich.) The second and third example are different for a similar reason. In the second example, with focal stress on *drummer*, we (may) find quantification over apartment complexes in which a drummer houses. The example then states that most apartment complexes where a drummer lives are usually half-empty. In the third example, where we find focal stress on *apartment complex*, the adverb may be taken to quantify over drummers. On this reading, the sentence says that most drummers that live in an apartment complex live in an half empty apartment complex.

Examples like these pose a problem for the *DRT* and *DPL* analysis of adverbs of quantification, a problem dubbed the 'proportion problem' in Kadmon 1987. Of course, it is easy to *annul* the introduction of certain variables in a restriction and to quantify over the values of the remaining variables. However, we don't really want to *annul* the introduction of certain variables in the restriction of a quantifying adverb when the values of some variables are selected for quantification, we just want the adverb to *neglect* the (alternative) valuations of the other ones. For this reason, Root 1986 suggests that asymmetric adverbial quantifiers do not discriminate between assignments that only differ from one another in their assignment to non-selected variables. This implies that asymmetric quantifiers in fact quantify over equivalence classes of assignments, each one of which consists of assignments that agree on the values of the selected variables. For instance, in the example *If a drummer lives in an* APARTMENT COMPLEX, *it is usually half empty*, the adverb *usually* selects (the variable associated with) *a drummer*, and in effect quantifies over a set of sets of assignments, each element of which is a set of pairs any first element of which is one and the same drummer and any second element any apartment complex the drummer lives in.

In *EDPL*, asymmetric adverbs more naturally fit into the general scheme of adverbial quantification. As we have seen unselective adverbial quantification is satisfactorily dealt with in *EDPL* in terms of quantification over the extensions of an assignment i in an input state s which are an element of the update of s with the restriction ϕ of an adverb. Clearly, the extensions in the update with the restriction are extensions of i to all variables introduced in the restriction and therefore unselective quantification involves quantification over the values of all these variables. Now, in order to account for asymmetric quantification, we do not consider the full extensions of i which are an element of the update with the restriction, but partial extensions which *survive* in the update with the restriction. An asymmetric quantifying adverb quantifies over extensions of i to the selected variables, viz., those that survive the update with the restriction of the adverb, and it relates this set of extensions to the set of extensions that also survive in further update with the nuclear scope of the adverb. In *EDPL* we can thus quantify over the values of *some* variables introduced in the restriction of an adverb, neglect the valuation of others, without denying these other variables the ability to serve as antecedents for anaphoric pronouns in the adverb's nuclear scope.

Asymmetric adverbs of quantification are assumed to come with a set of selection indices X the values of which the adverb quantifies over. It is required, again, on pain of undefinedness, that an asymmetric adverbial quantifier A_X effectively quantifies over the values of all the variables in X, so X must be a subset of the variables introduced in the restriction. Its interpretation is defined as follows:

Definition 10 (Adverbs of quantification (asymmetric)) If $X \subseteq iv(\phi)$
$$s[\![A_X(\phi)(\psi)]\!] = \{i \in s \mid [A](\{j \mid i \leq_X j \;\&\; j < s[\![\phi]\!]\})(\{j \mid j < s[\![\phi]\!][\![\psi]\!]\})\}$$

As an illustration, let us briefly consider two related examples.

(19) If a man gives her a PRESENT, she usually thanks him for it
$Usually_{\{y\}}(\exists y(My \wedge \exists z(Pz \wedge Gyzx)))(Txyz)$

This sentence requires of an assignment i in a state of evaluation s that most extensions j of i with a valuation for y such that $j(y)$ gives $i(x)$ a present, are also valuations such that $i(x)$ thanks $j(y)$ for a present $j(y)$ gives to her. So, interpreted in a state s, this example preserves all those i in s that assign x an individual that renders thanks to most men that give her a present, irrespective of the number of presents given.

(20) If a MAN gives her a present, she usually thanks him for it

$Usually_{\{z\}}(\exists y(My \wedge \exists z(Pz \wedge Gyzx)))(Txyz)$

When interpreted in a state s, this example preserves all those i in s that assign x an individual that renders thanks for most presents given by a man, irrespective of the number of men that give it.

