UvA-DARE (Digital Academic Repository)

Information in natural language

Kamp, H.; Stokhof, M.

Publication date
2008

Published in
Philosophy of information

[Link to publication](#)

Citation for published version (APA):

Kamp, H., & Stokhof, M. (2008). Information in natural language. In P. Adriaans, & J. van Benthem (Eds.), *Philosophy of information* (pp. 49-112). (Handbook of the philosophy of science; No. 8). North-Holland. <http://staff.science.uva.nl/~stokhof/inl.pdf>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Information in Natural Language

Hans Kamp* Martin Stokhof†

To appear in:

Pieter Adriaans and Johan van Benthem (eds), *Handbook of the Philosophy of Information*, Elsevier, Amsterdam

I Introduction

Natural languages are vehicles of information, arguably the most important, certainly the most ubiquitous that humans possess. Our everyday interactions with the world, with each other and with ourselves depend on them. And even where in the specialised contexts of science we use dedicated formalisms to convey information, their use is embedded in natural language.¹

This omnipresence of natural language is due in large part to its flexibility, which is almost always a virtue, sometimes a vice. Natural languages are able to carry information in a wide variety of ways, about a seemingly unlimited range of topics, which makes them both efficient and versatile, and hence useful in almost every circumstance. But sometimes, when pinpoint precision is what counts, this versatility can get in the way, and we make use of formal languages, such as those of mathematics.

The variety of ways in which the use of natural language involves information, reveals itself immediately if we look at the various functions that utterances of natural language expressions may have. First, many of the utterances we produce serve to directly impart information to our readers or listeners — usually information which we take to be new and of interest to them. We describe situations, stating what we take to be facts (‘Mary is in Paris’), or contemplating what we regard as possibilities (‘John might be accompanying her’). This declarative use of language is perhaps the most obvious way in which natural languages are used to convey information.

But, of course, this doesn’t hold for all utterances. We also ask questions (‘What time does the meeting start?’), in order to elicit information rather than to impart it; we give directives (‘Open a window’, ‘Stay away from her’), in order to get the other to do certain things, or to keep him from doing them; we issue warnings and threats (‘Look out, a bus!’), ‘If you do that, I’ll tell the boss’), we express regret and joy (‘I apologise

* Institut für maschinelle Sprachverarbeitung, Universität Stuttgart

† ILLC / Department of Philosophy, Universiteit van Amsterdam

1. We would like to thank the editors for their patient support and Menno Lievers for his detailed critical comments, which have led to many improvements. Obviously, only the authors are to be held responsible for the end result.

for the belated reply, ...', 'Congratulations!'), and so on. But in these cases, too, our utterances carry information, and that they do so is essential: a question must convey what information is requested; a directive must specify information about what is to be done, or to be refrained from; a warning or threat must identify a particular situation or event; and if we don't convey what it is that we regret, or what we are happy about, the point of our speech is lost.

These humdrum observations also illustrate a further point. Not only do natural language utterances involve information about a variety of types of situations: factual and possible, past, present and future; they also convey information about the specific attitudes that natural language users have concerning these situations: that the speaker takes them to be factual, or merely possible; that they are to be avoided, or to be realised, by the hearer; that the speaker regards them with regret, or with joy. Thus the information carrying capacity of a natural language encompasses not just what its expressions are about, but also the various attitudes that its users may have towards that.

Another way in which information might be conveyed is more indirect than in the examples above, where it is coded in the syntactic form or indicated by a particular expression or turn of phrase. Making use of the context in which an utterance is produced, for example by relying on the presence of certain expectations on the part of the hearer, we may also indirectly convey information about a certain situation. For example, when answering a question about the whereabouts of Jane by means of a disjunction ('She's either in Paris, or in London, with Mary'), we indicate that we do not know exactly where she is. This is information that is not explicitly stated, but only suggested. However, an addressee who expects the speaker to be as co-operative as he can will pick up this information without hesitation.

And it doesn't stop there. Besides utterances of the above kinds there are those which serve various social purposes: greeting someone, acknowledging a gesture or utterance she is making, expressing concern or empathy. Many utterances we use to such ends — 'Hi, how are you?', 'I am sorry to hear that', and so on — are formulaic. They carry information, not in virtue of being about something and expressing an attitude to that, but by being fixed through special conventions, which bypass the general mechanisms by which information is associated with linguistic form. But these utterances do carry information nonetheless, as is indicated by the fact that the purposes they serve can as a rule also be accomplished by means of other, non-formulaic utterances.

Yet another way in which information is conveyed by natural language is through mechanisms that relate a specific utterance to its linguistic context. After all, an utterance hardly ever occurs on its own, out of the blue; usually it is part of a larger whole, a text or a conversation, that serves a specific purpose and accordingly may have a specific form. When it is part of such a larger textual or conversational complex an individual utterance may contribute to the meaning of that complex as a whole through mechanisms that relate it to other parts of the complex. The use of pronouns to refer to an entity mentioned previously in a conversation (*A*: 'I met John the other day.' *B*: 'How's he doing?') is a simple example; the specific form of a question – answer dialogue, in which answers more often than not are fragments of complete sentences, yet do express complete propositions (*A*: 'Who's chairing the meeting?' *B*: 'Bill.'), provides

another.

As they stand, all these observations, with their repeated references to ‘information’, are, we take it, hardly controversial. But to say precisely what the information is of which they speak is not so easy. For one thing, it is not at all clear that we are dealing with a uniform concept: when trying to explain in what sense natural languages are information carriers, we may well find that it is necessary to distinguish various ‘kinds’ of information. And should that need arise, there will be the further task of saying exactly how these different notions are related to each other and how natural languages are able to handle various notions of information in such elegant and efficient ways. To outline some aspects of the current state of thinking about these issues is the goal of the present chapter.

But first we should make clear what the information concept that we talk about in this chapter is not. We are not concerned with information based on mere likelihood, according to which the information carried by a symbol or symbol string is some inverse function of the probability of its occurrence. Common to such a concept of information and the one that will be relevant in this chapter is the conception of individual events that are classified as each being of some particular event type. In our case the event types are the types of symbols and symbol strings and the individual events are particular occurrences (‘utterances’) of symbols and strings of them. The probability-based notion of information presupposes in addition to such a space of classifiable occurrences a probability distribution over possible occurrences, which assigns each occurrence of an individual event an a priori probability in terms of the classification-related properties it has. On the other hand, what is essential to the concept of information that will be discussed here is that symbols and symbol complexes have denotations, i.e., that they stand for, or represent, entities and situations, and that the information they carry is about those denotations.

