Chapter 3
Information Update and Support

The previous chapter on *PLA* gives an impression of the impact of putting a simple idea to work in a well-defined setting. It presents an extension of ordinary predicate logic with pronouns, which are taken to be anaphorically related to existential quantifiers and shows the consequences to be quite manageable. However, the reader may be concerned that such extensions generate complications at any further level of generalization. This chapter together with the following one serve to eliminate those concerns.

In this chapter I lift the extensional system from the previous chapter to an intensional level, showing that some welcome opportunities and applications result. Such a lift enables a study of information update and information support. I focus on an update algorithm and a support calculus and show that these are well-behaved in a very intuitive sense. The intensional lift turns out to be unproblematic, precisely because the *PLA* system does not confuse the semantics with the pragmatics of information exchange.

Before I begin, however, it is expedient to start with a concise overview of some previous systems which have extended *DPL*, in its indeed dynamic spirit, and which have inspired our treatment of the relevant phenomena in this chapter and the following, adopting a static notion of meaning, with a dynamic notion of conjunction.

3.1 Coreference and Modality

Interpretation in *DPL* (Sect. 2.2), and in many of its off-spring, has been characterized as an update function on states of information about the values of variables. The selling metaphor is that certain noun phrases, in particular indefinite descriptions, set up discourse referents and this is covered by saying that they 'declare variables', in computer science jargon. If an indefinite noun phrase is used, it declares a variable as its marker, and in subsequent interpretation one may learn more about the possible values of that variable, or, rather, put more constraints on the values.

P. J. E. Dekker, *Dynamic Semantics*, Studies in Linguistics and Philosophy 91,
DOI: 10.1007/978-94-007-4869-9_3,

The relational format of *DPL*, in which interpretation is modeled as a relation between variables assignments, is not directly fit for the intuitive format of update of information. Interpretation in *DPL* is a *relation*, not a *function*, and variable assignments characterize total information about the values of variables. This kind of information is not in any need of update at all, because it is total.

Simultaneously with the development of *DPL*, and of very similar interpretative architectures, insights from Robert Stalnaker and Irene Heim led to the development of formal systems of interpretation which also characterize 'update of information'. At first, however, these were not concerned with information about the values of variables, but about what concerns language users in the first place: about reality, or the world.

The basic idea taken from Stalnaker (1978) is that interlocutors in a discourse distinguish between relevant states of the world, which the interlocutors conceive possible, or desirable, and who seek to diminish this set of possibilities to (one of those of) the one the *actual* word could or should be. The idea is quite simple, indeed. If we look at situations of exchange of information, our partial state of information about the world can be modeled as a set of possibilities, which represent the world as being (like one) in that set; and information update consists in reducing that set, which effectively means reducing the set of alternative ways the world may be, and thus getting more information about what or how it actually is. As a matter of fact, this kind of approach to stating the interpretation of the sentences of natural language as 'update functions' on 'information states' has led to intriguing new perspectives on presupposition and epistemic modality. (Most prominently Heim 1991; Beaver 1995; Chierchia 1995; Groenendijk et al. 1996.)

Now it seems it would be an attractive thing, to combine the mentioned two notions of dynamic interpretation—at least it could look like a missed opportunity if one dynamic set-up would not, in a sense, generate descriptions, generalizations, and explanations by means of the same mechanisms. This turned out not to be an easy matter, though. Various attempts have been made (van Eijck and Cepparello 1994; Chierchia 1995; Groenendijk et al. 1996; Aloni 1998) all of whom develop idiosyncratic architectures, the differences between which hardly ever become substantial. One of the sources of this proliferation, and also of its solution, was the acknowledgment that the type of information related to the establishment of anaphoric relationships, is after all orthogonal to the type of information we are really interested in. The real world is, of course, real; but variables, and the values of particular variables, are virtual. Variables do not themselves show up in natural languages, and where they do in syntactic/semantic structures, they are interchangeable. (One should not talk about the real world in that, technical, way. One could do so, in a philosophical way.)

Coherent systems of interpretation that combine the two types of interpretation, or, rather, the two types of information, indeed distinguish the two carefully. As we have already seen, (Groenendijk et al. 1996) has it that: "Discourse information of this type looks more like a book-keeping device, than like real information." The other authors just mentioned also use distinct methods and tools to model update of information about variables, versus update of information about the world. In order

for the reader to get a glimpse of what is going on here, technically and critically, I give a tiny outline of the prominent system of Groenendijk et al. (1996) in the remainder of this section.

The task set out in Groenendijk et al. (1996) (from now on: *GSV*) is to develop a coherent system of *dynamic* interpretation that deals with anaphoric dependencies, epistemic modality, and, possibly, presupposition. Two types of information are involved, which are kept separate, conceptually, but which are technically intertwined nonetheless.

The first kind of information which the authors focus on is what they call discourse information, which is information about the (representation of) 'discourse items', a technical concept, relating to the series of subjects or objects that have been introduced in a discourse. Building on Vermeulen (1996), this type of information is encoded in so-called referent-systems. In the presentation of *GSV*, referent systems are partial functions (injections) from variables to natural numbers. The domains of these functions r typically consist of the variables x employed in (the transcription of) a discourse, and the numerical value $r(x)$ assigned to them. If the value $r(x)$ is i, then the variable is associated with the i-th discourse item in a currently evolving discourse. We have already touched upon the redundancy of this kind of information in Sect. 2.4, and even on it figuring as a source of logical and computational complications.

The second kind of information extends these referent systems with information about the world, and about the discourse items figuring in these worlds. This type of information can be summed up as follows. Once a discourse has introduced n discourse items, and attributed several properties to them, this kind of information is encoded as an 'n-ary relation in intension'. The relation consists of precisely those sequences $d_1, \ldots, d_n$ of n objects, paired with a possible world w, iff these sequences in w have the properties attributed to them in the discourse, and also stand in the asserted relations there.

It is not very useful to inspect the relevant definitions here, but the following illustration may gives a sense of the right idea. Consider the following middle-of-the-road utterance.

(36) She$_x$ insulted a$_y$ nurse.

Relative to the current state of the discourse, the referent system r assigns x something that has been introduced and which is thus assigned a discourse item, say the ith discourse item of the current discourse. The indefinite "a$_y$ nurse" introduces a new discourse item—an insulted nurse—which comes under the label y. If, and only if, the current stage of discourse is inhabited by n discourse items, this nurse is the $(n + 1)$th item, so $r(y) = (n + 1)$. At this point in the discourse $r(y)$ is a nurse who is said to be insulted by $r(x)$. This kind of information is modeled by means of the sequences of individuals $d_1, \ldots, d_{n+1}$, that, paired with a world w, have the properties and stand in the relations asserted of the n discourse items in the discourse, such that, in w, the $r(x)$th individual in the row (i.e. $d_{r(x)} = d_i$) stands in the insult-relation relation with the $r(y)$th individual in the row (i.e. $d_{r(y)} = d_{n+1}$). In this sketchily presented example we find two types of information, and two types

of updates. There is information about the 'discourse items', and there is information about their properties and the world they live in. The first can be updated by *adding*, or *introducing*, discourse items; the second by ascribing them, or the world they live in, certain properties.

Even though in the presentation of *GSV* thus far things look cumbersome, nevertheless, these details are well-motivated, intuitively. That is, as long as one buys the idea that updates in discourse involve updates of information about the values of specific variables. To be sure: nothing can be said against the invaluable use of variables in any theoretical or rational enterprise, be it physics, economics, humanities, linguistics, or even the question whether to go shopping or not. Here, however, the variables themselves, or their names, are getting subject of discussion, and this should not be.

In the first place, what index or name we assign to a parameter, or linguistic- or discourse- constituent, should not matter as long as we do it systematically. Variables are, and ought to be, replaceable by other variables. As argued in the previous chapter, however, they are not interchangeable in *DPL*; neither are they in *GSV*. In the second place, conflating variables or pronouns in actual discourse with the variables quantified or abstracted over in theoretical linguistics leads to troubling, and uninvited, questions about whether quantifiers do or do not bind them. Part of the reaons that made it take so long to develop comprehensive dynamic systems of interpretation had to do precisely with this issue.

Interpretation in *GSV* is presented in the following format. For any formula in the language—a language of first order predicate logic with an epistemic modality operator, and possibly an operator indicating presuppositional material—an interpretation is given that 'updates' 'information states'. For any such formula ϕ, and for any given 'input state' s, the update of s with ϕ is a new information state, written $(s)[\![\phi]\!]$; the result of applying the update function $[\![\phi]\!]$ to the information state s. The update function normally produces an *extension* of information, in the two ways. Updates may consist in a state s getting more informed about the discourse items already kept track of in s, or about the world in general; and it may consist in the introduction of new discourse items in state s. Of course, in general the two kinds, or aspects, of updating information proceed in tandem.

One of the most important, and most debated, clauses in the definition of the interpretation function is the one dealing with existentially quantified formulas, i.e., the definition of $(s)[\![\exists x\phi]\!]$. An interpretation that assumes existential quantifiers correspond to random assignments would interpret $\exists x\phi$ in s, first, as an extension $s[x]$ of s with a random assignment to x, and then an update of the state $s[x]$—state s with a new discourse item labeled x, about which there is no information—with ϕ. As a matter of fact, this is not the interpretation that *GSV* advocates, but it is worth spelling it out in some detail. The definition could be given as follows.

- $(s)[\![\exists x\phi]\!] = (\bigcup_{d\in D} s[x/d])[\![\phi]\!]$.

In this definition, $s[x/d]$ is the set of extensions of possibilities i in s, with a new discourse item x, evaluated as d. Given the referent system of s, this is one with a new, most recent discourse item, and the state $\bigcup_{d\in D} s[x/d]$ conceives of all theoretically

possibly interpretations $d \in D$ of that discourse item, relative to any conceived possibility in s. Thus, $\bigcup_{d \in D} s[x/d]$ contains no information whatsoever about x, and relative to this state of information of ignorance about x, ϕ provides a first update.

Notice that, upon the just now suggested definition of $\exists x\phi$, the existential quantifier can be conceived of as an atomic action, which could stand alone in a discourse. Upon this interpretation $\exists x\phi$ really says: "Take any x. Now ϕ." The definition that *GSV* actually supports does not allow for such a paraphrase. Formally the difference is subtle because only the round brackets are organized in a different way.

- $(s)[\![\exists x\phi]\!] = \bigcup_{d \in D}(s[x/d])[\![\phi]\!]$.

This time, the update provide by the embedded formula ϕ is carried out distributively over extensions $s[x/d]$ of the information state s with the information that x is a new discourse item, with the value d, for any d in the domain D. After all these updates have been carried out, the results are summed up, by means of the big union $\bigcup_{d \in D}$. The subtle difference with the previous definition is that, in the former, ϕ is interpreted, locally, with no information whatsoever about the value of x; while in the latter, ϕ is interpreted every time with *total* information about the value of x, that it is d—and this for any $d \in D$. Now if ϕ is an ordinary extensional formula, the two definitions will not make any difference. However, if ϕ contains epistemic operators or presuppositional devices, the updates may be different if these devices have a free occurrence of x in their scope.

Let us, like *GSV*, and for the time being, interpret $\Diamond$ as an epistemic operator, to the effect that $\Diamond\phi$ says that, according to the current information, it is possible that ϕ—that ϕ has not been excluded—formally, that an update $(s)[\![\phi]\!]$ with ϕ in the current state s does not lead to contradictions. In that case $\exists x \Diamond Hx$ ("Someone might be hiding.") may yield different results according to the two definitions. The first, the one not used by *GSV*, says that you take may a random interpretation of x, and then test whether you think it is possible that x be hiding. All in all, this says that you don't exclude that anybody is hiding—you conceive it is possible that anybody is hiding. According to the second definition, you test, as a value for the new discourse item labeled by x, whether any specific value d of x might be an individual which, for as far as you know, might be hiding. All in all, this gives you information about the identity of x: it can only be an individual you have not excluded to be hiding.

One of the examples *GSV* discusses, and where this difference matters, concerns the following two statements, with associated translations.