As expected, unselective quantification is a borderline case of asymmetric quantification:

Observation 20 $A_{iv(\phi)}(\phi)(\psi) \Leftrightarrow A(\phi)(\psi)$

Furthermore, we find the following equivalences:

Observation 21 If $X \subseteq iv(\phi)$

$$\begin{array}{llll} Sometimes_X(\phi)(\psi) & \Leftrightarrow & Sometimes(\phi)(\psi) & \\ Never_X(\phi)(\psi) & \Leftrightarrow & Never(\phi)(\psi) & \\ Always_\emptyset(\phi)(\psi) & \Leftrightarrow & (\phi \hookrightarrow \psi) & (\Leftrightarrow (\downarrow\phi \rightarrow (\phi \wedge \psi))) \\ Not\_Always_\emptyset(\phi)(\psi) & \Leftrightarrow & \neg(\phi \hookrightarrow \psi) & (\Leftrightarrow (\downarrow\phi \wedge (\phi \rightarrow \psi))) \end{array}$$

For the adverbs *sometimes* and *never* it makes no difference whether or not they quantify over a selection of the variables introduced in their restriction. This is as it should be, since there seems to be no evidence that there are distinct asymmetric readings of these adverbs. On the other hand, for the adverbs *usually* and *always*, it does make a difference whether or not they select variables for asymmetric quantification, and which variables they select.

The weak implication ($\hookrightarrow$) which has been addressed in section 2.3, appears to be a borderline case of asymmetric adverbial quantification, i.e., universal quantification over the values of an empty set of selection indices. So, it appears that this weak implication does not constitute a really different notion of implication of its own, but that it fits into the more general scheme of adverbial quantification as one of the many forms of asymmetric quantification. This can be further substantiated by slightly varying the 'dime' implication again:

(21) If I have a dime in my pocket, I throw it in the parking meter.
$(\exists y(Dy \wedge Piy) \hookrightarrow Tiy) \Leftrightarrow Always_\emptyset(\exists y(Dy \wedge Piy))(Tiy)$

(22) If a man has a dime in his pocket, he throws it in the parking meter.
$Always_{\{x\}}(\exists x(Mx \wedge \exists y(Dy \wedge Pxy)))(Txy)$

On its most natural reading, the first example says that if I have a dime in my pocket, then I throw one in the meter, and this reading is captured by interpreting the conditional sentence as one of weak implication (universal quantification over the values of an empty set of selection indices). However, the second example, which is a minor variant of the first, is most likely interpreted as stating that every man who has a dime, throws one in the meter. Neither the weak, nor the strong, reading of the implication gives us this. In fact, this example has a mixed strong and weak reading: strong with respect to the men who are faced with a parking meter and who have a dime in their pocket; weak with respect to the number of dimes they throw in. Now, if the weak implication is supposed to constitute a form of implication of its own and if it is used to account for the first example, we still need an explanation of this second, related, example, which, on the preferred reading, does not exemplify a purely weak or strong implication. If, on the other hand, we fit both these sentences into the scheme of asymmetric adverbial quantification, both get assigned proper readings in a uniform way.

To conclude, we see that both unselective and asymmetric adverbial quantification can be neatly accounted for in terms of quantification over partial assignments. The equivalence classes of assignments which Root needs to deal with asymmetric quantification can be properly conceived of as assignment restrictions, partial assignments, the kind of entities in terms of which the semantics of *EDPL* has been defined in the first place. Notice that the present definition of asymmetric quantification heavily relies on the notion of an assignment surviving the update with the restriction of an adverb, which only makes sense in a proper update semantics like *EDPL*.