On a simple-minded, purely causal conception of how symbols denote the two conceptions of information would be compatible. On such a view, the occurrence of symbols (both simple and complex) is prompted by the occurrence of their denotations. So the space of symbol occurrences maps onto a corresponding space of denotations, and the probability of a symbol occurrence is the direct reflection of the occurrence probability of the denotation that is its cause. In that case the information represented by the occurrence of a given symbol would be the occurrence of its denotation and the quantity of that information could be meaningfully assessed in terms of the probability that the denotation should have occurred. We will see, however, that in connection with natural languages such a conception of denotation is untenable. The simple causal nature of the denotation relation which it presupposes is belied, both at the level of the simple symbols of the language (its ‘words’) and at that of its complex symbols (its phrases, sentences, texts and conversations), by the way in which natural languages actually work.

The first, and most widely acknowledged difficulty with a purely causal conception of denotation concerns the denotation of complex symbols. The denotations of phrases and sentences are determined by their syntactic form and by the denota-

tions of the words from which they are made up. In principle the recursive process by which the denotations of complex symbols are built from the denotations of their constituents might have a purely causal grounding. But any serious exploration of the way in which natural language expressions denote soon reveals the extreme implausibility of this. Complex expressions denote what they do because of the denotation building rules that are part of the language as a conventional system; and speakers can use these expressions to refer to their denotations because they know those rules and thus know that a given complex phrase or sentence does have the denotation to which they want to refer. In other words, information as conveyed by natural language utterances depends on a conceptualisation of what the information is about that, at least to a large extent, is shared between the users of the language.²

The existence of a set of conventional rules for building complex expressions to denote complex entities or situations is something that natural languages share with the formal languages of logic, mathematics, and computer science. But, as we will argue in some detail below, there are also important differences between natural and formal languages. One of these is that in natural languages the principles which govern the building of expressions to denote complex things or situations are far more complex than the comparatively straightforward recursive principles that define formal languages (like those of the predicate calculus or the lambda-calculus). This greater complexity is connected with the remarkable flexibility and adaptability of natural languages, which makes it possible to use them for the purpose of conveying information about a vast and open-ended range of different subjects.

This is connected with another feature of natural languages, viz., that they can be used to speak about non-existent objects, unrealised situations. In some cases expressions and linguistic constructions are even meant to do just that, e.g., when we use a counterfactual sentence ('If I had left home earlier, I wouldn't have missed the five o' clock train'), in other cases the possibility is left open, e.g., when we utter a conditional sentence ('If John comes to the party too, Mary will be upset', where the actual appearance of John is neither affirmed nor denied). And then again we may be convinced that what we say is true, whereas in fact things are not as we assert them to be. 'Information', then, is used here in such a way that it can also be false (just as it can be misleading, or partial): the information provided by an utterance, i.e., what anybody who understands its linguistic meaning might assume to be the case, need not actually hold. By no means should this be considered as a defect of natural languages. In fact, it is an unavoidable consequence of the partial and fallible nature of human knowledge and our ability to imagine what we know or have reason to think is not the case on the one hand, and

2. More on this below, in section 3. It should be noted that this observation is not meant to rule out the possibility of 'non-conceptual content': it pertains to the information expressed by means of utterances of linguistic expressions only and remains neutral with respect to the question whether objects, events, situations, —including linguistic expressions and their use—, may also convey information of a different nature. Also note that we take utterances (i.e., the production of 'tokens') to be primary to expressions (conceived of as 'types') when it comes to what are entities that carry information. But in as much as there are systematic relations between the two, we sometimes also talk about expressions in that vain. We assume that no confusion will arise from this.

on the other the fact that natural languages are means of expressing not only what we think is the case, but also what we suspect may be the case, what we hope, fear, would wish to be the case, and so on.³

There is an obvious connection between denotation and meaning: the meaning of a linguistic expression is given by what it denotes, in actual situations or in non-actual ones. Since the notion of linguistic information we are after is also closely tied to denotation, there is an intimate connection between linguistic information and linguistic meaning. The fact that both linguistic meaning and linguistic information are connected with denotation entails an important moral for either. Both linguistic meaning and linguistic information are inherently relational concepts, both involve the form-governed relation between linguistic expressions and their denotations. This is a moral that some would consider too obvious to merit stating. But the relational nature of meaning is not something that has always been self-evident to everyone. In fact, the moderately clear picture of the ways in which the meanings of linguistic expressions are relational that we now possess is the outcome of a long process of philosophical analysis. Because the two notions are so closely intertwined the history of the concept of linguistic meaning is at the same time also the history of linguistic information. Therefore we will devote the first part of this chapter to tracing what from the perspective of one particular tradition, viz., that of formal semantics and its immediate predecessors, are seen as some of the salient stations in the historical development that has led up to the current state of thinking on these issues. Probably, from other angles different pictures would emerge, but it is beyond the scope of this chapter to sketch those as well. So the emphasis is on history as perceived by the discipline, not as it actually occurred (if there is such a thing). For it is the former, and not the latter, that best explains its development.

This concise historical overview also shows how the formal semantics tradition has struggled to come to grips with the variety of ways in which natural language utterances carry information that we briefly touched upon above. That process is one of both contraction and expansion, as is so often the case in the development of a scientific discipline. At one stage there is a focus on one specific aspect of a phenomenon, which often allows the use of formal tools and leads to precise accounts. At another stage the resulting notions are extended to deal with other aspects, or complemented by other notions, as the situation may require.

After the historical sketch we turn to an overview of various elements that would

3. It is this very potential of natural languages to be about non-actual objects and situations that according to Frege liberates the human mind and sets us apart from other animals. In his 'Über die wissenschaftliche Berechtigung einer Begriffsschrift' (Frege (1882); the translation is Bartlett's Frege (1964)) he writes:

Nonetheless, our imagery [...] would be limited to that which our hand could form, our voice intone, if it were not for the grand discovery of the symbol which calls to our mind that which is absent, out of sight or perhaps even unseeable.