(37) Someone is hiding in the closet. He might have broken the vase.
$((\exists x HCx \wedge \Diamond Dx).)$

(38) Someone who is hiding in the closet might have broken the vase.
$(\exists x(HCx \wedge \Diamond Dx).)$

In case you have no idea who is hiding in the closet, but do know that someone broke the vase, and that *he* might be hiding in the closet, then (37) is fine. However, if there are also some persons whom you know have not broken the vase, and if one

of them might be hiding in the closet, then (38) is not supported. In that case none of the ones who might have done it are in the closet. Or, so the argument goes.

If the above observation is correct, Egli's theorem has to be qualified again, after all. According to *GSV*, at best the following holds.

Observation 14 (Modified Egli's Theorem)

- In *GSV*, $(\exists x\phi \wedge \psi) \Leftrightarrow \exists x(\phi \wedge \psi)$, provided that x does not occur free in the scope of a modal operator $\Diamond$ in ψ.

Related observations can be made about the interaction between quantifiers and presuppositions. One way of formulating presuppositions is by using a sentential presupposition operator ∂, to the effect that $\partial\phi$ says that, in the current state of information s, ϕ is presupposed to hold (Beaver 1995). Quite a long time ago this conception was argued to lead to what has become known as Heim's problem. Consider the following example, with schematic translation.

(39) A fat man was pushing his bicycle.
$(\exists x(FMx \wedge \exists y(\partial BOxy \wedge Pxy))).)$

The first definition of $\exists x$ [the one not chosen by Groenendijk et al. (1996)] induces a random assignment. The effect of interpreting this formula will be as follows. First it says, take any random value of x, so go to a state of information, containing the previous information, and add a new discourse item, labeled by x, about which you have no information. Then add the information that x is a fat man. Next, perform another random assignment to a new discourse item labeled by y, and interpret $\partial BOxy$, i.e., that your current state already contains the information that a/any fat man x has this bicycle y. This presupposes that all fat men have bicycles, and even worse, that you know them, the men and the bicycles, which doesn't seem to be a reasonable presupposition of (39).

The interpretation of (39) gets much better when using the ordinary distributive notion of existential quantification that *GSV* adopts. Doing so, the interpretation involves local processing of the embedded clauses relative to specific possible values of x and y; once a specific value of x is known to be a man, and y is known to be a bike of his, we can test or update with the information that that fat man was pushing that bike. The only presupposition that remains after processing example (39) is that at least one man is known to have a bike, which is fairly acceptable.

The *GSV* system does get around a couple of annoying problems raised in the literature, but uses ad hoc devices to circumvent the main problem, which consists in conflating variables with pronouns. The heart of the problem, which is not answered anywhere in this type of dynamic semantics, is that this type is bound to fail α-conversion as a valid logical principle. If, in a logic with quantifiers, or variable binding operators, it *does* semantically matter which are the variables (names) actually bound, this will hamper all kind of logical reductions. The laws of *lambda*-conversion, the Church-Rosser property of the λ-calculus, but also completeness proofs of first order predicate logic, depend on α-conversion. It is pretty undigestible to throw this away only for the purpose of treating natural language pronouns.

Another unsolicited effect of conflating pronouns with variables in the *GSV* system is the update with formulas containing real, quantified, variables, different from the update with those containing pronouns, which are free variables in their system. Consider a formula $\Diamond BVx$. If, on the one hand, the variable x in this formula is eventually bound, the evaluation of the formula is rigid. It tests whether the current interpretation of x, a concrete individual in the domain, might possibly have the property BV. It tests particular individuals whether they, according to the current information, might be blamed for having broken the vase. If, on the other hand, the variable is 'free'—a pronoun, so to speak—the update is different. Because it is no longer about individuals, but about discourse items the formula tests whether whatever information is gathered about the discourse item labeled by x is consistent with that item, whoever he is, having broken the vase. An significant consequence of this is that the law of universal instantiation no longer holds: in the *GSV* system, $\forall x \Diamond BVx$ does not even entail $\Diamond BVx$.

In addition to the problems with the interpretation of, or update with, modal formulas of the *GSV*-language, the informational support people bring to bear upon their utterance can be questionable. Here is (a variant of) an example brought forward by Maria Aloni in 1994, with associated translation.

(40) Someone might have broken the vase. She didn't do it.
$((\exists x \Diamond BVx \wedge \neg BVx).)$

If 'might' in the first sentence expresses an alethic modality nothing is wrong with the sequence of sentences. But the idea is to interpret it as an epistemic modality, so that the first sentence says: "There is someone whom, as far as my information is concerned, is not excluded to not having broken the vase." The next sentence asserts that person didn't break the vase. This sequence of two sentences, under this interpretation, is pretty weird. A hearer might indeed have no problem in, first, acknowledging, that some person *might* be the culprit, and, next, accepting, that *that* person *did not* do it after all. The problem with this example in the *GSV* system is that one and the same speaker is attributed information that supports his utterance that someone might, for as far as he knows, done it, and also that he didn't do it. While this indeed sounds, intuitively, very inconsistent, it is not only consistent in the *GSV*-system, it can even be coherently supported in the very same system.

The *GSV* system is intriguing because it combines a dynamic treatment of anaphora and modality in a single system, with some good results. It can be disputed, however, that it is in any sense explanatory, or even satisfactory. The dynamics of establishing anaphoric relationships, the treatment of epistemic modalities and the conflation of variables with pronouns have not improved our understanding of the empirical issues involved, but rather raise problems of their own. The moral of the current findings is that anaphora, quantification, modality and presupposition are, or can be, or have to be, all related, but they are, after all, distinct subjects. We want to be able to study them in their own right, without any non-classical bias towards what a solid formalism is needed for stating generalized interpretations. As argued in Chap. 2, a classical and solid basis for the treatment of anaphoric relationships may

consist in a satisfaction semantics in the style of Tarski. We will pursue this idea in the remainder of this chapter.

In the remainder of this chapter I show how a well-motivated system of interpretation can be fine-tuned to accommodate talk about 'contents' of utterances, 'updates' of information and 'support' for assertions, without having to give up classical intuitions about meaning. As a matter of fact, I argue that this modest point of view allows for a perspicuous treatment of anaphoric dependencies which have been—mistakenly—taken to motivate a dynamic notion of meaning. In Chap. 4 I also show how this perspective lends itself to systematic extensions of its scope to quantification and modality, without the kind of ad hoc moves that rival systems appear to have propagated.

3.2 Update and Support

The primary aim of this section is to present an update semantics which makes *PLA* more closely resemble the familiar update systems from Heim (1982), Veltman (1996), Groenendijk et al. (1996), and others. Before this is done an intensional notion of the contents of *PLA*-formulas will be specified. This is achieved by lifting the extensional Tarskian satisfaction system to a Kripkean possible worlds model, and thus get our fingers at what one may label a formula's propositional or cognitive content.

The propositional contents of formulas are formalized as sets of possible worlds, those that are consistent with the information expressed by these formulas. In the style of *PLA* they are defined relative to what one may call the pragmatic parameter of witnesses, which are now not just sequences of individual witnesses, but sequences of individual *concepts*, that is, sequences of intensional witness functions. The specification of the contents of *PLA*-formulas then simply consists in a relativization of all ingredients of the *PLA* system to possible worlds, and corresponding abstractions over these worlds. This is by and large a notational exercise.

In the intensional setting a model $M = \langle W, D, I \rangle$ consists of a set of possibilities, or worlds W, a domain of individuals D, and an interpretation function I for the constants of the language. For predicate and n-ary relational constants R^n we have that $I(R^n) \in (\mathcal{P}(D^n))^W$, i.e., a function from worlds to sets of n-tuples of individuals; for individual constants c it is required that $I(c) \in D_w$, by means of which I indicate the domain of *partial* functions from W to D. Individual constants, and variables and pronouns are interpreted as partial functions because they may fail to have a denotation in some worlds. Interpretation is defined relative to variable assignments g such that for any variable x: $g(x) \in D_w$, and sequences $\widehat{\epsilon}$ of witnesses $\epsilon_1 \ldots \epsilon_n \in D_w^n$. Everything is, thus, made parametric upon the set of worlds. (A formal side-remark. The domain of functions $(D^W)^n$ is isomorphic, of course, to the domain $(D^n)^W$, so that any witness sequence of total individual concepts can always be identified with a total function from worlds to witnessing sequences of individuals. It may sometimes be helpful to think of these sequences this way.)

The interpretation of terms (constants c, variables x and pronouns $\mathfrak{p}_i$) is always world dependent. They are given as follows.

- $[c]_{M,w,g,\widehat{\epsilon}} = I(c)(w)$, $[x]_{M,w,g,\widehat{\epsilon}} = g(x)(w)$, $[\mathfrak{p}_i]_{M,w,g,\widehat{\epsilon}} = \widehat{\epsilon}_i(w)$.

In what follows the elementary satisfaction clause is abbreviated as follows.

- $M, w, g, \widehat{\epsilon} \models Rt_1 \ldots t_n$ iff $\langle [t_1]_{M,w,g,\widehat{\epsilon}}, \ldots, [t_n]_{M,w,g,\widehat{\epsilon}} \rangle \in I(R)(w)$,

and similarly for identity statements $t_1 = t_2$, of course. The notation conventions for $\widehat{e}, \widehat{c}, \widehat{ce}, b\widehat{ce}, \widehat{a}$, and $\widehat{ace}$ from Sect.2.2 now apply to intensional objects $\widehat{\epsilon}, \widehat{\gamma}, \widehat{\gamma\epsilon}, \beta\widehat{\gamma\epsilon}, \widehat{\alpha}$, and $\widehat{\alpha\gamma\epsilon}$. With these notational conventions reformulating the *PLA*-satisfaction semantics in the intensional setting is fairly easy. All parameters are functionally dependent upon a world of evaluation. The contents of *PLA*-formulas, specified as sets of possible worlds, are recursively defined as follows. ($[\![\phi]\!]_{M,g,\widehat{\epsilon}}$ here specifies the contents of ϕ in a model M relative to assignment g and sequence of witnessfunctions $\widehat{\epsilon}$.)

Definition 1 *(PLA Contents)*

- $[\![Rt_1 \ldots t_n]\!]_{M,g,\widehat{\epsilon}} = \{w \in W \mid M, w, g, \widehat{\epsilon} \models Rt_1 \ldots t_n\}$;
  $[\![\neg\phi]\!]_{M,g,\widehat{\epsilon}} = W \setminus (\bigcup_{\widehat{\gamma} \in D_w^{n(\phi)}} [\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}})$;
  $[\![\exists x\phi]\!]_{M,g,\beta\widehat{\gamma\epsilon}} = [\![\phi]\!]_{M,g[x/\beta],\widehat{\gamma\epsilon}}$;
  $[\![(\phi \wedge \psi)]\!]_{M,g,\widehat{\alpha\gamma\epsilon}} = [\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}} \cap [\![\psi]\!]_{M,g,\widehat{\alpha\gamma\epsilon}}$

The content of an atomic formula is as usual given by the set of worlds in which it is satisfied. The ons significant addition is that the associated sequences of witness functions must provide the satisfying sequences of witness in each world. Relative to each particular world, satisfaction is identical to satisfaction in the extensional system. The content of a negated formula also captures that of the extensional system, but in a different format. According to the definition $\neg\phi$ gives us those worlds w for which no satisfying sequence of witnessing functions $\widehat{\gamma}$ for ϕ can be found, that is, relative to the given sequence $\widehat{\epsilon}$. The definition for the existential quantifier is virtually identical to that in the extensional system. The content of ϕ relative to an interpretation of x as individual concept β, is that of $\exists x\phi$ with the concept β as its witness. Where witnesses in the extensional set up figured as satisfiers, or truth makers, they are now conceived of as concept builders. The content of $\exists x\phi$ is the content of ϕ relative to a particular, contextually given, concept of x.

The definition for conjunction is the most remarkable one, since it stands in contrast with standard *dynamic* notions of conjunction: it is classic. Conjunction boils down to the intersection of information content. The operation is still dynamic, of course, like that of extensional *PLA*, but the dynamics does not reside in the contents, but in the contextual parameter and the fregean treatment of witnessing function sequences. A quick and easy comparison reveals that this treatment of the witness parameters is exactly like that in extensional *PLA*. The contents of the second conjunct ψ are given by its own sequence of witness functions $\widehat{\alpha}$, relative to the sequence of witness functions $\widehat{\gamma}$ for the first conjunct ϕ, when it itself is evaluated relative to $\widehat{\epsilon}$ in its turn.