4.3 Adnominal quantification

The most general scheme of quantification so far is that of asymmetric adverbial quantification. As appears from fact 20, symmetric adverbial quantification is a borderline case of asymmetric quantification, viz., a form of quantification in which all the variables of an adverb's restriction are selected. It may be interesting to notice that several recent accounts of dynamic *adnominal* quantification can be covered by the asymmetric scheme of quantification, too. Adnominal quantification can be conceived of as an asymmetric form of quantification in which one variable is selected, the one bound by the determiner. Again, I will only be concerned with the internal dynamics of the quantified structures at issue, that is, with anaphoric relationships between indefinites occurring in the restriction of a quantifier and anaphoric pronouns in its nuclear scope. For a treatment of the external dynamics of adnominal quantifiers see, for instance, van den Berg 1994.

The following definition of dynamic adnominal quantification fits in the general scheme of quantification developed so far. The interpretation defined here corresponds to the so-called $\exists$ or 'weak' interpretation of quantifiers (as presented in Chierchia 1992; van Eijck and de Vries 1992, respectively). Let D be a binary quantifier which has $[D]$ as its usual set-theoretic interpretation, then:

Definition 11 (Binary quantifiers) If $x \notin D(s)$

$$s[\![Dx(\phi)(\psi)]\!] = \{i \in s \mid [D](\{j \mid i \leq_{\{x\}} j \;\&\; j < s[x][\![\phi]\!]\})(\{j \mid j < s[x][\![\phi]\!][\![\psi]\!]\})\}$$

On this definition of adnominal quantification, a binary quantifier Dx quantifies over the possible valuations of a single variable x and, hence, over the individuals that satisfy the requirements on x in the restriction of D. (That is, on the assumption of extension, quantity and conservativity, cf., the remarks in section 4.1). Let us consider one example:

(23) Most men who gave her a present had packed it up
$Most\,y(\exists z(Pz \wedge Gyzx))(Uyz)$

The interpretation of this example in a state s yields a state consisting of assignments $i \in s$ such that most extensions of i to y under which the value of y is a man who gave a present to the value of x, are extensions under which the value of y has packed up a present he gave to the value of x. Put more simply, the interpretation of this example in a state s preserves those assignments i in s which assign an individual z to x such that most men who gave a present to z gave her a present they had packed up.

As has already been indicated, adnominal quantification corresponds to asymmetric quantification over a singleton set of selection indices. Suppose that adverb A and determiner D have the same set-theoretical interpretation. Then:

Observation 22 $A_{\{x\}}(\exists x\phi)(\psi) \Leftrightarrow\ D\,x(\phi)(\psi)$

We also find the following relations between our binary interpretations of the classical four quantifiers and their standard unary paraphrases:

Observation 23

$$\begin{array}{ll}
An\,x(\phi)(\psi) & \Leftrightarrow {\downarrow}\exists x(\phi \wedge \psi) \\
No\,x(\phi)(\psi) & \Leftrightarrow \neg\exists x(\phi \wedge \psi) \\
Every\,x(\phi)(\psi) & \Leftrightarrow \forall x(\phi \hookrightarrow \psi) \\
Not_Every\,x(\phi)(\psi) & \Leftrightarrow \neg\forall x(\phi \hookrightarrow \psi)
\end{array}$$

We see that the binary determiners *a(n)* and *no* have the same truth-conditional content as their usual first order paraphrases. However, the determiners *every* and *not_every* give rise to two different interpretations. If we translate *every* by means of the first order universal quantifier, as with $\forall x(\phi \rightarrow \psi)$, we get the strong interpretation of the donkey sentence *Every farmer who owns a donkey beats it*, with truth-conditions requiring every farmer to beat every donkey he owns. If *every* is treated as a binary determiner, the weak reading results. Under this interpretation, a sentence *Every man who has a dime puts in the parking meter* claims that every man who has a dime puts one in the meter. Obviously, both readings are required.