And George Steiner regards the very possibility it gives us to talk about the non-actual as fundamental for human language (Steiner, 1975, page 215):

Hypotheticals, 'imaginaries', conditionals, the syntax of counterfactuality and contingency may very well be the generative centres of human speech.

be needed in an adequate theory of natural language information. Such a theory must do several things. First of all, it should give an account of how expressions of natural language come to have meaning, and of the ways in which meaning depends on the context of utterance. This involves, among others things, coming up with a suitable set of formal concepts that can be used to define adequate representations of natural language meanings, to model the relevant features of context and to characterise the way in which these interact. Second, the scope of an adequate theory should encompass the fact that natural languages are used in interactive situations: natural languages are used to convey meaning for a reason, and that reason lies in information exchange, broadly conceived. Thus, the information conveying capabilities of natural languages are tailored to their use in a discourse context, be it dialogical, textual, or of some other form. As a matter of fact, many features of these capabilities depend on structural features of such discourses, which, hence, need to be modelled. Third, it is the language users that need to be taken into account: what information the utterance of a natural language expression conveys, and how it does that, obviously depends also on the language users involved, both as speakers and as hearers. Modelling them, then, is yet another essential ingredient of an adequate theory of natural language information. The overview we will be giving follows the broad outlines of an approach that has become well-established in empirical work in natural language semantics over the last couple of decades. But we should emphasise that it is not this particular approach we want to propagate, but rather the underlying ideas that together form a theme that allows for many variations: some of these we will refer to when appropriate.

2 A tale of two developments

As is so often the case, a systematic discipline has a somewhat distorted picture of its own history, one that usually takes ‘a bird’s eye view’ and focuses on those aspects that retain a certain present relevance. For us, taking such a bird’s eye view of what we see as the important stages in philosophical and linguistic thinking about the concepts of information and meaning, one development that stands out is that from ‘thick’ and fairly concrete notions of meaning, closely tied to (perceptual) experience, judgement and application, to rather ‘thin’ and abstract conceptions and (ultimately) to a view of natural languages as purely information coding and information transferring devices.⁴ This line of development is complemented, however, by another one that extends the restricted, ‘descriptive’ conception of linguistic meaning that is the outcome of the former, and tries to enrich it by encompassing a wider variety of aspects and by reinstating connections with other components of human cognition. The first development is assumed to have its starting point in traditional philosophical thinking,⁵ it gains momentum with

4. This is not unlike Göran Sundholm’s Sundholm (to appear) view on development of logic: from a theory about judgements and reasoning as psychological acts to (ultimately) formal symbol manipulation. The distinction between ‘thick’ and ‘thin’ concepts is taken from Bernard Williams, who developed it with regard to ethical concepts.

5. In this short sketch we limit ourselves to developments in Western philosophy. That is not to deny that very interesting theories and views, that are highly relevant from a systematic point of view, have been

the development of formal logic at the end of the nineteenth century, comes to fruition in the 1970's and 1980's, and still remains strong until the present day. The latter development is partly a reaction to the former and mainly dates from the last two or three decades, after systematic thinking about linguistic meaning developed into a distinct discipline.

2.1 Uncovering structure

An, admittedly very rough, sketch of the first line of development distinguishes the following stages. At the first stage, which is assumed to start in classical philosophy and to extend right up to the rise of modern philosophy in the sixteenth and seventeenth century, thinking about the concept of meaning usually is intimately related with metaphysical and epistemological concerns. The latter obviously take precedence, and meaning, and language more generally, as such are by and large not distinct and independent topics of concern. Language and linguistic meaning are viewed and analysed primarily as means to express judgements, and it is the origin, content and justification of judgements that most philosophers are interested in.

For example, Plato's discussion of the possibility of false statements in the *Sophist* is motivated by a metaphysical concern about the possibility of knowledge and the relation between thought and reality, not by any autonomous interest in natural language meaning.⁶ Similarly, the main motivation behind the work of the scholastics on language and logic is metaphysical (and some of it theological). And the 'idea theories of meaning' of the classical empiricism and rationalism of the sixteenth and seventeenth centuries are mainly motivated by questions and problems in epistemology.⁷ From a modern, systematic perspective there seems to be no such thing as a separate philosophy of language in this period, nor is there a distinct and substantial empirical discipline of linguistics that is concerned with the analysis of natural language meaning for its own sake.⁸ It would take a century or so for linguistics to really come into its own, with the

developed in other traditions. Especially in India there is a rich tradition of sophisticated thinking about language, as is witnessed by the great works of Pāṇini and other Indian grammarians (cf., Cardona (1988, 2nd ed 1997)). However, historically these have not played a major role in shaping present day theories in semantics, and it is for that reason that we feel it is justified to leave them out.

6. Thus the discussion of word and sentence meaning and of truth and falsity, in the *Sophist*, 261c6–264b3 Plato (1921), ends as follows:

Then because speech, we saw, is true and false, and thinking is a dialogue of the mind with itself, and opinion is that completion of thought, and what we say by "it seems" is a combination of perception and opinion, it must be that because all of these are like speech, some thinking and opinion must also be false.

Evidently, the linguistic analysis is subservient to the metaphysical point that Plato wants to make.

7. Hence Ian Hacking Hacking (1975) called idea theories 'nobody's theory of meaning': since meaning as such is not a separate concern, nobody had a *theory* about it, or even felt the need to come up with one.

8. Which is not to say that no work was being done that we could call 'linguistic' or 'semantic', for there certainly was. There is a whole tradition of thinking about grammar that goes back to Hellenistic times, at least, and within logic, there are penetrating analyses of the functions of expressions in, e.g., the Stoic school and in medieval scholastic thinking. The point is that in many cases the analyses developed are subservient to different goals, and that both the results and the ways in which these are argued for, are rather different from how the issues are approached in modern times. But, of course, that does not mean

seminal work of Humboldt and others, and almost another one for language to become a separate and central topic in philosophy.

Nevertheless, a case can be made that this is the stage that most closely resembles a ‘common sense theory of meaning’. From a common sense perspective it seems plausible that the meaning of a declarative sentence and the judgement that it serves to express are the same.⁹ No strict separation between ‘propositional content’ and ‘illocutionary force’ seems to be called for. Also, what the sentence means, the judgement it expresses, and what in reality justifies that judgement seem to be not really distinguished: how language relates to reality, the question that haunts much of the later philosophical thinking, thus never comes into proper focus. The way in which we form judgements about reality, — be it either in empiristic fashion, by receiving impressions through the senses and manipulating them, or more rationalistically, with more of the content being innate to the mind — is the way in which language ‘hooks up’ with it, except that it needs no hooks, since the relation is immediate.