One argument against systems of dynamic semantics is that they fail to provide an independent notion of information content (cf., e.g., Moltmann (2006, p. 228ff), who calls this the "propositional content problem." Brasoveanu (2010) makes a similar observation.). We now have a system here that deals with the dynamics of interpretation and whose independently motivated contents have been demonstrated, by being defined. It does not serve a great purpose to illustrate the assignment of contents in this system with worked out examples, because they are easily derived from the extensional formulation of *PLA* in the previous chapter, as shown in the following observation. Let M_w be the extensional model corresponding to w in M, g_w the assignment such that for all x: $g_w(x) = g(x)(w)$, and $\widehat{\epsilon}_w = \epsilon_1(w) \ldots \epsilon_n(w)$ if $\widehat{\epsilon} = \epsilon_1 \ldots \epsilon_n$. Thus, $M_w, g_w, \widehat{\gamma\epsilon}_w$ constitutes a sequence of parameters for extensional *PLA* satisfaction.

Observation 15 (Proper Contents)

- $[\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}} = \{w \in W \mid M_w, g_w, \widehat{\gamma\epsilon}_w \models \phi\}$.

Although $[\![\bullet]\!]$ (contents) and $\models$ (satisfaction) are each given an independent compositional definition, we now see they amount to exactly the same thing. The intensional content of ϕ consists of the extensional models for ϕ, relative to the right parameters, of course. One may observe that, like I announced before, nothing unusual is required in the intensional lift of the extensional system, neither does anything unexpected happen when doing so. The dynamics of interpretation, which has been seen at work in the Sect. 2.2 on the extensional satisfaction system, has been transposed to an intensional level, and we see again, and maybe more clearly, that the dynamics essentially resides in the dynamic conjunction of information contents, and not in the notion of contents themselves. Of course, this rather conservative moral does not prevent us from making a rather standard shift to a 'dynamic', or 'update' notion of meaning, a task which I now turn to.

Apart from the question whether lifting *PLA* is necessary there are some benefits of turning *PLA* into an update semantics in the style of Stalnaker (1978), Heim (1982) and Veltman (1996), where update semantics seems to have received its name. The basic idea is to give a recursive specification of what happens to an information state on a particular occasion if it is made to accept what has been said on that occasion. Interestingly, although the notion of an update is every now and then assumed to be the primitive, basic, notion, it always gets explained in terms of the update of a state *with the news or information conveyed by an utterance*. Conceptually, then, something more primitive underlies the update, and, conservatively speaking, that notion can very well be equated with our notion of information content, as we will see shortly. The information state may be or represent any hearer's state or a model of the common ground, a matter of debate which I will come back to later.

An update semantics for *PLA*-formulas is given by the specification of $(\tau)[\![\phi]\!]$, which is the information state that results from accepting ϕ in information state τ. It is specified here relative to the by now familiar parameters of a model, assignment, and a sequence of witness functions.

Definition 2 *(PLA Update)*

- $(\tau)[\![Rt_1 \ldots t_n]\!]_{M,g,\widehat{\epsilon}} = \{w \in \tau \mid M, w, g, \widehat{\epsilon} \models Rt_1 \ldots t_n\}$;
  $(\tau)[\![\neg\phi]\!]_{M,g,\widehat{\epsilon}} = \tau \setminus (\bigcup_{\widehat{\gamma} \in D_w^{n(\phi)}} (\tau)[\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}})$;
  $(\tau)[\![\exists x\phi]\!]_{M,g,\beta\widehat{\gamma\epsilon}} = (\tau)[\![\phi]\!]_{M,g[x/\beta],\widehat{\gamma\epsilon}}$;
  $(\tau)[\![(\phi \wedge \psi)]\!]_{M,g,\widehat{\alpha\gamma\epsilon}} = ((\tau)[\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}})[\![\psi]\!]_{M,g,\widehat{\alpha\gamma\epsilon}}$.

Information states are conceived of as set of possible worlds, like the contents of formulas, where sets of worlds characterize information, beliefs, and also desires in modal logical frameworks. A set of possible worlds models some agent's information if they are consistent with what that agent believes, or, put differently, if each of them might be the actual world according to what that agent knows and believes. All the things the agent believes must be true in the worlds in such a set, and if all these beliefs *are* true in a world, then that world belongs to the set. With this intuitive understanding, the first two clauses of the update definition are fairly obvious. If an atomic formula is asserted, and accepted, then one rules out those worlds, initially considered possible, in which that formula is not satisfied. Only those worlds that do satisfy the atomic formula are preserved. This is standard in an update semantics, but here, as well, we have to update relative to the right contextual parameters: sequences of witness functions. An update with a negated formula $\neg\phi$ involves ruling out those worlds which would have been preserved if a satisfying sequence of witness functions had been found for ϕ. If the update under no such sequence turns out to preserve a world w from the initial state τ, then w is preserved in the update of that state with $\neg\phi$.

The clause for the existential quantifier is basically copied from the content format to the update format. Relative to a witness function β, $\exists x\phi$ amounts to the update with ϕ with x rendered as β. Update with such a formula does not need a satisfying or instantiating witness, but a satisfying or supporting witness function, and the update then proceeds smoothly. Of course, in actual situations a hearer will often fail to know what *is* the witness function at issue, or what kind of individual concept an indefinite description is associated with. As we will see in the next section, the speaker can or should be taken to know this, but she is not required to provide the hearer with a clue to the witness. The hearer thus may be ignorant of the witness function and have to take into consideration a whole array of possible witnessing concepts, and consequently allow for a whole array of corresponding updates, one for each of these concepts. A more realistic notion of an update is therefore not deterministic, or it ought not to be conceived of as a function, but rather as a relation. [As a matter of fact, this moral has, for similar reasons, already been drawn in Groenendijk et al. (1996), Aloni (2000), van den Berg (1996), and Brasoveanu (2008).]

For a system dealing with the dynamics of interpretation, the notion of conjunction is the most obvious one. Updating with a conjunction $(\phi \wedge \psi)$ amounts to updating with ϕ first, and ψ next. The conjunction of two update functions, thus, consists in their composition.

We could again inspect a couple of examples to see this update system at work, but it would not really serve an independent purpose. For, as we already have seen in

the slogan motivating an update semantics, an update of an information state τ with a formula ϕ consists in the update of τ with the information conveyed by ϕ. This is made precise in the following observation.

Observation 16 (Proper Update)

- $(\tau)[\![\phi]\!]_{M,g,\widehat{\gamma_\epsilon}} = (\tau \cap [\![\phi]\!]_{M,g,\widehat{\gamma_\epsilon}})$.

This observation shows that we can define updates in terms of contents, but it follows that we can also define contents in terms of updates, for the observation implies $(W)[\![\phi]\!]_{M,g,\widehat{\gamma_\epsilon}} = (W \cap [\![\phi]\!]_{M,g,\widehat{\gamma_\epsilon}})$, where the latter of course equals $[\![\phi]\!]_{M,g,\widehat{\gamma_\epsilon}}$.

Two interesting points arise from the last observation, which formally maybe trivial, but conceptually are significant. In the first place information update is persistent. Whatever information is there in an information state, an update of that state always contains at least the same information, or more. For any state τ and formula ϕ: $(\tau)[\![\phi]\!]_{M,g,\widehat{\gamma_\epsilon}} \subseteq \tau$. In the second place, while updates and contents can be (that is to say: are) independently specified in a compositional manner, the present observation shows that they bring us the same result. Whatever reason one may have for formatting one's semantics in a standard static way (by defining contents), or in a standard dynamic way (by defining updates), the above observation shows that for the purposes at hand the choice is immaterial. Robert Stalnaker's analysis of assertion has already been based on this idea (Stalnaker 1978; Stalnaker 1998), and Max Cresswell has made it an explicit issue (Cresswell 2002). For the purpose of this chapter, the question whether conjunction should be thought of as intersection or as composition is no more than a frame question.

Support for Statements

A hearer's update of his information state with the contents of utterances made does not make much sense if there was any reason to suspect the information were false. Since nature shows itself, but doesn't talk, we have to rely on the agents that make utterances, and if we have sufficient reason to trust them, we have good reason to believe what they say. This observation motivates a shift of attention to the speaker's role in the exchange of information. Of course we cannot blame a speaker for saying something false if she is deluded, and sincerely believes what she says. A minimal demand is that she only provides information which she has. We may require her information state to support what she says. (In actual exchanges, and also in the exchange situations set out in Stalnaker (1978), the situation is a bit less demanding. People can be justified to only pretend, or act as if, their information supports what they say. The only reasonable requirement then would be that a speaker at least *could have had* a non-absurd information state to support all the things that she says, and then the baseline is one should not be inconsistent, as Groenendijk (2007) formulates the requirement. Even if this is the requirement, however, it does build on the underlying idea of an information state supporting what is said, which will be developed in the remainder of this section.)

In order to make this precise I use the simplifying assumption that information, the information of a speaker this time, can be represented by a set of non-excluded possibilities. In terms of such information states it is easy to define a support relation for the *PLA* language, $\sigma \models \phi$, which holds if σ contains the information which ϕ provides. (This, again, relative to a model M, variable assignment g and sequence of witness functions $\widehat{\epsilon}$, of course.) This relation is recursively defined again. The relevant parameters are those of a model, an assignment, and a sequence of witness functions.

Definition 3 *(PLA Support)*

- $\sigma \models_{M,g,\widehat{\epsilon}} Rt_1 \ldots t_n$ *iff for all* $w \in \sigma$*:* $M, w, g, \widehat{\epsilon} \models Rt_1 \ldots t_n$*;*
  $\sigma \models_{M,g,\widehat{\epsilon}} \neg\phi$ *iff there is no* $\emptyset \subset \rho \subseteq \sigma, \widehat{\gamma} \in D_w^{n(\phi)}$*:* $\rho \models_{M,g,\widehat{\gamma\epsilon}} \phi$*;*
  $\sigma \models_{M,g,\beta\widehat{\gamma\epsilon}} \exists x\phi$ *iff* $\sigma \models_{M,g[x/\beta],\widehat{\gamma\epsilon}} \phi$*;*
  $\sigma \models_{M,g,\widehat{\alpha\gamma\epsilon}} (\phi \wedge \psi)$ *iff* $\sigma \models_{M,g,\widehat{\gamma\epsilon}} \phi$ *and* $\sigma \models_{M,g,\widehat{\alpha\gamma\epsilon}} \psi$.

An atomic formula is supported by an information state iff it is supported throughout the state, relative to the sequence of witness functions. This requires that the formula is satisfied in each possibility in the state by the corresponding witness sequence. A negation $\neg\phi$ is supported by a state if no update or extension of the state supports the negated formula ϕ by means of any sequence of witness functions, except, that is, the absurd, or empty state. Support for $\exists x\phi$ relative to a witness function β consists in support for ϕ under a valuation of x as β. Support for a conjunction mimicks a conjunction's satisfaction. If the second conjunct ϕ can be supported with a sequence of witness functions $\widehat{\alpha}$ relative to a sequence $\widehat{\gamma\epsilon}$, of which the first part $\widehat{\gamma}$ helps to support the first conjunct relative to sequence $\widehat{\epsilon}$, then the whole $\widehat{\alpha\gamma\epsilon}$ serves to support $(\phi \wedge \psi)$ relative to $\widehat{\epsilon}$. The format of this clause is virtually the same as the one defining the satisfaction of a conjunction in the extensional set up of *PLA*.

Once again, it does not really serve a purpose to illustrate the above definition with some worked out examples, because the outcomes can be derived from the propositional contents or the update system, that is, ultimately, from the extensional satisfaction relation discussed above. In the following observation the various systems show to be appropriately related to each other.

Observation 17 (Proper Support)

- $\sigma \models_{M,g,\widehat{\gamma\epsilon}} \phi$ iff $\sigma = (\sigma)[\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}}$ iff $\sigma \subseteq [\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}}$.