The ambivalence of *every* is reflected in its negation *not_every*. If *every* is read strongly, a quantified structure *Not every A B* translates as $\neg\forall x(A'x \rightarrow B'x) \Leftrightarrow {\downarrow}\exists x(A'x \wedge \neg B'x)$. This seems appropriate for the negation of the strong donkey sentence. The sentence *Not every man who owns a donkey beats it* has the truth conditions that there is a man who owns a donkey which he does not beat. On the other hand, if *every* is read weakly, as a binary determiner, then *Not every A B* translates as $Not\_every\,x(A'x)(B'x)$, which is equivalent to $\neg\forall x(A'x \hookrightarrow B'x) \Leftrightarrow {\downarrow}\exists x({\downarrow}A'x \wedge (A'x \rightarrow B'x))$. This seems appropriate for the negation of the weak dime sentence. For the sentence *Not every man who had a dime threw it in the meter* to be true, there must be a man who did have one or more dimes but who didn't throw any one of them in the meter. For the dime sentence, this 'weak' reading seems the intuitive one.

The different readings of quantified donkey sentences which have been called 'weak' and 'strong' readings have been widely discussed in the literature, cf., among many others, Bäuerle 1985; Root 1986; Rooth 1987; Schubert and Pelletier 1989; Heim 1990; van Eijck and de Vries 1992; Chierchia 1992; Kanazawa 1994. The last three of these papers are representative for three principled kinds of analyses of the phenomena at issue. I want to end this section by bringing out the main differences between these three approaches and mine. Although the labels 'strong' and 'weak' are misleading to some extent (cf., Kanazawa 1994), I will adapt to what seems to be current practice and keep on using them.

van Eijck and de Vries 1992 presents the most liberal proposal for dealing with quantified donkey sentences. (In terms of empirical coverage, this approach is very similar to the one already dubbed "the best solution currently available" in Heim

1990, p. 163.) In van Eijck and de Vries' system, any quantified structure is assigned a strong and a weak reading. This, however, is too liberal an approach, since, for instance, the determiners *no* and *an* do not appear to have a strong reading. Upon the 'strong' reading of the determiner, a sentence *An/No farmer who owns a donkey beats it* would be assigned the truth conditions that there is a/no farmer who owns a donkey and who beats every donkey he owns. Especially for the determiner *no* these are quite counterintuitive results (cf., also, Chierchia 1992, p. 154). On its strong reading the sentence *No farmer who owns a donkey beats it* would be *true* in a situation where there are lots of farmers owning lots of donkeys, and where each farmer beats all of his donkeys except one. Common opinion seems to have it that the sentence should be *false* then.

Restrictions on the distribution of weak and strong readings of quantified donkey sentences readings have been discussed in Rooth 1987; Chierchia 1992; Kanazawa 1994, the last two of which present more or less principled restrictions. The problem with these last two proposals is that they seem to give us too *few* readings. Chierchia 1992 employs the conservativity of natural language determiners to define a dynamic interpretation of determiners which can be conceived of as conservative at a dynamic level of interpretation. It so happens, however, that upon Chierchia's conservativity restriction, he only gets us the weak readings of determiners. The strong reading of the universally quantified donkey sentence, which, one might say, lies at the heart of the development of *DRT* and *FCS*, falls beyond the scope of this analysis. We may have to add here that Chierchia does take pains to get strong readings of quantified structures, too, this by proposing an additional E-type interpretation procedure for pronouns. However, as Kanazawa 1994, p. 156–7 also remarks, it appears that it is going to give us, again, the *unwanted* strong readings of quantifiers which are upward monotonic in both arguments (like the indefinite determiner *an*).

Kanazawa 1994 relates the possible distribution of weak and strong readings to the monotonicity properties of the determiners involved. On his, deliberately biased account, determiners require a weak interpretation if they are uniformly upward monotonic ($\uparrow$MON$\uparrow$) or uniformly downward monotonic ($\downarrow$MON$\downarrow$).[†] Determiners with opposite monotonicity properties ($\downarrow$MON$\uparrow$ or $\uparrow$MON$\downarrow$) favor a strong interpretation. As can be read off his scheme for defining weak and strong readings, only on this distribution of weak and strong readings dynamic determiners inherit the monotonicity properties of their static counterparts. As far as empirical coverage is concerned we find the following difference between Kanazawa's proposal and the one put forward in this paper. According to Kanazawa's principle a binary determiner with different left and right monotonicity properties has a strong preference for the strong dynamic interpretation, whereas on my proposal they are uniformly assigned weak readings. However, this difference may seem to be more substantial than it in fact is. On the one hand, notice, first, that, as Kanazawa himself observes, there are only few determiners of the kind mentioned, and, second, that the two most notorious examples (viz., *every* and *not_every*) in our system have already been predicted to have a strong interpretation too, viz., the one they get when they are conceived of