The second stage in the development towards a more independent and more abstract conception of linguistic information is characterised by the rise of ‘meaning proper’ in the wake of the development of modern logic, mainly through the work of Frege, Russell, and early Wittgenstein. One of the hallmarks of Frege’s philosophy of logic is his anti-psychologism: in order to give logic its proper due, he claims, we need to separate it from ‘psychology’, i.e., we need to distinguish the subject of logic, viz., the systematic explication of the validity of inference, from the empirical study of actual judgements and actual reasoning. In his logical theorising Frege developed his position gradually. In the *Begriffsschrift* Frege (1879) he distinguishes between judgement and content, noting that the content of, e.g., an hypothetical judgement cannot be expressed in terms of the judgement of the antecedent and that of the consequent, but has to be defined in terms of their respective contents. However, he does still formulate his logic using a separate sign, the ‘judgement stroke’, for the judgement as such. Later on, he states that only the content of a judgement, but not the actual act of judging that content as true (or false, or plausible, or . . .), plays a role in the normative theory of valid deductive reasoning.¹⁰ As Frege states in ‘Der Gedanke’, which dates from 1918: logic is concerned with the ‘laws of truth’, and these laws can be regarded also as the ‘laws of thought’, but *not* in the sense of laws covering ‘general features of thinking as a mental occurrence’.¹¹ His argument, characteristically concise, is that ‘error and superstition have causes just as much as correct cognition’ and a study of actual thought would need to treat them on a par with correct judgement and valid reasoning, which contravenes the true task of logic. Rather than being a description of how we actually think and

that no interesting insights were developed along the way. Cf., Robins (1990) for an overview. Seuren (1998) is an example of an approach that is not purely historical, but attempts to connect the development of linguistics with modern systematic theories.

9. Which is not to say that it can not be, and has not been, challenged. Cf., e.g., (Dummett, 2004, page 1) for a dissenting opinion from an anti-realistic point of view.

10. A point concisely expressed by Wittgenstein in the *Tractatus*, where, referring to the *Begriffsschrift* he remarks parenthetically in 4.412: (‘Frege’s “judgement stroke” “|—” is logically quite meaningless’).

11. Cf., Frege (1918–19); quotations are taken from the English translation by Peter Geach in Frege (1977).

reason, logic is a normative theory that states how we should. Similarly, where in the early *Begriffsschrift* Frege uses the term ‘Vorstellungsinhalt’ (lit., ‘content of imagination’) to refer to contents of judgements, he later acknowledges that this term may lead to confusion since it is (also) used in a psychological sense, and instead settles on the term ‘Gedanke’ (‘thought’), which is supposed not to carry such connotations.¹² No doubt also inspired by Frege, Wittgenstein claimed in the *Tractatus* that ‘psychology is no more closely related to philosophy than any other natural science’, immediately following up with the claim that epistemology is ‘the philosophy of psychology’.¹³ Thus the idea that it is possible to treat language and meaning separately from questions regarding judgement and justification is gaining ground, and with that, the contours of modern philosophy of language become visible.

The separation from epistemology did not carry with it a similar move away from metaphysical concerns: the analysis of language and meaning remained strongly related to ontology. In part this is due to the particular philosophical aims to which people at the time made the analysis of meaning subservient. Stimulated by the success of the use of formal languages in the ‘new logic’, the age-old quest for a philosophically transparent (‘ideal’) language gained new momentum. This time it would be a strictly formal one, and it would provide philosophers with an analytic tool that could be used with scientific precision and mathematical rigour. This ‘linguistic turn’ put language centre stage in philosophy, and consequently turned philosophy of language into a distinct and central discipline.¹⁴ This is not the place to trace what happened to the idea of linguistic analysis as a philosophical tool employed outside the analysis of meaning proper.¹⁵ What is relevant here is how it influenced subsequent theories about natural language information and natural language meaning. And from that perspective it is important to briefly point out three general characteristics that have been very influential, viz., ‘universalism’, ‘intensional referentialism’, and ‘compositionality’.

‘Universalism’ refers to the nature of the task that philosophical analysis sets itself, viz., to give an account of ‘how language operates’ in general, with no reference to any specific features of any specific language in particular. What is of interest is not the way a certain language works, but what underlies the possibility of any language to express meaning. A straightforward feature, perhaps, of any philosophical analysis worth its salt, but one that will turn out to have repercussions for the form and the application of theories that are subsequently based on this idea. For in the application to concrete,

12. Cf., Frege’s comment from 1910 to Jourdain, who had written a summary of the *Begriffsschrift* in a paper on the history of logic: ‘For this word I now simply say ‘Gedanke’. The word ‘Vorstellungsinhalt’ is used now in a psychological, now in a logical sense. Since this creates obscurities, I think it is best not to use this word at all in logic.’ (Frege, 1879, page 11).

13. (Wittgenstein, 1960, 4.1121). Cf., also Wittgenstein’s attempt to give an extensional analysis of so-called ‘propositional attitude’ statements in 5.541 ff.

14. It is interesting to note that in the history of phenomenology, associated with the work of Husserl, Heidegger, Ricoeur, Merleau-Ponty and others, a similar development took place, but without the strict separation from epistemology that is characteristic for analytic philosophy. Cf., Dummett (1996) for more details.

15. There are a large number of studies dealing with this topic; cf., Biletzki & Matar (1998), Soames (2003).

empirical cases the universalistic and a prioristic features of these theories do not simply disappear. In many cases they become consolidated in the use of certain formal tools and in the adherence to particular basic methodological principles that are applied ‘across the board’ and that are even taken for granted as defining characteristics of the enterprise.

‘Intensional referentialism’ indicates the central role of the notions of reference and truth in the analysis of meaning, combined with the use of an intensional ontology consisting of possible situations and the entities of which such situations consist. Together these two assumptions, or requirements, tend to favour a fairly abstract notion of meaning, one that is grounded in the possibility of words having referential relations to objects, properties and relations in the world, where the relation of reference is understood as a truly intensional concept, not in any way restricted to reality as we know it: ‘the world’ can be any one from a set of logically possible ones.

Meanings of complex expressions, including sentences, are then assumed to be somehow constructed from these basic referential relations, which means that compositionality is assigned a key role.¹⁶ The result is an approach to meaning that is detached from actual reality and actual language use, one that works in a bottom up fashion, constructing complex meanings from basic ones, and that assigns the basic meanings a fairly independent status: they are ‘self-sufficient’ in so far as they have determinate and thoroughly non-contextual identity conditions.