This observation shows that an information state supports a formula *iff* it already has incorporated the update with it, that is, *iff* it subsumes its contents. As with the observation about proper updates, the observation about proper support raises a couple of interesting points, which, again, are formally trivial but conceptually significant. In the first place, the above observation implies that what we can do with an independently and compositionally defined notion of speaker support, we can do with propositional contents or updates as well. This, in the second place, means that complicated existentially quantified statements, the satisfaction conditions and contents of which can be laid out in a step by step fashion, can also be supported in

a step by step manner. This is a truly non-trivial point, to which I return in the next section.

The close connection between update and support also gives us the following equivalences.

- $\sigma \models_{M,g,\widehat{\epsilon}} \neg\phi$ iff for all $\widehat{\gamma} \in D_w^{n(\phi)}$:
  - $(\sigma)[\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}} = \emptyset$;
- $\sigma \models_{M,g,\widehat{\epsilon}} (\phi \rightarrow \psi)$ iff for all $\widehat{\gamma} \in D_w^{n(\phi)}$ there is $\widehat{\alpha} \in D_w^{n(\psi)}$:
  - $(\sigma)[\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}} \models_{M,g,\widehat{\alpha\gamma\epsilon}} \psi$.

These equivalences agree with intuition of course. An information state supports the negation of a formula iff any update with the formula is ruled out; and it supports a conditional sentence $(\phi \rightarrow \psi)$ if an update of it with ϕ yields a state that supports ψ. Indeed, the matter boils down to the so-called 'Ramsey-test' for conditional sentences.

A final observation here puts the above-developed notions of update and support 'en rapport' in situations of information exchange.

Observation 17 (Supported Updates)

- If $\sigma \models_{M,g,\widehat{\gamma\epsilon}} \phi$ then $(\sigma \cap \tau) \subseteq (\tau)[\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}}$.

Given the previous observations, this result is formally trivial. For if we take $\upsilon = [\![\phi]\!]_{M,g,\widehat{\gamma\epsilon}}$, then by Observations (16) and (17) Observation (18) claims that if $\sigma \subseteq \upsilon$ then $(\sigma \cap \tau) \subseteq (\tau \cap \upsilon)$, an instance of a distribution law for $\subseteq$. Conceptually it is extremely appealing though, because it may counts as a real felicity condition on ordinary discourse. The observation says that the update with supported utterances does not corrupt the information exchanged. If the information of speaker and hearer is correct, that is, if the actual world w_0 is excluded by neither of them, then it is also not excluded in any update they accept provided the update is supported. If I believe what you tell me, and you believe what you tell me, then we can end up wrong only if at least one of us was already wrong in the first place.

It may be noticed that this feature is typically threatened when it comes to discussions about epistemic modality. An epistemic possibility or necessity can be made correctly, and be supported by the speaker's information state, but taken in a different way by the hearer, so that almost automatically misunderstanding and conflict arise. Not so, as we will see, in the treatment of epistemic modalities in the final part of this monograph. What should be emphasized, here, is that the confusion from statements of epistemic modality does not arise from the differences in content a sentence may have on different occasions. Rather the confusion arises by the use of a sentence in a very specific circumstance which may lead to divergent assessments by speaker and hearer, respectively. Such a confusion is due to a misguided analysis, I claim.

The requirement that information may get corrupted through information change of course cannot be expelled throughout. For one thing, Stalnaker (1978) states very explicitly that assertions are acts which change the world. Therefore, something which was true before an assertion may be false afterwards. (For instance, the fact

that no assertion has been made.) As soon as we allow assertions to be about assertions themselves, or about their consequences, then the very assertion of one true statement may make itself false. Here is a natural example.

(41) I haven't told you before, but your wife is cheating on you.

An assertion made with this sentence may be true and supported, but once it is made, it is rendered false. After you have told me that my wife is cheating on me, it is no longer true that you haven't told it. However, as long as we talk about independent facts of the world, and not about the peculiarities of our own discourse and information, the supported update fact stated in Observation (18) is preserved by all means.

The exposition about update and support could stop here, but this would leave neglected an issue which will turn out to be more and more relevant in the following sections. The three notions of content, update and support are defined relative to arbitrary sequences of witness functions, that is, relative to sequences of arbitrary individual concepts. The full domain of individual concepts is normally unreasonably large. This doesn't matter so much for the *PLA* notion of content. For the *PLA* notion of an update, such a large domain may already seem to be a bit too liberal to be reasonable, but that's maybe to be decided by the person who has to act out the update. For the notion of support, however, unrestricted access to all individual concepts is awkward.

In actual situations of information exchange, quite a few things are conceptually organized and coordinated, and among these is a common assumption about the domain under discussion. Not only do people involved in an exchange situation assume a common domain, but also a common conception of it. Interlocutors do not act as if the things they talk about happen to be the same things, as if by coincidence. This implies, among other things, that when we are concerned with the actual support one agent has for an utterance, or the update another one extracts from it, both the support and the update are assumed to be based on witness functions which naturally belong to a shared, suitable and accessible conception of the domain.

These deliberations imply most concretely that if we want to say that a speaker's information state supports the utterance of a certain formula, the sequence of witness functions is not arbitrary, but these witness functions should be taken to belong to a suitable set of individual concepts. An appealing characterization of such a set has been given by Maria Aloni in terms of what she aptly dubbed 'conceptual covers' (Aloni 2000). Conceptual covers are special sets of individual concepts. The concepts in the sets constitute a way of 'seeing' a domain in the sense that each individual in any relevant situation, is covered by exactly one concept. Formally they are defined as follows. In the following definition W is the set of worlds assumed in a model, and relative to a subset U of them we have a set of concepts C, and we define whether this set C constitutes a conceptual cover for, at least, a subset V of U, which can be taken to constitute a characterization of what the world is taken to be like. (The intermediary subset U is needed to allow concepts to be defined for possibilities not believed to be actual.)0

Definition 4 *(Conceptual Covers)*

- *A set of concepts* $C \subseteq D^U$ *is a conceptual cover for* $V \subseteq U \subseteq W$ $(V \neq \emptyset)$ *iff for all* $c, c' \in C$*, if* $c \neq c'$ *then for all* $v \in V$*:* $c(v) \neq c'(v)$*;*
- *a conceptual cover for* V *covers* $C_u = \{c(u) \mid c \in C\}$ *in* $u \in V$*.*

The above definition of a conceptual cover is more general than Aloni's because it allows covers of proper subsets of the domain and for proper subsets of worlds. (As a matter of fact, for practical purposes, the total covers from Aloni, which are defined for all possible worlds, and which cover all individuals, are quite unrealistic. We will come back to this issue below.)

A conceptual cover is a set of individual concepts, at least defined for a particular set of possible worlds V—an information state, or a situation if you want. They serve to uniquely identify a number of individuals, in each possible world in the set V. Different concepts are assumed to identify different individuals in every world in V, even though we could have different concepts that, counterfactually, do identify the same individual in a world $u \notin V$. (In that case, however, the two extended concepts could not both be concepts in one and the same cover for $V \cup \{u\}$.)

If we take any world v from V, then a cover C for V will determine a set of individuals, which are the 'real' individuals identified by C in v. These may be different sets of individuals in different worlds v and v', even though, even if $c(v)$ and $c(v')$ are different, they 'count as the same', from the perspective of the cover, that is. It is like the temperature in my house, which may change.

Before we discuss an example some formal properties of conceptual covers will be given which follow from the definition above.

Observation 19 (Restricting and Extending Conceptual Covers)

1. If C covers C_u in $u \in V$ then $|C| = |C_u|$;
2. a cover C for V, is a cover for $X \subseteq V$ $(X \neq \emptyset)$;
3. if C covers C_u in u then $\{c \in C \mid c(u) \in C'_u\}$ covers $C'_u \subseteq C_u$ in u;
4. a cover C' for $V' \supseteq V$ is an extension of a cover C for V iff for all $c \in C$ there is $c' \in C'$ such that for all $v \in V$: $c(v) = c'(v)$;
5. every cover of C_u for $u \in V$ has an extension which covers D for $u \in W$.

According to the first point in Observation (19), the number of concepts in a cover C for V equals the number of objects that it covers in any world $v \in V$. Relative to an information state or situation V, a cover, thus, is numerically well-behaved. The other points have to do with extensions and restrictions of the information states and the domains covered. If we restrict the set of worlds V, so if we gain or add more information, then we can keep the cover, for a more limited domain of worlds. Due to the distinctness clause in the definition of a cover, the number of concepts will remain the same. (That is, as long as we don't take U to be the empty set, which counts as an absurd state of information.) What does change, under such an 'update' is not the number of concepts, but the number of alternative individual values that each concept may take. That is to say, if we reach a state $\{v\} \subseteq V$, then the restriction to $\{v\}$ of a cover C for V will be a cover that totally defines C_v. Actually, in such

an ideal state of total information no further question exists about the identity of the individuals covered.

The third item shows that one can take any restriction of a domain of individuals covered by a cover C, and get a cover $C' \subseteq C$ that covers exactly that domain. The fourth item is actually a definition of a cover-extension. An extension of a cover C for V is a cover C' for $V' \supseteq V$ consisting of concepts which are defined for a possibly larger domain of worlds. It may also cover a larger domain of individuals in every world $v \in V$, but, of course, only as long as C' still is a cover satisfying distinctness. The final line states that every partial cover can be extended to a total cover, for all possible worlds, and such that it covers the whole domain of individuals D in each world. Thus, Aloni's own total covers are, directly, or indirectly still in play.

While it is true that one conceptual cover for V may cover different sets of individuals in different worlds in V—so that for some $u \neq v$: $C_u \neq C_v$–we can as well have covers for V that cover exactly the same domain $C_u = C_v \subseteq D$, for any u and v in V. In that case one might say that such a cover is one cover of, or perspectives on, a specific subset of the domain of individuals.

Let us now turn to some covers in action. A typical covering situation may arise in a soccer stadium, e.g., Nou Camp in Barcelona. Suppose you are there, you know Barça plays in red/blue shirts and this team consists of 11 players. You see indeed 11 players in the right colours, but they are too far away for you to even distinguish them. Also, you are given a list of the names of the players: Valdés, Alves, Piqué, Mascherato, Abidal, Sergio B., Xavi, Iniesta, Pedro, Villa, Messi. You are also given a list of the squad numbers, a subset of the numbers ranging from 1 to 22. In addition to this you are carrying a photograph of the team playing (winning) today, August 14, 2011. Focusing on the single domain of players seen at the field, you can conceive of them from various perspectives, each corresponding with a way of 'seeing' that domain. The list of eleven names corresponds one-to-one with the eleven figures seen on the field, even though you may fail the right or complete mapping—or even be mistaken about it. What you do suppose, though, is that any two names name a different person on the field, and every two persons in the field have a different name. Believing what you are told, the same holds between the 11 squad numbers, and the names of the players. (Just FYI, the association runs as follows: 1 Valdés, 2 Alves, 3 Piqué, 14 Mascherato, 22 Abidalm, 16 Sergio B., 6 Xavi, 8 Iniesta, 17 Pedro, 7 Villa, 10 Messi. Of course, one may fail to know part of this information.) Consequently, the squad numbers are another way of uniquely determining the domain—even though you still need not be able to tell who is who. You also believe the persons depicted in the photograph are mapped one-to-one to the players on the field, the list of names and the list of numbers.

Now all this information is captured by an information state σ that only consists of worlds that satisfy the stated assumptions. Then we can formally identify (at least) four relevant conceptual covers for σ, all of which cover the eleven players in the field.

1. A direct perception cover C_1, let's call it a 'Kaplan cover'. This cover consists of eleven constant functions. For any individual d among the eleven players in the field, there is a concept assigning d to any world in σ.
2. A naming cover C_2, let's call it a 'Kripke cover'. This cover consist of eleven functions which individuate the domain by their names. So there is a concept assigning to each world $w \in \sigma$ the person named 'Valdés in w, and similarly for the other ten names. Notice, that given the stated assumptions, the individuals assigned to the names are always one of the persons on the field.
3. The squad number cover C_3, let's call it an 'Evans cover'. This is the cover with functions which, for any squad number, e.g., 17, assigns every world in σ the person who is carrying number 17 on his back. (Which you might be unable to see.)
4. The picture cover C_4, let's call it a 'Lewis cover'. This cover identifies a person by his resemblance with a person on the picture. Since the faces on the picture are dissimilar enough, and this information is present in σ, in every world there is exactly one among the players in the field who resembles, for instance, the person in the middle of the front row.