[†] A binary quantifier D is left upward (downward) monotonic, $\uparrow$ ($\downarrow$), iff always if $D(A)(C)$ and $A \subseteq B$ ($A \supseteq B$) then $D(B)(C)$; it is right upward (downward) monotonic, $\uparrow$ ($\downarrow$), iff always if $D(A)(B)$ and $B \subseteq C$ ($B \supseteq C$) then $D(A)(C)$.

as 'unary' quantifiers. On the other hand, Kanazawa does not completely exclude weak readings of these determiners. Rather, he conceives of the preference for strong readings here as a *default*, which can be overruled by pragmatic and other factors.

To summarize this section, let me first observe that I haven't made any observations on 'weak' and 'strong' readings which are not already found in the literature, and, furthermore, that no principled analysis of the distribution of 'weak' and 'strong' readings of quantified donkey sentences is offered here. It may even be conceived a coincidence that my unary (in terms of $\exists x$) and binary (according to the general scheme of adverbial quantification) quantifiers together get us precisely the sets of readings which have been put forward in the literature. But then it is striking indeed that my naive approach does so well, empirically speaking. Principled theories like those of van Eijck and de Vries, Chierchia and Kanazawa either give us too many readings ('strong' readings of *no* and *an*), or too few (not the strong, or not the weak reading of *every* and *not_every*). Of course, this doesn't settle the matter in any definitive way. To begin with, as much as there is no pertinent evidence for the non-existence of readings which the present system of *EDPL* does *not* derive (e.g., 'strong' readings of non-monotonic determiners), just as much it is not completely settled that the readings which *EDPL does* derive (their 'weak' readings) are real. Given this, then, we leave the subject for further research.

5 Conclusion

In this paper I have presented a reformulation of dynamic predicate logic into an update semantics along Heimian lines. I have developed a proper notion of information about the values of variables, and a semantics for the language of predicate logic which interprets predicate logical formulas as functions updating information about the values of variables. Throughout I have been dealing with extensional notions. Models of interpretation are extensional models, and, likewise, the type of information we have dealt with is extensional, too.

A last issue I want to raise in this paper is this. In Dekker 1993a I have shown that a system of dynamic semantics (in particular, Groenendijk and Stokhof's dynamic Montague grammar) provides a suitable framework for dealing with implicit arguments of nouns and verbs. Implicit arguments can be conceived of as existentially closed argument slots of the relevant lexical items, and by the dynamics of the existential quantifier they remain available for further specification. Thus, for instance, a relational noun like *sister* can be treated structurally in the same way as ordinary nouns are, and, still, a modifying phrase *of John* can be used to specify the implicit argument by its 'existential disclosure'. In the paper I gave programmatic formulations of the phenomena at issue, hoping to show that "any compositional reformulation of *DRT*" could accommodate them. The point, now, is that this is not too obvious for the one offered in this paper. Uncritical adoption of the treatment of implicit arguments in an *EDPL*-style system, might render undefined iterated use of one and the same lexical item, that is, if the item carries an (existentially closed) implicit argument.

Things are not that bad, though. Although my treatment of implicit arguments, as such, cannot be carried over to *EDPL*, one can think of several ways of merging the two proposals. For instance, we can propose to use the lexicon in a more dynamic fashion. By taking 'appropriate copies' of the translations of lexical items the required variable conventions can be observed. A more principled solution, however, appears to be the following. In Dekker 1994, I have developed an *EDPL*-style information system in which discourse referents are not labeled by variables, but which can be used equally well to model the update of information about the possible values of potential antecedent terms. The interesting thing is that in the ensuing system of *PLA*, no fresh-variable conventions are required. In fact, in that system, variables only play the role they play in ordinary predicate logic. Now, in this system, which has the advantageous semantic properties of *EDPL*, but which does not render iterated quantification over the value of one and the same variable undefined, my Dekker 1993a treatment of implicit arguments can be accommodated without problem. Clearly, we aren't any longer talking about the possible values of variables then.