The third stage that we need to distinguish in this brief historic sketch is that in which semantics arises as a separate discipline. This happened in the late 1960’s, early 1970’s, when developments in philosophy, logic and linguistics came together and gave rise to the idea that a formal theory of meaning can be developed and applied in the description of actual natural languages. This is the time in which people like Donald Davidson, Richard Montague, David Lewis, Max Cresswell, and a great many others did their seminal work.¹⁷ This time is the heyday of ‘Montague grammar’ (and its various variations and rivals) as a grand unifying framework, in which the conception of meaning that was developed mainly from a philosophical perspective at an earlier stage, was formalised using various logical techniques (borrowed from model theory, modal logic, type theory, tense logic, etc.), and applied in the description of natural languages. This approach to natural language semantics, aptly dubbed ‘formal semantics’, proved very successful and was the dominant one for quite some time, in particular in

16. Compositionality extends the expressive power of a language — the range of different meanings it is able to express — beyond that of its simplest expressions (its ‘words’). How far it does, depends on the kind of compositionality that the language allows. It is commonly assumed that most (and presumably all) human languages display a kind of compositionality that is genuinely recursive and that permits the construction of infinitely many expressions of unbounded complexity from a finite vocabulary. This sets human languages, as well as many formal languages, such as that of the predicate calculus, apart from simple signalling systems, in which each of a certain finite set of signs corresponds to one state of the system’s fixed application domain (like, say, the set of traffic signs of the traffic code of a given country), and also from language-like systems with limited forms of compositionality, such as the ‘language’ of the bee-dance or the languages used by chimpanzees who have acquired the ability to produce the sign for ‘green banana’ on the basis of having separately learnt the sign for ‘banana’ and that for ‘green’.

17. Cf., Davidson (1967); Montague (1973); Lewis (1970); Cresswell (1973). Other seminal work was done by Barbara Partee (1973). Though less directed to natural language the work done by Jaakko Hintikka, David Kaplan and Saul Kripke in that period was also of fundamental importance.

philosophy, less so in linguistics at large.¹⁸

One thing that is important from the perspective of this chapter is that through the extensive use of formal languages as tools for modelling natural language meaning yet another shift in that concept occurs: meanings now are first and foremost formal constructs, and theories of meaning are primarily differentiated in terms of the formal machinery one deems necessary for the description of semantic features of natural languages:¹⁹ concerns with epistemology or ontology become less and less important as semantics becomes more and more autonomous, and the nature of the concept of meaning reflects this. Montague's claim, in 'Universal Grammar' Montague (1970b), that 'there is in my opinion no important theoretical difference between natural languages and the formal languages of logicians' and that therefore 'it [is] possible to comprehend the syntax and semantics of both kinds of languages within a single natural and mathematically precise theory' testifies to this shift. The consequences are far-reaching. For one thing, although Montague seems to think of logic and semantics as some kind of 'equal partners', the practice is less symmetrical: it is formal languages that are used as models for natural languages, and this implies a sharpened focus on those aspects of meaning that can indeed be dealt with using existing logical techniques, and a proportionate neglect of those that can't. The distinction between 'structural semantics' and 'lexical semantics', arguably one that is not in any sense inherent in meaning itself but rather an artifact of the kind of instruments one wants to use, is maximally exploited and the resulting concept of meaning becomes both more formal and more 'thin'.

At this third stage the three features identified above are very much present, although not always explicitly so. Looking at the abundance of descriptions of various semantic phenomena in a wide variety of languages produced in the 1970's and 1980's, one might think that 'universalism', the idea that a proper semantic theory deals with natural language semantics as such, isn't something that people subscribed to. And indeed, the very fact that semanticists deal with actual phenomena, some of which are specific to a particular language, indicates that their concern is not that of the philosophers at an earlier stage. Nevertheless, the use of a unified framework has universalistic consequences, whether intended or not. The point is that the framework itself embodies assumptions about what meanings are, how they are related to each other, how they are expressed, and so on. So right in the framework itself there is a conceptual structure, explicated by means of the formal properties of the concepts and languages that are used, that shapes a concept of natural language meaning that is independent of any concrete manifestation in any concrete natural language.²⁰

18. In particular within the Chomskyan tradition people tended to reject the use of model-theoretic techniques, and pursued a different approach, that is more in line with Chomsky's idea that linguistics is a branch of cognitive psychology, and, ultimately, of biology. Cf., further below.

19. One could say that as a result of this shift semantics deals with an altogether different type of phenomena. Although this may seem exaggerated — and it probably is — it does point to a curious and slightly worrisome fact, viz., that there seems to be no theory-independent agreement about what exactly the domain of semantics consists of. This is reinforced when one takes a closer look, e.g., at the kind of arguments that formal semanticists and Chomskyan semanticists bring to bear on their dispute. Cf., Stokhof (2002) for some more discussion.

20. This is particularly clear in the work of Donald Davidson, who actually uses the logical structure

The other two features, pertaining to the central role of reference and truth and the use of an intensional framework, and to compositionality as the basic principle for dealing with semantic complexity and creativity, are less hidden and more explicitly adhered to. Despite discussion about the kinds and the number of intensional concepts that one needs to employ, the common denominator is the use of a formal framework that models ‘the world’, —i.e., that to which the expressions of the language bear a referential relation and in terms of which the concept of truth for the language is defined,— in an abstract, and, one might almost be tempted to say, ‘detached’ way. ‘The world’ is reduced to the bare minimum of components and structure that is needed to define what kinds of things the referents of various types of basic expressions are, compositionality being understood to take care of the rest. It is important to note that it is not actual reference that is defined or explicated, it is only the formal type of relationship involved that is being accounted for.

The resulting picture, which for a long time served as the classical model for semantics of natural language and which we will refer to as such in what follows, in many ways comes close to that of a natural language as a formal language, —significantly, a formal language without a concrete application. It portrays natural languages as information carrying devices in the fairly abstract sense in which the same characterisation can be given of many other information carrying systems, ranging from signalling systems to mathematical notation. But, as was already indicated in the introductory section, if we look more closely at the various ways in which natural languages convey information, at what kind of information that is and what it is about, we encounter a much richer structure, and one that is tied more closely to the actual world that we live and use our language in than is accounted for in this approach. Small wonder, then, that after its initial success and broad acceptance the classical model became gradually discredited. At first one tried to augment it with additions that put more semantic flesh on its formal bones; later it was supplanted altogether by approaches in which the flesh is taken as seriously as the bones.

2.2 Reinstating content

The ‘counter current’ that contributed to a much more balanced picture of the specific characteristics of how natural languages act as information carrying devices does not represent one, homogeneous conception of meaning, rather it springs from a number of sources. These do have one thing in common, though, which is a profound dissatisfaction with the conception of linguistic meaning that informs the formal semantics of the 1970’s. Different people addressed different aspects of that dissatisfaction; together they effected a shift in the orientation of natural language semantics that is still taking place today. Again, we should note that this development primarily is a reaction to a ‘self-styled’ history, which only partly covers what actually occurred in philosophical

of a semantic theory, which according to him takes the shape of a Tarski-style theory of truth, in a transcendental argument against ‘the very idea of a conceptual scheme’, arguing that because the semantics of any language can only be described by means of such a theory and because the very framework of that theory implicates substantial properties of the meanings expressed in the language, all languages are essentially translatable into each other Davidson (1974).

thinking about language and meaning. Obviously, there is the work of a number of authors who already early on explored different directions that implicitly challenged some of the basic assumptions of the classical model, e.g., the ‘linguistic phenomenology of J.L. Austin, H.P. Grice’s work on meaning and intention, and the work on speech acts of John Searle’²¹, much of which was inspired by Wittgenstein’s later work.²² But this work only became influential after formal semantics had gone through an autonomous development, and even then it was taken up not in semantics proper, but mainly in a theory of pragmatics, which was supposed to complement it.