The main point about, and motivation for, these conceptual covers, is that one may know, for instance, who scored, if one sees a person on the field making a goal, and one employs the Kaplan cover; at the same time, you may as well wonder who actually scored, not knowing this from the perspective of the Kripke cover. At the same time one may know again who did, from an Evans cover, because one could see his squad number. Yet, one may again not know, from the Lewis cover, which of persons on the picture scored—whether it was this most handsome guy, or the guy with the funny hair. All kinds of possibilities of knowing who and not knowing who are possible. In Sect. 4.2 we will come back to these issues in more detail.

A note-worthy point that anticipates some issues to be discussed in Sect. 4.2 is, in a slogan, that conceptual covers cover each other. [(This point is also, formally, addressed in Aloni (2005).] First, observe, that conceptual covers are initially assigned the status of a perspective on a given domain of individuals, *without* any relevant properties of themselves. They are all 'Dinge an Sich', just elements of any set. These objects, however can be equated with the concepts given in a direct perception cover, since for any set of individuals $E \subseteq D$, we can establish an isomorphism with the set of *constant* functions from any non-empty set of worlds V to E. It is then, conceptually, a simple step to conceive of any cover (C) of the same domain, as a cover of the set of corresponding set of constant concepts, i.e., as a cover of the direct perception cover (DPC). For, if two concepts in C are different, then, for any world v in the set of worlds covered will have a different value, and are also associated with different constant concepts from DPC in v. And this holds in general. For any two covers C_1 and C_2 for sets of worlds V, which both cover $E \subseteq D$, we find that for any $v \in V$, there is a one-to-one mapping from C_1 to C_2, relating $c_1 \in C_1$ with $c_2 \in C_2$ iff $c_1(v) = c_2(v)$. Assuming, for the sake of the argument, cover C_1 to be 'basic', we may assume C_2 to be a cover of a domain consisting of the concepts in C_1, and so that in any world v, the value of c_2 is, not the individual C_2, but the concept c_1 from C_1

that has the same value as c_2 in v. Of course we can also conceive of C_1 as providing a cover of C_2. In this sense, any pair of covers C_1 and C_2 for the same set of worlds, and with the same cover, can also be seen to cover each other. Philosophically this may be a delicate point, although, eventually, it boils down to the observation that no logic or semantics can distinguish between two models which are isomorphic to it.

Returning now to the issue of *PLA*-support, we could implement the idea that the individual concepts relevant for the support of linguistic utterances all come from one and the same contextually given conceptual cover of the domain. This, however, would prevent a flexible use of them which appears to be required to model the type of ignorance expressed by *But I don't know who is who.* (I will come back to this type of locution in the Chap. 4.) The solution is to allow each quantifier to draw its concepts from its own cover. Formally this requires us to label quantifiers with an index specifying which covers they draw from, and for the intensional system of *PLA* this solution requires us to add the qualification that the witness for an existential quantifier be in the cover associated with the quantifier. The relevant clauses in the definition now read as follows.

Definition 5 *(PLA Satisfaction and Support)*

- $M, w, g, \beta\widehat{\gamma\epsilon} \models \exists x_C \phi$ *iff* $\beta \in C$ *and* $M, w, g[x/\beta], \widehat{\gamma\epsilon} \models \phi$*;*
- $\sigma \models_{M,g,\beta\widehat{\gamma\epsilon}} \exists x_C \phi$ *iff* C *is a cover for* σ*,* $\beta \in C$*, and* $\sigma \models_{M,g[x/\beta],\widehat{\gamma\epsilon}} \phi$*.*

The only relevant change here is that the witnesses for an existentially quantified formula, are witness functions from a given conceptual cover C. This may invoke contextually restricted quantification, depending on which set of individuals is covered by C. As indicated above, if needed we can always assume this to be the whole domain D. Notice that I also have not specified which set of worlds the cover C is supposed to be a cover for. In the definition of the support relation of an existentially quantified formula by a state σ it is supposed to be a cover for at least σ, but this is not always sufficient, for instance, if x eventually has to be evaluated in a counterfactual world. However, as we have seen above, we can always assume such a cover to be 'large enough', to make sure that the individual concepts involved will never be undefined for any world they have to apply to, including counterfactual ones.

The definition of quantification relative to conceptual covers is still very liberal, yet it enables us to state further requirements on a conceptual cover whenever it shows up. In general, for instance, we can assume the employed conceptual covers to be 'natural covers'. Because, most probably, there *are* no really natural concepts, it is inappropriate to require specific individual concepts to be natural. However, to require particular concepts to be taken from a natural cover does make sense. (This is like asking whether a cover fits in a natural conception of the domain, which might be given independent motivation. Compare this, for instance, with compass points. Southwest is a fine concept if is assumed to belong to a natural 'cover' which includes N, E, S and W, as well as NE, SE, and NW; an unnatural cover, however, can be argued to be one which hosts north, north northeast, northeast, east northeast, east and southwest.) In the sequel I will silently assume such a 'naturalness' assumption,

without making it very specific, though. The benefits of doing so, and, hence, of the present adaptation, will become clear in the subsequent sections.

3.3 Information Exchange

Even though the preceding sections were not entirely trivial, one may wonder why we went through this exercise of defining alternative but basically equivalent semantic notions like that of content, update and support for *PLA*-formulas. In the first place, this exercise shows, by demonstration I claim, that it is immaterial whether the treatment of singular anaphoric pronouns is or should be stated in the fashion of a satisfaction semantics, an assignment of propositional contents, an update algorithm, or a support calculus. Second, it shows a real payoff when we start talking about contents in multi-agent situations, typically the kind of situations in which information is exchanged by linguistic means. In this section I want to apply the concepts developed above to a couple of phenomena that seem to resist a treatment in purely semantic terms.

While, intuitively, the role of sequences of individuals in the extensional setting consists in providing satisfying witnesses for truth, the role of sequences of individual concepts in the intensional support setting is to embody the referential intentions instrumental in a formula's support. When terms, among which I include indefinite noun phrases, are used in discourse, it is assumed they are used with referential intentions. This claim is actually quite old, and it has relatively recently been picked up in van Rooy (1997), Kamp (1990), Stalnaker (1998), and Zimmermann (1999). Let me substantiate it a bit further first, before we see how it fits in the picture developed in this section.

A first and compelling example for the assumption of referential intentions dates from Peter Strawson 1952. Strawson mentions two possible responses, (43) and (44), to statement (42) (Strawson 1952, p. 187, in the 1960 edition).

(42) A man fell over the edge.
(43) He didn't fall; he jumped.
(44) It wasn't a man; it was a woman wearing trousers.

These examples pose a challenge to standard syntactic and semantic approaches, because a speaker here starts talking about an individual with an indefinite noun phrase, which the respondent refers back to by means of a pronoun. What information is exchanged here, and how remains a mystery if the relation between the indefinite noun phrase and the anaphoric pronoun is one of syntactic or semantic binding. However, on the simple assumption that the first speaker intends his indefinite noun phrase to speak of a certain man (or woman) in the vicinity of both interlocutors, what the little exchange is about, and what the exchanged contents are, is clear. The actually present man (or woman) can be used to resolve the pronoun. Notice that this analysis perfectly fits in Cresswell (2002)'s proposal, at least in concept. The opening statement of this little discourse would be read as "A man, namely x, fell

over the edge." and the response could be "He, namely x, did not fall, etc." On this interpretation x gets a value from the non-linguistic context of the discussion, and two different, together inconsistent, attributions are made about this value. Notice, though, that this analysis does not easily fit Cresswell's formal rendering of *DPL*, as it does not fit in the system of *DPL*. For, the analysis of the reply "He, namely y, didn't fall, etc." then cannot be achieved without binding y to a man (or woman) who is said to have fell.

One might think that the actual presence of the man (or woman) in Strawson's example may be used to explain the case as one of direct reference. However, indefinites can also be used to report on individuals which are not actually present. Imagine the following situation. One person, Liz, was visited by two students yesterday, who both asked for the secretary's office. One of them (Wilburt) did, and the other one (Norbert) did not wear pink pumps. Liz is fully aware of all this, except that she never knows the students' names. Zil knows that Wilburt always bothers people with nonsense questions, and that he wears pink pumps. Now Liz starts the following exchange with Zil:

(45) *Liz:* Yesterday, a student ran into my office who inquired after the secretary's office.
Zil: Was he wearing pink pumps?
Liz: He was indeed.

So far the conversation seems perfectly fine, if Liz had indeed started out talking about Wilburt. But notice that Zil can ask questions about the student Liz talked about, without (yet) having an idea of who it was. Yet, there is a proper answer to the question. Notice, too, that if Liz had simply meant to assert that the set of students that visited her yesterday was not empty, then Zil's question would have been odd indeed, or impossible for Liz to answer. For suppose that her reply to Zil's question was the following. (This version of the scenario is due to Ede Zimmermann, p.c.)

(46) *Liz*: I don't know. What do you mean? If it was the first one who came in, he was wearing pink pumps; if it was the other, he was not.

Liz could motivate her 'ignorance', by claiming that, when she earlier said that a student ran in, she had not yet decided which of the two she was talking about. It seems obvious that this would be a very odd reply. Liz is expected to have had someone in mind when she started to talk about a student. So, if one starts talking about an individual, who is introduced by means of an indefinite description, there still is a fact of the matter of who it is, about whom we can sensibly ask and answer questions, like whom it was about.

Such questions can also be genuine identity questions. Imagine the following situation. Weird rumours have it that young girls get battered in Gotham city, and that city representatives like, e.g., Bernard J. Ortcutt are involved. Bor and Cor exchange the latest findings in the tabloids.

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

confessed?
Bor: I don't know.
Cor: Do you know who confessed?
Bor: No idea, the police didn't disclose his identity.

There is nothing incoherent about a dialogue like this. It is an exchange of information between two inhabitants of Gotham city, who—perhaps naively, but not incoherently—take their tabloids seriously. The example shows that the two interlocutors can exchange information about an individual without knowing exactly whom they are talking about. Yet, the two interlocutors apparently assume the discourse to be about a certain individual, otherwise the question *who he is* wouldn't make sense in the first place. The explanation of this issue, like any issue of identity, is that a person or object may be given in a certain way, for instance as 'the person whom the tabloids are reporting on'. While one fails other means to identify that person, there still is a referential intention involved, viz., the intention to relate to precisely that person.

The upshot of these examples is that regular indefinite noun phrases are used with referential intentions, even when legitimate questions about the identity of the actual referents remain, but that their use is justified conceptually. The concept may be intensional, in the sense that it invokes the concept of an individual someone else intended to refer to, and its use may engage in a causal intentional chain in the sense of Kripke (1972), Chastain (1975), Donnellan (1978), and Evans (1982). In any case, and that is relevant here, it is a concept which can always be rendered, formally, as a function from worlds to individuals.

The preceding observations then neatly fit the compositional notion of support developed above. Like I said, the sequences of witness functions embody the referential intentions with which the indefinite noun phrases or corresponding existentially quantified formulas may be used. So, follow up on example (42) from Strawson above, the speaker says that *a man fell over the edge*, we may assume that the speaker's state σ supports an indexical individual concept β of a person who goes down the edge in the speaker's presence. Notice that such a concept makes perfect sense in a natural conception of the situation, a conceptual cover indicated as S here.

- $\sigma \models_{M,g,\beta} \exists x_S(Mx \wedge Fx)$.

The hearer, in this situation, may be able to pick up that individual concept, and notice that the propositional content expressed, $[\![\exists x_S(Mx \wedge Fx)]\!]_{M,g,\beta}$ is false according to his information state τ.

- $(\tau)[\![\exists x_S(Mx \wedge Fx)]\!]_{M,g,\beta} = \emptyset$.