***

To conclude, Groenendijk and Stokhof have shown that the interpretation of simple singular anaphoric relationships does not enforce a representational level of interpretation upon us. They have presented a dynamic, non-representational, interpretation of the language of predicate logic, which is concerned with information about the values of variables in discourse ('discourse markers'). The upshot of this paper is that it is advantageous and elucidating to incorporate Heimian variable conventions in a dynamic semantics which deals with information about the values of variables. I have been mainly concerned with the implications which the adoption of Heim's felicity conditions has for the type of information structure which is required to deal with anaphora, and for the logic of the updates of these information structures. As I hope to have shown, Heim's felicity conditions are more than a convenient pragmatic principle, more than something that forces 'practical good luck', so to speak. The pragma-/syntactico kind of felicity which Heim started out to work with, and which corresponds to Kamp's practice of always introducing 'fresh' variables, in fact delineates a neat semantical landscape, which accommodates logically attractive forms of information, update and deduction.

Appendix

In this appendix I will show that, under the assumption of definedness, *EDPL* interpretation can be defined in terms of the product and ground operator $\underline{\wedge}$ and $\vee$, which have been presented in section 1.2, and ordinary complementation $\overline{}$ (for $s \in S^X$, the complement $\overline{s}$ of s is $\top^X \setminus s$). Thus we will see that, under the definedness assumption, *EDPL* information update corresponds to the (commutative) meet of the lattice of *EDPL* information states. For the sake of convenience, an existentially quantified formulas $\exists x \phi$ is here conceived of as a conjunction $\exists x \wedge \phi$, with $s[\![\exists x]\!] = s[x]$ (if $x \notin D(s)$, undefined otherwise). Thus, both atomic formulas and quantifiers are conceived of as atoms, which we can assign (static) denotations. The denotations of complex formulas can be defined in terms of these:

Definition 12 (EDPL denotations)

$$
\begin{array}{ll}
[Rx_1 \ldots x_n] & = \{i \in D^{\{x_1,\ldots,x_n\}} \mid \langle i(x_1), \ldots, i(x_n)\rangle \in F(R)\} \\
[x = y] & = \{i \in D^{\{x,y\}} \mid i(x) = i(y)\} \\
[\exists x] & = \top^{\{x\}} \\
[\neg\phi] & = \overline{\bot^{fv(\phi)} \vee [\phi]} \\
[\phi \wedge \psi] & = [\phi] \wedge [\psi]
\end{array}
$$

An atomic formula with free variables $x_1, \ldots, x_n$ denotes a state with a domain $\{x_1, \ldots, x_n\}$, and which consists of assignments i which verify the formula. The quantifier $\exists x$ denotes the state of no information about the value of x. A negation $\neg\phi$ denotes a state of information about the free variables in ϕ. The common ground of $\bot^{FV(\phi)}$ and $[\phi]$ restricts the assignments in $[\phi]$ to the assignments to the free variables in ϕ which have an extension in $[\phi]$. The complement of this set is the set of assignments to the free variables in ϕ which have no extension in $[\phi]$. Conjunction is associated with our product operation. Now we can state the following fact:

Observation 24 $s[\![\phi]\!] = s \wedge [\phi]$ if defined

Proof: by induction on the complexity of ϕ. If ϕ is an atom the case is straightforward. If ϕ is a conjunction, $s[\![\phi]\!] = s \wedge [\phi]$ by the induction hypothesis and the fact that the product operation is associative. The only non-trivial step is the one dealing with negation. First observe that if ϕ is defined, then $D([\phi]) = av(\phi)$ and $fv(\phi) \subseteq D([\phi])$, so $i \in [\neg\phi]$ iff $i \in D^{fv(\phi)}$ & $i \not\leqslant [\phi]$ iff (by def. $\wedge$) $i \in D^{fv(\phi)}$ & $i \not\leqslant (\{i\} \wedge [\phi])$ iff (using the induction hypothesis) $i \in D^{fv(\phi)}$ & $i \not\leqslant \{i\}[\![\phi]\!]$. Now, if ϕ is defined for s, $s[\![\neg\phi]\!] = \{j \in s \mid j \not\leqslant \{j\}[\![\phi]\!]\}$. Notice that for any $j \in s$, j falsifies ϕ iff the restriction i of j to the free variables of ϕ does, so