The conception of meaning that people reacted against can be dubbed ‘classical descriptivism’. Central to this conception is the essentially Fregean principle that the meaning of an expression determines its reference by providing a specification of the conditions that something needs to satisfy in order to count as being the referent. The Fregean concept of ‘Sinn’ is explicated formally by reconstructing it as a function that takes a possible world (or other such intensional construct) as its argument and delivers an entity (an individual, set of individuals, or set of n -tuples of individuals, as the case may be) that acts as the referent in that world. In line with the above-mentioned distinction between structural and lexical semantics, the actual specification of these functions for concrete expressions was by and large considered not to belong to the subject matter of semantics. Instead one focused on the various ways in which these largely unspecified functions can be combined to form appropriate meanings for larger expressions, in particular sentences, yet another illustration of the pivotal role of compositionality.

By thus reducing both the ‘world’ (that which natural languages are about) and the meanings of particular words and phrases to formal structures many questions were bracketed out that both linguists and philosophers would consider it their task to answer: questions as to how concrete expressions actually refer to concrete objects or properties, how such referential relations arise, what role contingent features of the way the world is have to play in that process, how considerations regarding the communicative functions of natural language utterances might interfere, how the use of language interacts with other cognitive functions, how utterances employ features of the linguistic and non-linguistic context in which they are produced, and a host of others. It is to the neglect of such questions that people reacted and which motivated them to develop alternative approaches.

Consequently, we can, admittedly somewhat arbitrarily, identify four separate sources of this counter current, one that is concerned with the role of the world, another that focuses on the variety of communicative uses, a third that insists on taking indexicality and the linguistic context seriously, and a fourth that investigates the cognitive status of language and its relations to other cognitive structures and functions. Of course, these divisions are to some extent artificial, but they serve to indicate major trends.

The first source of dissatisfaction with classical descriptivism relates to the minimal role that it assigns to the world and our interactions with it. One central question

21. Cf., e.g., Austin (1962), Grice (1957), Searle (1969).

22. Primarily his *Philosophical Investigations* Wittgenstein (1958).

here is how linguistic meaning comes about, a question that actually reinstates the connection with traditional, basic epistemological concerns. And, as in the tradition, there are two main points of view, an internalistic and an externalistic one.²³ The central claim of semantic externalism is that meaning derives from the world, at least substantially.²⁴ It is from our environment and our interactions with it that natural language expressions get their meanings, and to a large extent the processes involved are of a causal nature. Hence this view is often also referred to as a ‘causal theory of reference’.

According to this externalistic view natural language meanings can, to a large extent at least, be naturalised: the contents of many natural language expressions can be identified with real situations, events, objects, properties and relations — entities belonging to the external world but with which the language user can interact through perception and action. The most explicit realisation of this viewpoint within formal semantic theorising, and the most systematic formal attempt to restore the relationship between meanings and their naturalistic determinants,²⁵ are situation semantics and its logical-philosophical foundation, situation theory.²⁶

Taken to its extreme, radical externalism involves naturalising all aspects of meaning. One reason why someone might think that such a radically externalistic account of linguistic meaning ought to be possible is that, arguably, all our concepts are ultimately the product of our interactions with the world in which we live, and thus are, in some fashion, reflections of the ways in which that world imposes itself upon us in the course of those interactions. But this consideration overlooks the fact that even where experience of the world is causally involved in the construction of the kind of information that linguistic expressions convey, this information cannot be equated with that experience.²⁷ There is no a priori reason to suppose that the world, our experience of it, and how we conceptualise and express it in natural language have the same fine structure. Rather, there are a number of good reasons to doubt that this is the case. For one thing, there is the moulding role that our cognitive system may exert on the form and structure of the experience. And many of our expressions refer to things in the world that exist at least partly because of our shared language and the way we use it. In fact, for all we know, linguistic representation makes its own contributions to the ontology of natural languages, which includes entities the reality of which is confined to aspects

23. What is called ‘internalism’ and ‘externalism’ here, in the context of semantics, should not be confused with the ‘internalism – externalism’ debate in the philosophy of mind and in epistemology, although there are connections, of course. In the philosophy of mind internalism and externalism are rival views on the nature of mental content, centring around the question whether mental content can be completely understood in terms of internal mental representations, and, ultimately, perhaps entirely in terms of brain states..

24. This is a crude generalisation, of course. There are many, often subtle variations on this theme that are lumped together here under the one heading ‘externalism’. Cf., McGinn (1989) for an overview. The locus classicus of semantic externalism is Putnam’s ‘The meaning of “meaning”’ Putnam (1975), which is also one of the classic sources of the theory of direct reference to which Kripke, Donnellan and others contributed.

25. And thereby also traditional connections between philosophy of language, epistemology, and psychology of a particular bend, viz., naturalistic and empiricist psychology,

26. Cf., Barwise & Perry (1983); Barwise & Seligman (1997).

27. Cf., also Dretske’s analysis of the concept of information in epistemology, in this volume.

of the ‘world’ that our language projects, and which have no right of being in any language-independent sense. So it seems that although experience allows information, — in the sense of enabling it by anchoring the terms of our language in an external world, thereby creating the possibility of objective reference and truth — , it does not determine it completely: in general, the information that is conveyed by means of natural language is the product of more factors than experience alone. And that entails that a complete naturalising of meaning is not possible.