The hearer corrects the speaker by saying that *he didn't fall; he jumped*, as in (43), because that's the information he has.

- $\tau \models_{M,g,\beta} \neg F\mathfrak{p}_1$ and $\tau \models_{M,g,\beta} J\mathfrak{p}_1$.

The information of speaker and hearer, and the communicated contents, are clearly centered around the same individual $\beta(w_0)$ in the actual world w_0, if any such person

indeed went down the edge. Here we see that the notions of information content, information support and information update fit well together in the analysis of a situation in which one speaker employs a pronoun to refer to an individual introduced in the discourse by another. The other scenarios we have discussed above fit our notion of support in an analogous fashion. The point made here about supporting witness functions for indefinite noun phrases and pronouns, also applies to names (individual constants), of course. While it may be hard to sort out which specific witness function a speaker has in mind when she uses an indefinite noun phrase, for a proper name this process is often easy, for each individual constant comes with its own individual concept which is a very natural concept to use it for. This is not to say, though, that we always use a name this way; not when we say, for instance, that Fred might not have been called 'Fred'.

To conclude this section, I apply the concept of support to illuminate a puzzle raised by Charles Saunders Peirce. (Also the following analysis neatly fits in the proposal made in Cresswell (2002); this time formally as well.) The puzzle, very concisely, is this. The following two formulas are equivalent in first order predicate logic.
(α) $(\exists x Wx \leftarrow \forall x TPx)$;
(β) $\exists x(Wx \leftarrow TPx)$.
In Read (1992), Stephen Read presented a situation to evaluate two natural sentences, the meanings of which appear to be captured by the above two formulas. The situation is that a sweepstakes takes place, a thousand people are in the position to participate, and each participant brings in one dollar. So if n people participate, one of them wins n dollars. Now consider the following two statements.

(a) Someone wins \$1,000 if everyone takes part.
(b) Someone wins \$1,000 if he takes part.

Statement (a) is obviously true in the situation as described. Statement (b), however, may raise people's eyebrows. This statement provides or suggests inside information that does not directly follow from the description of the situation. The puzzle is, how can that be, if (α) and (β) render the meanings of (a) and (b)? If the first two are equivalent, then the second two should be equivalent, too. But they don't seem to be. What is going wrong here? Our judgment of the sentences (a) and (b)? Their translation into (α) and (β)? Or perhaps the equivalence of the latter?

A *PLA* style analysis of the situation reveals what is going on. First observe that an utterance of (b), but not of (a), naturally raises the question: *Who?* For it suggests that there is someone with the special property that if he takes part, he wins. (The property is even more special, since his taking part would make all other take part as well, otherwise he doesn't win the full \$1,000.) If we know who has this special property, we have good enough reason to convince him to take part. Where does this suggestion come from? Consider what concept δ of the speaker could make her information state σ support her utterance of (b), rendered as (β) above.

- $\sigma \models_{M,g,\delta} \exists x(Wx \leftarrow TPx)$ iff
  $\sigma \models_{M,g[x/\delta]} (Wx \leftarrow TPx)$ iff
  $\forall w \in \sigma$: $\delta(w) \in I(W)(w)$ if $\delta(w) \in I(TP)(w)$.

We find that the speaker should indeed have a concept δ at her disposal of someone who, according to her information, wins \$1,000 if he takes part. What concept can that be? It can be the concept of a person of which the speaker thinks he is not going to take part anyway. It can also be the concept of a person of which the speaker thinks he is going to win \$1,000 anyway. Gricean reasoning disallows these explanations, because if they provide the speaker's evidence for his assertion, the assertion would violate Grice's quantity maxim. Another concept is one of a person of which the speaker doesn't know whether he will participate or not, but of whom she does believe that if he will participate, he will win this \$1,000. Here is our lucky bird. The speaker seems to know of a predestined winner, or she has inside information that the whole sweepstakes is set up.

The story does not happily end here, though. Read's sweepstakes situation supplies another concept that enables an information state to support formula (β) or example (b). This is a hybrid concept that corresponds to the concept of the winner in all possibilities in which everybody participates, and to the concept of an arbitrary drop out if at least one person doesn't take part. This witness function renders (β) supported by any information state that incorporates the description of the sweepstakes situation. However, even though the concept itself cannot be ruled out as unnatural, it doesn't seem to fit in a reasonable conception of the sweepstakes situation, that is, it does not belong to any natural conceptual cover of the situation. The amended version of the support relation thus rules this out, so the story does happily end after all.

The overal conclusion of this discussion is that (α) and (β) do render the meanings of examples (a) and (b), and that the two formulas do have the same truth conditions, in *PL* and, hence, in *PLA*. In *PLA* however, (β) and (b) have additional satisfaction conditions. They require a supporting witness together with the cover it is drawn from, and on the level of information support these combine with pragmatic information to yield the stronger reading which Pierce and Read have observed.

3.4 On the Contextualist Debate

With our use of witnesses, the notion of a dynamic conjunction, and our focus on issues of update, support, and information exchange, we have obviously, and deliberately, incorporated arguably pragmatic aspects of meaning in the interpretational architecture. 'Dubious business', some may claim, as for instance those who, like Gareth Evans maybe, live in Plato's, Frege's, Russell's or Carnap's heaven. Going this way, you spoil the autonomy of semantics, of logic, and put much of philosophy up for grabs. Others may be less discomforted, like those who have felt uneasy about the old-fashioned practice of throwing all unsolvable syntactic/semantic problems in the proverbial pragmatic wastebasket (Bar-Hillel 1971). One may rejoice that we continue cleaning up that bin. Yet others again may welcome this pragmatic move, and welcome me at their side, arguing this is just a first step into seeing that every-

thing is pragmatic after all, and once this is properly acknowledged, we all see there is no room left for semantics proper.

With the last reaction we enter the so-called 'Contextualist debate'. I say 'so-called', because the discussion does not consist much of arguments supporting conclusions. The debate is often presented as one between contextualists on the one hand, who claim that the interpretation of linguistic utterances is so inherently context-dependent that there is no sensible room left for any context-independent semantics proper; and on the other hand so-called minimalists, who fight for the right of something like a formal semantics, delivering minimal propositions as core meanings associated with linguistic constructions. In what follows I will attempt a constructive reply to the contextualists' arguments. I accept almost all of their arguments, without any inclination to agree with their destructivist conclusions.

The contextualist challenge consists in series of arguments showing to what a massive extent pragmatic and contextual phenomena determine the interpretation of utterances on all syntactic levels. Let us review some of these types of observations with regards to the crucial realms of reference, predication, quantification, and construction.

Reference There exists quite established consensus, since Donnellan (1966, 1978), that the following example can be used to convey information about a man who hasn't murdered Smith.

(48) Smith's murderer is insane.

If uttered in the right context, with people sharing the same, mistaken, assumptions, a person who is not Smith's murderer can be taken to be so, and thus serve as the referent of the predication that he is insane. Saul Kripke, not convinced that the phenomenon addressed by Donnellan is semantically relevant, agrees that an utterance can be used in such a way, and he adds two more relevant examples (Kripke 1979b):

(49) Jones is raking the leaves.
(50) *A:* Her husband is nice to her.
B: He is nice to her, but he is not her husband.

The man actually said to be raking the leaves can be mistakenly identified as Jones, but even so an utterance of (49) may successfully communicate that *he*, not Jones, is raking the leaves. Example (50) can be used in relation to a wandering couple, whom the speaker assumes to be married, while, apparently, the hearer knows better. Yet, what is communicated, and what the hearer agrees upon, clearly, is that the man who is not her husband is nice to her. What a semantic theory would say about the meaning of the uttered sentences in these examples is orthogonal to some extent to what is communicated. The intended reference is accomplished, but not, or not obviously, by semantic means. An even more striking case is one from Nunberg (1979):

(51) The ham sandwich wants to pay.

The definite noun phrase "the ham sandwich" is used to pick out the person who ordered a ham sandwich in a restaurant, while, of course, strictly speaking the sentence itself is not about a person, but expresses a somewhat unexpected inclination

of a ham sandwich. If language is a means for communication, then, in the above examples, context plays a larger role in determining the referents spoken about than say the literal meanings of the denoting words employed.

Predication Consider the following utterance, at the edge of a lake, on a very hot August afternoon.

(52) The water is polluted.

It may be clear what the referent of "the water" is (the water in the lake). The intention of the utterance may also be clear: you better not go swimming in there, let alone drink from it. All of this can be very clear, even though many questions can be asked as to what the speaker means by saying that the water to be polluted. Is every bit of the water polluted? Or most, or just an unknown part of it? Or is there just a couple of barrels filled with chemicals at the bottom of the lake? The speaker may fail to know the answer, and may not even care a bit about these questions. Does this mean that the speaker doesn't know what she says? Or is what she says the disjunction of all affirmative responses to these questions, a proposition like "The water contains some substance which might be dangerous to your health"? Prima facie, this proposition is too weak to make any sense, and, on second thought, it may seem to stand in need of further analysis itself. Still, even without deciding upon a specific proposition, the purpose of the message is clear enough: don't go in the water and don't drink it. Something essentially similar holds for a statement like:

(53) The mountains are dangerous.

One may ask, "dangerous, in what sense?" "Are the slopes slippery? Are there mountain lions? Will the desert rocks have a fatal effect on my desperate mood?" In a given context, any such thing can be meant, or, as in the case of the polluted water, not even such a single thing. Still, again, the message can be clear: stay in the valley.

A typical example often discussed in the literature concerns the apparent 'simple' qualification of something as red.

(54) The car is red.

François Recanati remarks:

> "(…) in most cases the following question will arise: what is it for the thing talked about to count as having that colour? *Unless that question is answered, the utterance ascribing redness to the thing talked about (John's car, say) will not be truth-evaluable*. It is not enough to know the colour that is in question (red) and the thing to which that colour is ascribed (John's car). To fix the utterance's truth-conditions, we need to know something more—something which the meanings of the words do not and cannot give us: we need to know *what it is for that thing (or for that sort of thing) to count as being that colour*. What is it for a car, a bird, a house, a pen, or a pair of shoes to count as red? To answer such questions, we need to appeal to background assumptions and world knowledge." (Recanati 2005, p. 183)

The issue is pushed further in an example from Anne Bezuidenhout.

(55) Chris has been sorting red apples.

> "We're at a county fair picking through a barrel of assorted apples. My son says 'Here's a red one,' and what he says is true if the apple is indeed red. But what counts as being red in this context? For apples, being red generally means having a red skin, which is different from what we normally mean by calling a watermelon, or a leaf, or a star, or hair, red. But even when it is an apple that is in question, other understandings of what it is to call it 'red' are possible, given suitable circumstances. For instance, suppose now that we're sorting through a barrel of apples to find those that have been afflicted with a horrible fungal disease. This fungus grows out from the core and stains the flesh of the apple red. My son slices each apple open and puts the good ones in a cooking pot. The bad ones he hands to me. Cutting open an apple he remarks: 'Here's a red one.' What he says is true if the apple has red flesh, even if it also happens to be a Granny Smith apple." (Bezuidenhout 2002, p. 107)

Now one may wonder, can we say that Chris was sorting apples yesterday, and today he did it again? Can we say that what he did today was the same thing that he did yesterday? Surely, one can say he has been sorting apples both days, so, yes, it was the same thing, or the same kind of thing that he was doing. But if it comes to his sorting *red* apples, are we still inclined to say he has been doing that both days? It doesn't seem so. It also doesn't seem to be appropriate to say he was doing the same thing in a different way. It seems he was really doing a different thing these two days, but both activities are appropriately described as sorting red apples. Shouldn't we conclude from this, then, that 'sorting red apples' is ambiguous, at least? And when it comes to the core term in this example, 'red', shouldn't we conclude that it is ambiguous between at least the readings 'red on the outside' and 'red on the inside'? And, so the arguments goes, shouldn't we conclude that 'red' is ambiguous in many unexpected ways, so that, eventually, it cannot be assigned a meaning?

Quantification Quantifiers are typically context-sensitive as well.

(56) Some philosophers are not from France.