$$
\begin{array}{ll}
s[\![\neg\phi]\!] = & \{j \in s \mid \exists i \leq j \colon i \in D^{fv(\phi)} \;\&\; i \not\leqslant \{i\}[\![\phi]\!]\} = \text{(as we just have seen)} \\
& \{j \in s \mid \exists i \leq j \colon i \in [\neg\phi]\} = \text{(by def. } \geqslant) \\
& \{j \in s \mid j \geqslant [\neg\phi]\} = \text{(since } fv(\phi) \subseteq D(s)\text{, by def. } \wedge) \\
& s \wedge [\neg\phi]
\end{array}
$$

The present result may serve to show that definedness prevents *EDPL* from collapsing into a static system of interpretation. But given this much, we might be tempted to reason as follows. Since the product operation is defined for all pairs of states, and since all *EDPL* atoms have a defined, static, denotation, why don't we adopt the definition of $[\![\ ]\!]$ in terms of $[\ ]$ and $\wedge$ and simply drop the definedness conditions? If we did, $\exists x_j Rx_1 \ldots x_j \ldots x_n$ and $Rx_1 \ldots x_j \ldots x_n$ would turn out to be fully equivalent (since, if $x \in D([\phi])$, $[\exists x\phi] = (\top^{\{x\}} \wedge [\phi]) = [\phi]$). So, in that case indefinites would really *be* free variables, which is reminiscent of "Lewis' philosophy".

Still, this doesn't seem to be an attractive alternative. For, firstly, if definedness conditions are dropped the free variables of a formula ϕ cannot be determined any longer on semantic grounds only, and, hence, the denotation of $\neg\phi$ is not any longer properly defined in terms of that of ϕ. So, if, on the same condition, we want to keep a proper, compositional notion of negation, we need to find another mechanism to determine the free variables of a formula (like, for instance, the one presented in Zeevat 1989). Secondly, and more importantly, dropping definedness conditions inevitably renders all three formulas $(\exists xFx \wedge \exists xGx)$, $(\exists xFx \wedge Gx)$, $(Fx \wedge \exists xGx)$ equivalent to $\exists x(Fx \wedge Gx)$. Now, if the objective is to translate natural language indefinites using existential quantifiers, and conceive of anaphoric pronouns as variables, this would indeed be quite a dubious way to proceed. Definedness ensures that variables are introduced *before* they are referred back to.

Acknowledgments

I owe many thanks to David Beaver, Jeroen Groenendijk, Herman Hendriks, Martin Stokhof, Frank Veltman and two anonymous referees of *Linguistics and Philosophy* for stimulating discussions and constructive comments. Earlier versions of this paper have been presented at the ASL/LSA conference in Santa Cruz and at the Eighth Amsterdam Colloquium, both in 1991, and at several smaller workshops in '92 and '93. The paper itself is drawn from my thesis Dekker 1993b, Ch. 5. Preliminary versions have appeared in Dekker, Paul and Martin Stokhof (eds.), 1992, *Proceedings of the Eighth Amsterdam Colloquium*, ILLC, University of Amsterdam, and in Cooper, Robin H. (ed.), 1993, *Integrating Semantic Theories*, ILLC, University of Amsterdam. The research for this paper was supported by the ESPRIT Basic Research Project 6852, Dynamic Interpretation of Natural Language, DYANA-2.