The next question then is what else might be needed for meaning. Several answers are possible, one of which is provided by internalism. Like externalism, internalism aims to enrich the meaning content of expressions. But it does so via a different route, viz., through an appeal to substantial contents and structures that are supposed to be resident in the human mind. From the internalistic perspective the mind contains a rich repertoire of basic contents, in the form of innate concepts and features and of structural operations, that together allow for the formation of the huge variety of actual meaning contents that we find expressed in natural languages. As such, internalism naturally allies with the equally rationalistic conception of grammar that is so characteristic for the Chomskyan paradigm in linguistics. Nevertheless, internalism, too, faces some difficult questions, some of which are conceptual: ‘What explains the origins of all these mind contents?’, ‘How can we account for the application of mental content to reality?’, others empirical: ‘What is the explanation of semantic variety across languages?’.²⁸

Also, it should be noted that externalism (usually) and internalism (by definition) are individualistic: they take the individual human being as their point of reference when discussing linguistic and mental content and its relation to the world. This is in accordance with much of the main stream thinking in philosophy of language, philosophy of mind, semantics, and linguistics. Again, a reflection of this is the central role that is played by compositionality. From an individualistic perspective what is often called the ‘creativity’ of language, viz., the potential infinity of structures and meanings that together make up a language, poses a serious problem. How can individuals, being finite creatures with finite memory and finite computational resources, be considered competent users of their language? Compositionality comes to the rescue: it not only characterises languages conceived as formal objects, but is also posited as an inherent feature of human linguistic competence.²⁹

Nevertheless, there remains the interesting question whether the individualism that characterises both externalism and internalism makes these accounts too restrictive. Internalism seems to have a hard time accounting for the availability and contents of concepts that rely on the existence of social institutions, and faces serious problems when dealing with phenomena such as distributed information and reliance on expert knowledge.³⁰ That we could locate the concepts involved in such phenomena exclu-

28. Cf., Farkas (2006) for a recent overview of externalistic and internalistic perspectives in the philosophy of language.

29. Cf. Groenendijk & Stokhof (2005) for some discussion about how these various elements are usually linked up, and for some discussion of possible alternative ways of accounting for competence.

30. Which is one of the central points in Putnam’s original 1975 paper (cf., footnote 24). For attempts to account for such issues in terms of the distinction between ‘broad’ , externally determined content and

sively ‘in the mind’ seems improbable. For the externalistic perspective individualism becomes problematic when it is robustly physicalistic. A lot of mental content and linguistic meaning seems to defy a straightforward reduction to physicalistic causes. Note that the problem here is not one for physicalism as a doctrine concerning the nature of scientific explanation. Whether or not that is a tenable position does not depend on the possibility of giving a physicalistic account of all of linguistic meaning, for one could argue that some such meanings simply have no role to play in an ultimate scientific account of the world. But from a semantic point of view this is different, since we obviously want a semantic theory to account for all linguistic meaning, including the meanings of those parts of the language that certain views on scientific explanation would consider irrelevant to their concerns. This does not rule out externalism *per se*, but it does indicate that an externalistic account of natural language meaning needs to take into account that whatever causal relations are involved in producing it, they are not monostratal and uniform, but rather play at different levels and are of different types; that they involve radically different kinds of entities, including various sorts of social entities; and that they work in both directions, from the world to meaning and vice versa. The structure of meaning is partly due to the structure of the world, but the structure of our world is also partly a linguistic one.

Such an approach transcends the structure of radical externalism as we characterised it above. In particular, the causal processes which it would take into account are not simply just the ones that govern the perceptions of individual speakers. This applies specifically to those processes that are needed to account for the meanings of social terms, among them those that pertain to the interactions between verbally communicating speakers of a given language.³¹ One effect of the impact of these additional causal relations, which connect the members of a given (speech) community rather than any one of them to a particular content, is that these linguistic meanings aren’t the private property of individual speakers, but rather a shared possession of the language community as a whole. For such expressions the ultimate linguistic competence rests with the community, and the competence of any particular member of that community is determined by the degree to which he partakes in that common good. Such a move away from mainstream individualism could also account for the real diversity of experience and the diversity of information, not necessarily parallel to the first, that we find across individual language users. Viewed from the perspective of a community, experience is heterogeneous, but connected, and the same holds for information. It is precisely this diversity that is one of the main reasons why humans use such complicated, expressive languages as they do.

The last observation is connected with the second source of the counter current to the classical model, which is a concern for the complexity and the variety of the communicative uses that are made of natural languages. In the introduction we hinted at this by giving some simple examples of other uses than the straightforwardly declarative use. Quite in line with its ancestry in logic and philosophical analysis the classical

‘narrow’, internal and individual content, cf., Fodor (1987).

31. This could also be called a form of externalism, viz., ‘social externalism’. Cf., e.g., Burge (1990).

model focuses on declarative utterances. Actually, just as the ‘judging’ element from the traditional notion of a judgement was first isolated and then dropped by Frege, leaving only the contents of judgements as the material to which logic was supposed to apply, analogously looking just at declarative utterances made it easy to first isolate the ‘use’ part of an utterance and then focus exclusively on the resulting content, turning formal semantics into a theory of pure contents, radically dissociated from the various ways in which these can be used. Such a separation between what is often called ‘mood’ (or ‘illocutionary force’) and ‘radical’ (i.e., propositional content) goes back to Frege and was taken up later in various forms by people like Austin, Stenius, and Searle.³² The resulting speech act theory made this distinction into one of its basic principles. Thus a division of labour arose between formal semantics as an account of propositional content and speech act theory, or pragmatics in a wider sense, as a theory of the use that is made of these contents.

However, some have questioned whether this strategy will work. For one thing the variety of uses we make of natural language expressions does not seem to be one-to-one related to the mood-radical distinctions we can make within these expressions, be it on the basis of syntactic form (interrogative, indicative, . . .), the presence of lexical items, or a combination thereof. And then there are aspects of meaning, i.e., information conveyed through a speaker’s utterance to other interlocutors, that are not in any obvious way coded into the expressions uttered, but that arise from the interplay between the context in which the utterance occurs, the intentions and expectations of the various speech participants, and other meaning elements. These ‘implicatures’, as they are called, have been studied extensively; and they have given rise to serious questions about the tenability of the classical model. Like in the case of speech acts, the initial approach towards an account of such indirectly conveyed meaning depended on a division of labour, in this case between semantics as conceived in the classical model and a pragmatic theory called the ‘logic of conversation’, developed by H.P. Grice.³³ Grice’s central idea was that language use is a cooperative task and that therefore language users can be expected to obey certain rational principles of communication, such as telling the truth (as they see it), giving sufficient but no superfluous information, and so on.

One problem with Grice’s original approach concerns one of its starting points: one of Grice’s main motivations was to show that certain aspects of the meaning of natural language connectives that are not captured by their extensional two-valued counterparts (for example, the order sensitivity of natural language conjunction) can be accounted for by an appeal to cooperative principles. A closer look at the apparent meanings of ‘and’ co-ordinations in English (the same also applies to other languages) reveals that their meaning depends on factors that go beyond the mere truth table of classical logical conjunction and are also different from the conversational principles Grice invokes. In particular, the order-sensitivity of ‘and’ co-ordinations is largely the effect of the mechanisms of interclausal temporal anaphora, mechanisms that are operative also where no ‘and’ is in sight, and that any theory of natural language meaning

32. Cf., Austin (1962); Stenius (1967); Searle (1969).

33. Cf., Grice (1975); Levinson (1983) is an excellent introduction to this and related subjects.

and information will have to account for in any case.