Besides the probably not uniquely answerable question who or what to count as a philosopher, a statement like this is often used with a specific domain of quantification in mind, so that it means something like "some philosophers attending the present conference," or "some philosophers we have been studying lately." The point here is, or so it has been argued, that one doesn't seem to construe or understand a most general and context independent proposition first, like "Some philosophers in the entire universe are not from France," and then derive from this a more specific, and more informative one, to the effect, for instance, that some of the philosophers at this conference are not from France. Rather, it seems, that in a given context one directly interprets the phrase "some philosophers" as "some philosophers at this conference," or any of the other specifications. But then, contextualists ask, what role would the most general, and trivial proposition have to play here in the first place?

While many facts about natural language determiners are basically well-understood, with the work originating from Barwise and Cooper (1981) and many subsequent studies, these very same expressions also interact with their linguistic context in ways which have not been explored and clarified fully yet. Here are two examples, a careful inspection of which requires an adaptation of the respective head determiners "All," "Five," "No," and "Most":

(57) { All / Five / No } boy scouts teased girl scouts and covered each other.

(58) Most orphans who got a present from their tutor thanked him for it.

Without going in the details a suitable interpretation of (57) has to adapt the basic interpretation of the noun phrase(s) "{ All / Five / No } boy scouts" to something which they may have done group-wise, and such that "cover each other" can be predicated of them. Likewise, the head determiner "Most" in example (58) must allow the restrictive property of being an orphan who gets a present from his or her tutor, to interact, or co-vary, with the scopal property of thanking the tutor for the present. This is not the place to actually see how these adaptations work out, but just to note that various proposals for sophisticated modified interpretations of these determiners have been made, by Kamp and Reyle (1993), Chierchia (1995), Nouwen (2003), and Brasoveanu (2008). These proposals do not converge on a most basic meaning of these determiners, but, rather, if combined, they create highly complicated interpretations apt for worst case scenarios, in which plurals and anaphora conspire, with who knows what else types of phenomena. It seems unlikely that under such an adaptive approach the 'Real Meanings' of the determiners are going to be uncovered. But then, or so it is argued, the sophisticated adapted interpretations should be properly conceived of as contextual adaptations themselves. Not shying the word, as 'pragmatically infected' interpretations.

Construction Simple numeral determiner phrases can be combined with a simple transitive verb and produce an equally simple cumulative reading:

(59) Twenty-five Dutch firms own three-hundred fourty-six Irish cows.

The, rather likely, cumulative, interpretation is that, if we look at the relation between x and y such that x are Dutch firms owning Irish cows y, then we count 25 owners, and 346 ownees. The example, and its interpretation, has been presented in Scha (1984) and Keenan (1992) has argued that it guides us beyond the Frege boundary—very roughly speaking, beyond first order logic. What this means, essentially, is that we have a construction, or context, in which the basic, or literal, noun phrase meanings are provably insufficient to produce the intended interpretation of the whole construction.

A related point arises when we reflect upon an example from Hendriks (1993), which is, deliberately it seems, rather unusual:

(60) Thomas seeks a fish or a bike.

This relatively simple seven word sentence can be assigned a large number of interpretations, using various tools developed and motivated in formal semantic systems. Here I list five of them, with corresponding glosses, but more can be generated.

- $SEEK(t, \lambda P(A\_FISH(P) \vee A\_BIKE(P)))$ ("Thomas is looking for whatever is a fish or bike; any such thing suits his purpose; hell knows why.");
- $(A\_FISH(\lambda x SEEK(t, \lambda P\ P(x))) \vee A\_BIKE(\lambda y SEEK(t, \lambda P\ P(y))))$ ("There is particular fish Thomas is looking for, or a particular bike; I don't know which of the two he has set set his mind on.");
- $(SEEK(t, A\_FISH) \vee SEEK(t, A\_BIKE))$ ("Thomas is looking for whatever fish, or he is looking for whatever bike, I don't really know which of the two.");

- $A\_FISH(\lambda x\, A\_BIKE(\lambda y\, (SEEK(t, \lambda P\, (P(x) \vee P(y))))))$ ("There is a (particular) fish and a (particular) bike such that Thomas seeks whatever is either one of the two. God moves in a mysterious way.");
- $(SEEK(t, A\_FISH) \vee A\_BIKE(\lambda y\, SEEK(t, \lambda P\, P(y))))$ ("Thomas seeks whatever fish, or a particular bike. I am so stoned.")

No matter how unlikely the sentence or one of its glosses is, with a bit of creative imagination one can think up a situation exactly described by the gloss, and the idea then is that one might be tempted to report that situation by uttering (60). If so, and I believe I could do so, one might be tempted to, after all, conclude, that the sentence (60) is so multiply ambiguous, and that it has all these readings. This would be correct, if by "multiply ambiguous" we mean "subject to a multitude of interpretations." However, it seems very hard indeed to believe that such a, relatively simple, sentence has all these readings—that knowing its meaning boils down to realizing all these readings, and that in any context of utterance we have to select an appropriate one of them. Rather, or so the contextualist can argue, the sentence has only one analysis, e.g.:

(61) $[_S[_{NP}$ THOMAS$][_{VP}$ SEEK$[_{NP}[_{NP}$ A_FISH$]$ OR $[_{NP}$ A_BIKE$]]]]$

It is only in a specific context that the building blocks get put together, in a specific way, which hearer and speaker may quarrel about and negotiate.

Considerations like those given above, and many others can be supplied at will, lead, unavoidably it seems, to the conclusion that there is an indefinite number of fairly unanswerable questions about the core mechanisms of meaning in actual practice—viz., reference, predication, quantification, and construction; moreover, assuming that language derives its meaning, not because it is a God given heavenly instrument, but from its actual use, to claim that these meanings are well established and firmly given seems preposterous and it seems that in certain, or maybe all, contexts their specific details escape our notice, or are derived *despite* these firmly given meanings.

As a matter of fact, rather destructivist conclusions have been drawn from such observations. "The evidence in favour of contextualism is provided by indefinitely many examples in which the same sentence, which does not seem to be ambiguous, is used in different contexts to say different things (Recanati 1994, p. 9)". "According to these philosophers, sentences can never express complete propositions independent of context, however explicit speakers try to be. In other words, content is always under-determined by the linguistic material (Recanati 2006, pp. 22–23)." Emma Borg summarizes the current feeling as follows: "These days, the natural descendent of the formal approach, known as minimalism, has been consigned to the margins: not everyone rejects minimalism, but lots of people do. Minimalism is rejected in favour of contextualism: roughly, the idea that pragmatic effects are endemic throughout truth-evaluable semantic content (Borg 2007, p. 339)."

Even minimalists who defend the formal semantic semantic enterprise tend to agree with these negative conclusions. Emma Borg, for instance, submits that "according to the minimalist (as I construe her) there is an entirely formal route to

meaning. This means not only that every contextual contribution to semantic content must be grammatically marked but also that those features contributed by the context must themselves be formally tractable (Borg 2007, p. 358).” This amounts to saying that there is hardly any formal meaningful structure which is not pragmatic. Cappelen and Lepore, who also defend semantic minimalism, as well as ‘Speech Act Pluralism’ (‘SPAP’), conclude that “SPAP is not really a theory; it’s a collection of observations, one of which is that there can be no *systematic theory* of speech act content ((Cappelen and Lepore, 2005)[p. 190]).” In the same spirit, Pagin and Pelletier (2007) aims to preserve a classical compositional architecture of interpretation, but one that allows Recanati’s pragmatic processes of enrichment, loosening, and semantic transfer to interfere on every node in the construction tree, or in the corresponding meaning-tree. This, obviously, means taking the contextualist’s arguments to heart, and surrender to the pragmatic hegemony which practically may overrule anything that has been called semantic ever before in the ensuing architecture.

Giving things a thought, however, I don’t believe the situation is so dramatic. The preceding discussion may paint a bleak future for semanticists, if one believes, along with the early Wittgenstein and other analytic philosophers, in something like the ultimate analysis of a sentence. Some such thing may speak from the following propositions [(from Wittgenstein (1922)]:

> 2.0201 Jede Aussage über Komplexe lässt sich in eine Aussage über deren Bestandteile und in diejenigen Sätze zerlegen, welche die Komplexe vollständig beschreiben.
> 2.0211 Hätte die Welt keine Substanz, so würde, ob ein Satz Sinn hat, davon abhängen, ob ein anderer Satz wahr ist.
> 2.0212 Es wäre dann unmöglich, ein Bild der Welt (wahr oder falsch) zu entwerfen.

The reduction ad absurdum in proposition 2.0211 relies on the assumption, not disputed there, that sentences in the end have meanings, independent of further analysis—that is, independent of other sentences being true or not. The later Wittgenstein, has quite nicely reformulated this intuition, in order to seriously qualify it next. In §99, Wittgenstein restates one of the main tenets from the Tractatus as follows:

> §99 Der Sinn des Satzes – möchte man sagen – kann freilich dies oder das offen lassen, aber der Satz muß doch einen bestimmten Sinn haben. Ein unbestimmter Sinn, – das wäre eigentlich gar kein Sinn. – Das ist wie: Eine unscharfe Begrenzung, das ist eigentlich gar keine Begrenzung. Man denkt da etwa so: Wenn ich sage “ich habe den Mann fest im Zimmer eingeschlossen – nur eine Tür ist offen geblieben” – so habe ich ihn eben gar nicht eingeschlossen. Er ist nur zum Schein eingeschlossen. Man wäre geneigt, hier zu sagen: “also hast du damit garnichts getan”. Eine Umgrenzung, die ein Loch hat, ist so gut, wie gar keine—Aber ist das denn wahr?

But is this true? Eventually, he conceives of it as a misconception in the *Philosophische Untersuchungen*. His picture of the road indicator in paragraph §85 is revealing in this context.

§85 Also kann ich sagen, der Wegweiser läßt doch keinen Zweifel offen. Oder vielmehr: er läßt manchmal einen Zweifel offen, manchmal nicht. Und dies ist nun kein philosophischer Satz mehr, sondern ein Erfahrungssatz.

Many questions can be raised as to what a Wegweiser (road indicator) exactly indicates, what direction it shows exactly, and with what amount of freedom. Reflecting on all kinds of deviant contexts, it seems that we don't have a full answer to the question what the road indicator in all of these contexts *really* means. (For instance, what is its meaning in the middle of the desert, or in five-dimensional space?) But as a matter of fact, in all normal circumstances, any non-deviant road indicator is perfectly clear, and we all know how to read it, and how to act upon what it signals in all or most of our run of the mill activities in everyday life. The observation that the indicator works well as long as it is embedded in our regular practices, is maybe all that needs to be stated. It doesn't require a Jack-in-the-Box meaning that makes clear to us what its function in a particular context is. Already in the first paragraph of the *Philosophical Investigations*, Wittgenstein has pointed out that every analysis comes to an end, not because it is the final analysis, but because we simply stop there. "Nun, ich nehme an, er handelt, wie ich es beschrieben habe. Die Erklärungen haben irgendwo ein Ende (Wittgenstein 1953, §1)."

Jason Stanley defends the formal semantic paradigm as follows:

> Competent English speakers know the meanings of the words in the sentence 'Some philosophers are from New York.' They also know how to combine the meanings of each of the words in this sentence to arrive at what is said by the utterance of the sentence, 'Some philosophers are from New York.' It is that linguistic competence that seems to be the source of their ability to report correctly about the truth of what is said by that sentence relative to different possible circumstances (…). (Stanley 2005, p. 221)

In response to the observation that more enters into the meanings of utterances, Stanley notes: "The first response (…) is to attempt to preserve the clear and elegant explanation in the face of the apparently recalcitrant data. The second is to abandon the clear and elegant explanation of the source of our truth-conditional intuitions in favor of a different one (Stanley 2005, p. 222)." Clearly, he prefers the first, not wanting to abandon "the project of giving a systematic explanation of the source of our intuitions" and instead "appeal to unconstrained and non-explanatory notions or processes." Even with the contextualist observations in the back of our minds to keep on carrying out such a program makes sense.