References

Barwise, J. and Cooper, R.: 1981, Generalized quantifiers in natural language, *Linguistics and Philosophy* 4, 159–220

Bäuerle, Rainer & Egli, U.: 1985, *Anapher, Nominalphrase und Eselsätze*, Sonderforschungsbericht 99, 105, Universität Konstanz, Konstanz

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

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

Dekker, P.: 1993a, Existential disclosure, *Linguistics and Philosophy* 16, 561–88

Dekker, P.: 1993b, *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 M. Harvey and L. Santelmann (eds.), *Proceedings from Semantics and Linguistic Theory IV*, pp 79–95, Department of Modern Languages and Linguistics, Cornell University

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

Groenendijk, J., Stokhof, M., and Veltman, F.: 1994, Update semantic for modal predicate logic, in R. H. Cooper and J. A. Groenendijk (eds.), *Integrating Semantic Theories II*, pp 135–226, ILLC, Amsterdam, Dyana-2 report R2.1.B

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, File change semantics and the familiarity theory of definiteness, in R. Bäuerle, C. Schwarze, and A. von Stechow (eds.), *Meaning, Use, and Interpretation of Language*, pp 164–89, de Gruyter, Berlin

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

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

Kadmon, N.: 1990, Uniqueness, *Linguistics and Philosophy* 13, 273–324
Kamp, H.: 1981, A theory of truth and semantic representation, in J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Formal Methods in the Study of Language*, pp 277–322, Mathematical Centre, Amsterdam, Reprinted in J. Groenendijk, T. Janssen, and M. Stokhof (eds.), *Truth, Interpretation and Information*. Foris, Dordrecht, 1984, pp. 2–41
Kamp, H. and Reyle, U.: 1993, *From Discourse to Logic*, Kluwer, Dordrecht
Kanazawa, M.: 1994, Weak vs. strong readings of donkey sentences and monotonicity inference in a dynamic setting, *Linguistics and Philosophy* 17, 109–158
Karttunen, L.: 1968a, *What Do Referential Indices Refer To*, paper P-3854, The RAND Corporation, Santa Monica, CA
Karttunen, L.: 1968b, *What Makes Definite Noun Phrases Definite?*, paper P-3871, The RAND Corporation, Santa Monica, CA
Karttunen, L.: 1976, Discourse referents, in J. D. McCawley (ed.), *Notes from the Linguistic Underground. Syntax and Semantics 7*, pp 363–85, Academic Press, New York
Landman, F.: 1986, *Towards a Theory of Information*, Foris, Dordrecht
Lewis, D.: 1975, Adverbs of quantification, in E. L. Keenan (ed.), *Formal Semantics*, pp 3–15, Cambridge University Press, Cambridge, U.K.
Root, R.: 1986, *The Semantics of Anaphora in Discourse*, Ph.D. thesis, University of Texas, Austin
Rooth, M.: 1987, NP interpretation in Montague grammar, File Change Semantics and Situation Semantics, in P. Gärdenfors (ed.), *Generalized Quantifiers*, pp 237–69, Kluwer, Dordrecht
Schubert, L. K. and Pelletier, F. J.: 1989, Generically speaking, or, using Discourse Representation Theory to interpret generics, in G. Chierchia, B. H. Partee, and R. Turner (eds.), *Properties, Types and Meaning II, Semantic Issues*, pp 193–268, Kluwer, Dordrecht
van Benthem, J.: 1986, *Essays in Logical Semantics*, Reidel, Dordrecht
van den Berg, M. H.: 1994, A direct definition of generalized dynamic qantifiers, in M. Stokhof and P. Dekker (eds.), *Proceedings of the Ninth Amsterdam Colloquium*, pp 121–40, ILLC, University of Amsterdam, Amsterdam
van Eijck, J. and de Vries, F.-J.: 1992, Dynamic interpretation and Hoare deduction, *Journal of Logic, Language and Information* 1, 1–44
Veltman, F.: 1995, Defaults in update semantics, *Journal of Philosophical Logic* 24, In Press
Vermeulen, C. F.: 1993, Sequence semantics for dynamic predicate logic, *Journal of Logic, Language and Information* 2, 217–54
Zeevat, H.: 1989, A compositional approach to discourse representation theory, *Linguistics and Philosophy* 12, 95–131

ILLC/Department of Philosophy
University of Amsterdam
The Netherlands
dekker@illc.uva.nl