We all have an idea of objects we refer to, even though we are unable to define what an object is; we all know how to predicate stuff of things, even though we fail to explain what any specific predication exactly amounts to; we all know what it means to say that everybody, or most of my fellow students, have a certain property, even though we cannot say for each occasion what it exactly says that the property applies to all or most members of a certain group. In formal semantics one tries and lays out the things which are known, without committing oneself to the things which are unknown. Nobody denies that there are all kinds of structural properties of languages, also on the level of meaning, 'meaning' intuitively understood here. A semanticist can be justified to claim to work on revealing such structural properties,

by simulating meaningful languages in formal models. This is not to say that these models, or some entities or constructions in there, depict the 'Real Meanings' of natural language utterances. Rather, they are coarse-grained simulations of parts of meaningful language, displaying structural properties of meaning found in real languages, or so we hope. Nothing in this picture entails that they should be like basic meanings, which language users have to grasp, in order to put them into practice. And why should one be tempted to do so? Surely one should not want to draw analogous conclusions from the models economists, or sociologists, or biologists, build to simulate, more or less realistically, economical, sociological, or biological processes.

Language is an organism, which is alive, in action and in evolution, and which is inherently social. Meaning, originally and eventually, lies in use. From a semantic point of view, where semantics is conceived of as the (study of the) theory of meaning, we cannot seriously close our eyes to pragmatic intrusions of meaning. No autonomous semantics of natural language seems to be forthcoming, which can serve as the input to a pragmatic theory of interpretation. But then, even if we adopt this (radical) pragmatic perspective, we are still confronted with all kinds of systematic, structural, meaningful aspects of natural language. Devices which systematically (but not unavoidably) serve certain purposes, and certain structural purposes. Even from this pragmatic perspective, bringing such aspects of use to light, and bring them together in an orderly system is certainly worthwhile. Indeed, no system accomplishes this better than a formal system, and, then preferably, a system that reflects, or models those systematic aspects of use. This reflection, set up in a fully general form, by itself is not a theory of meaning, as it is not a system of interpretation. It does not mimick or aim to mimick the ways in which agents produce and interpret language, or their capacity to do so. It does, however, bring to light these systematic properties, no matter how non-robust, contingent, and non-persistent they may be in reality. Yet, bringing these structures to light, brings important aspects of language to light. Conceived in this manner, semantics does not constitute a theory of meaning; but, turning the table, I wouldn't see either how pragmatics would ever offer anything like a theory of the envisaged kind, if it has not extrapolated a job for the semanticists.

The upshot of this whole discussion is that from the old-fashioned semanticists' point of view, the pragmaticists may have won, or are in the process of winning the battle, and the semanticist may have to surrender. But surely the pragmaticists, who claim victory, like they did in the fifties and sixties of the previous century, have nothing on offer instead. Besides case studies of single expressions in singular contexts, they cannot supply us with any more than entirely contingent descriptions, which are always subjective, or at least not objectifiable. If any real work is to be done, they will have to agree that the old fashioned semanticists are the ones that do the job, and build models, which do generalize to some extent. With them only, theories will come off the ground. Upon further reflection the situation is even worse. Herman Cappelen and Ernest Lepore argue that radical contextualism appears to be internally inconsistent. "To interpret the sentences that express RC [Radical

Contextualism, PD] you have to assume that RC is *not* true." (Cappelen and Lepore, 2005)[extensively discussed on pp. 128–139]

In this whole picture dynamic semantics plays a peculiar role. It has been put forward as a radical departure from old-fashioned theories of meaning, it has been advertised as the new semantics which steals ground from the pragmaticists, and cultivates it. This may be correct, historically speaking, but the situation is conceived from a different perspective now. All semantics finds its roots in pragmatics, and flourishes from there. Yet, there is good reason to say that there are semantic structures and regularities to be revealed. Dynamic semantics can then be conceived of as the kind of *real* semantics which wears its pragmatic roots on its sleeves.

References

Aloni, Maria. (2000). Conceptual Covers in Dynamic Semantics. In Nick Braisby Lawrence Cavedon, Patrick Blackburn, & Atsushi Shimojima (Eds.), *Logic* (pp. 23–48). CSLI: Language and Computation, Vol III, Stanford, CA.

Aloni, Maria. (2005). Individual Concepts in Modal Predicate Logic. *Journal of Philosophical Logic, 34*, 1–64.

Aloni, Maria. (1998). Quantification in Dynamic Semantics. In Jerry Seligman & Patrick Blackburn (Eds.), *Proceedings of ITALLC'98* (pp. 82–92). Chiayi, Taiwan: National Chung Cheng University.

Bar-Hillel, Yehoshua. (1971). Out of the pragmatic waste-basket. *Linguistic Inquiry, 2*, 401–07.

Barwise, Jon, & Cooper, Robin. (1981). Generalized Quantifiers in Natural Language. *Linguistics and Philosophy, 4*(1), 159–220.

Beaver, David 1995. Presupposition and Assertion in Dynamic Semantics. Ph.D. thesis, CCS, Edinburgh. Published in 2001 by CSLI Publications, Stanford.

Bezuidenhout, Anne. (2002). Truth-conditional pragmatics. *Philosophical Perspectives, 16*, 105–34.

Borg, Emma. (2007). Minimalism versus contextualism in semantics. In Gerhard Preyer & Georg Peter (Eds.), *Context-Sensitivity and Semantic Minimalism* (pp. 339–360). Oxford: Oxford University Press.

Brasoveanu, Adrian. (2010). Decomposing Modal Quantification. *Journal of Semantics, 27*, 437–527.

Brasoveanu, Adrian. (2008). Donkey pluralities: plural information states versus non-atomic individuals. *Linguistics and Philosophy, 31*, 129–209.

Cappelen, Herman, & Lepore, Ernie. (2005). *Insensitive Semantics*. Oxford: Blackwell.

Chastain, Charles. (1975). Reference and Context. In Keith Gunderson (Ed.), *Language, Mind, and Knowledge* (pp. 194–269). Minneapolis: University of Minnesota Press.

Chierchia, Gennaro 1995. Dynamics of Meaning. Anaphora, Presupposition, and the Theory of Grammar. Chicago: The University of Chicago Press.

Cresswell, Max J. (2002). Static Semantics for Dynamic Discourse. *Linguistics and Philosophy, 25*, 545–571.

Donnellan, Keith. (1978). Speaker Reference, Descriptions and Anaphora. In Peter Cole (Ed.), *Syntax and Semantics* (Vol. 9, pp. 47–68). Pragmatics, New York: Academic Press.

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

van Eijck, Jan, & Cepparello, Giovanna. (1994). Dynamic Modal Predicate Logic. In Makoto Kanazawa & Chris Piñon (Eds.), *Dynamics* (pp. 251–276). CSLI: Polarity and Quantification, Stanford.

Evans, Gareth. (1982). *The Varieties of Reference*. Oxford: Clarendon Press. Edited by John McDowell.

Groenendijk, Jeroen, & Stokhof, Martin. (1991). Dynamic Predicate Logic. *Linguistics and Philosophy, 14*(1), 39–100.

Groenendijk, Jeroen, Stokhof, Martin, & Veltman, Frank. (1996). Coreference and Modality. In Shalom Lappin (Ed.), *The Handbook of Contemporary Semantic Theory* (pp. 179–213). Oxford: Blackwell.

Groenendijk, Jeroen 2007. The Logic of Interrogation. In: Maria Aloni, Alastair Butler & Paul Dekker (eds.) Questions in Dynamic Semantics. CRiSPI 17, Amsterdam: Elsevier, pp. 43–62. Also appeared in T. Matthews and D. Strolovitch (eds.), 1999, Proceedings of SALT IX, CLC Publications.

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

Heim, Irene, 1991. On the Projection Problem for Presuppositions. In: Steven Davis (ed.) Pragmatics: A Reader, Oxford: Oxford University Press. pp. 397–405. Also appeared in Barlow, M. and D. Flickinger and M. Wescoat, (Eds.). (1983). *Proceedings of WCCFL II*. Stanford: CA.

Hendriks, Herman 1993. Studied Flexibility. Ph.D. thesis, University of Amsterdam.

Kamp, Hans, & Reyle, Uwe. (1993). *From Discourse to Logic*. Dordrecht: Kluwer.

Kamp, Hans 1990. Prolegomena to a Structural Theory of Belief and Other Attitudes. In: C. Anthony Anderson & Joseph Owens (eds.) Propositional Attitudes, Stanford: CSLI. pp. 27–90.

Keenan, Ed. (1992). Beyond the Frege boundary. *Linguistics and Philosophy, 15*, 199–221.

Kripke, Saul 1972. Naming and Necessity. In: Donald Davidson & Gilbert Harman (eds.) Semantics of Natural Languages, Dordrecht: Reidel. pp. 254–355, 763–769.

Kripke, Saul 1979b. Speaker's Reference and Semantic Reference. In: Peter A. French, Theodore E. Uehling & Howard K. Wettstein (eds.) Contemporary Perspectives in the Philosophy of Language, Minneapolis: University of Minnesota Press. 6–27.

Moltmann, Friederike. (2006). Unbound anaphoric pronouns: E-type, dynamic, and structured-propositions approaches. *Synthese, 153*, 199–260.

Nouwen, Rick 2003. Plural pronominal anaphora in context: dynamic aspects of quantification. Ph.D. thesis, UiL-OTS, Utrecht. LOT dissertation series, No. 84.

Nunberg, Geoffrey. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy, 3*, 143–184.

Pagin, Peter, & Pelletier, Francis Jeffry. (2007). Content, context, and composition. In Gerhard Preyer & Georg Peter (Eds.), *Context-Sensitivity and Semantic Minimalism* (pp. 25–62). Oxford: Oxford University Press.

Read, Stephen. (1992). Conditionals are not truth-functional: an argument from Peirce. *Analysis, 52*, 5–12.

Recanati, François. (2005). Literalism and contextualism: Some varieties. In Gerhard Preyer & Georg Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (pp. 171–196). Oxford: Oxford University Press.

Recanati, François. (1994). Contextualism and Anti-Contextualism in the Philosophy of Language. In S. L. Tsohatzidis (Ed.), *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives* (pp. 156–66). London and New York: Routledge.

Recanati, François. (2006). Crazy Minimalism. *Mind and Language, 21*(1), 21–30.

Scha, Remko J.H. 1984. Distributive, Collective and Cumulative Quantification. In: Jeroen Groenendijk, Theo M.V. Janssen & Martin Stokhof (eds.) Truth, Interpretation and Information, Dordrecht: Foris. pp. 131–158.

Stalnaker, Robert. (1978). *Syntax and Semantics 9—Pragmatics* (p. 332). New York: Academic Press.

Stalnaker, Robert. (1998). On the Representation of Context. *Journal of Logic, Language and Information, 7*(1), 3–19.

Stanley, Jason. (2005). Semantics in context. In Gerhard Preyer & Georg Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (pp. 221–253). Oxford: Oxford University Press.

Strawson, Peter F. (1952). *Introduction to Logical Theory*. London: Methuen & Co.
van den Berg, Martin H. 1996. The Internal Structure of Discourse. Ph.D. thesis, Universiteit van Amsterdam, Amsterdam.
van Rooy, Robert 1997. Attitudes and Changing Contexts. Ph.D. thesis, IMS, Stuttgart.
Veltman, Frank. (1996). Defaults in Update Semantics. *Journal of Philosophical Logic, 25*(3), 221–261.
Vermeulen, Cornelis F. M. (1996). Merging Without Mystery or: Variables in Dynamic Semantics. *Journal of Philosophical Logic, 24*(4), 405–450.
Wittgenstein, Ludwig. (1953). *Philosophische Untersuchungen/Philosophical Investigations*. Oxford: Blackwell. Translation by G.E.M. Anscombe.
Wittgenstein, Ludwig (1922). Tractatus Logico-Philosophicus. Oxford: Routledge and Kegan. Originally appeared in 1921 as Logisch-Philosophische Abhandlung in the Annalen der Naturphilosophie, 14.
Zimmermann, Ede. (1999). Remarks on the Epistemic Rôle of Discourse Referents. In Lawrence S. Moss, Jonathan Ginzburg, & Maarten de Rijke (Eds.), *Logic, Language and Computation* (Vol. II, pp. 346–368). Dordrecht: Kluwer.