What goes for ‘and’ goes for most applications to which Gricean conversation theory has been put: The principles of the theory are important and indispensable, but so are other principles, which also transcend the restricted conception of meaning that is part of the classical model. And again and again it has been found that deciding which of these principles should be counted as semantic and which as pragmatic is possible only on theory-internal grounds.³⁴ This has led to the view that the demarcation between semantic and extra-semantic (= pragmatic) aspects of meaning is to a considerable extent arbitrary, and has thereby undermined another fundamental assumption of the classical model.

What has thus emerged in lieu of the classical model is a far more complex account in which a great variety of principles and mechanisms collaborate in the construction of utterance meanings out of the meanings of the words contained in them. Some have taken the reasoning that has led to this line of development one step further and argued that even the concept of ‘literal meaning’ that the classical model, speech act theory and Gricean pragmatics all rely on is a myth. In a theory of literal and non-literal meaning the words of the language have literal meanings, which are encoded in the lexicon. These serve as a starting point for the derivation, via inferential processes that take various pragmatic factors into account, of other, non-literal meanings, and, on the basis of these, of the specifications of individual utterance contents. But here too, it is argued, we are dealing with a distinction — that between the literal meanings of words and their non-literal meanings — which proves to be slippery and hard to draw except on theory-internal grounds. One major empirical problem is the existence of (productive) polysemy. The assumption of literal meaning forces one to try to account for the various meanings of, e.g., ‘running’ as it occurs in ‘The tap is running’, ‘John is running’, ‘The program is running’, ‘My nose is running’, etc., by picking one meaning as the core, or ‘literal’ one and then accounting for the others on the basis of some contextual derivational process. A more plausible alternative is to forgo the choice and account for this type of variability by making lexical meanings themselves contextual and flexible, in effect viewing linguistic meaning as something that is the result of interaction between a language user and his environment.³⁵

Emphasis on interaction with the environment, especially the communicative environment, consisting of other speech participants, conversational goals, information about the world (individual and shared), and so on, is characteristic for the third source of the counter current, the one that focuses on context in this broad sense. An important shift in the way meaning is viewed that is characteristic for this development is the result of a turn away from the exclusively descriptive orientation, with its emphasis on the language – world relation, that is a central feature of the classical model, to a perspective on language and language use that analyses them primarily in terms of information and information exchange.³⁶ The resulting view is one in which the primary

34. Cf., the discussions in Recanati (2004), van Rooij (2004b), Stanley (2005), and the contributions in Szabó (2005).

35. Cf., Bartsch (1996) for more discussion and a concrete model along these lines.

36. Stalnaker’s work on presupposition and assertion is an early representative of this conceptual turn. Cf.,

focus is on the ‘horizontal’ relation between language users engaged in an information exchange discourse, with the ‘vertical’ relation of language to world entering only indirectly, and no longer playing the lead role. The information exchanged in a discourse can be quite diverse: usually, part of it will be information about the world, but at least as important is information of speech participants about each other, and information about the discourse itself. When engaging in a conversation, but also when reading a text or listening to a speech, what the participants know, or think they know, about each other plays a crucial role in interpretation, and, derivatively, also in production. (A speaker will choose the expression she utters so that it will lead, to the best of her knowledge, her audience to assign to it the interpretation she intends, given the total package of information, about world, antecedent discourse and her own state of mind, that she assumes is available to that audience.) Stalnaker’s notion of ‘common ground’, i.e., the information that the speech participants assume they share, is an important element in this, since it provides them with common resources for picking out individuals, properties and situations, solving (co)referential relationships, and so on. But the common ground will normally also include information of all the different kinds we have mentioned, not just assumptions that directly concern the topic of conversation.

In addition to what is being said by whom to whom, i.e., content in the narrow sense, it is also form that matters for determining what information gets exchanged. Among the natural language devices that serve this purpose we find: anaphoric expressions of various kinds, among them pronouns, tenses, and certain temporal and spatial adverbs, which permit resumption of entities previously introduced into the discourse; presupposition-inducing expressions, that enrich and structure the common ground; the order in which the events that make up a narrated episode are described, which usually indicates the temporal ordering of those events; and so on. These and other devices help the hearer to relate the information conveyed by an utterance to the information he already has, and thus to identify exactly what the new information is. As such they are an integral part of what linguistic meaning is and how linguistic expressions convey information. At yet another level, not so much concerned with linguistic form or narrow content, there is information about the aims with which speech participants have entered the conversation, their rhetorical strategies, and other features of their linguistic personae. This type of information is crucial for the detection of irony or sarcasm, the appreciation of a verbal sleight of hand or a clever play on words, and for the recognition of an implicit reproach or a concealed request. These aspects of discourse, too, are factors that enter into the way in which natural language utterances play their information conveying role.

These considerations have given rise to a variety of alternatives to the classical model. In as much as all these models share the shift from the descriptive to the information exchange perspective, along with a shift from the sentential to the discourse level, they can be captured under a common denominator, that of ‘dynamic theories of meaning’.³⁷ These theories take the development outlined above one step further

the two seminal papers Stalnaker (1974) and Stalnaker (1979).

37. Thus the original model of discourse representation theory developed by Kamp in Kamp (1981) (cf.,

and change the notion of meaning itself: the descriptive and referential and hence truth oriented perspective of the classical model is replaced by a dynamic one that views the meaning of expressions in terms of what is called their ‘context change potential’, with information being one, central aspect of the context. This further shifts, or rather blurs the original distinction between semantics and pragmatics (i.e., the distinction between what is supposed to be a matter of meaning proper and what belongs to the realm of use). Accordingly the focus of research in these theories is no longer on the referential and logical features of linguistic meaning but on issues involving information structure (topic – focus, presupposition, anaphoric relations, intonation and prosody) as linguistic devices that can be used to link a new sentence in a text or a new utterance in a conversation to what went before, or to prepare the ground for what comes next. This increasing focus on information exchange and information structure also weakens the link with ontology that in the classical model was secured through the central role of reference and truth. In a dynamic perspective truth becomes a mere limit concept of the more general notion of acceptance by the speech participants of information that is being exchanged.³⁸

Integral to the dynamic view on meaning as context change potential is a renewed interest in the cognitive function of meaning. This ties in with the fourth source of the counter current that we discerned above, viz., a renewed interest in the cognitive aspects of language and its relations to other cognitive systems. The development of the classical model in the 1970’s brought along a new and somewhat problematic relationship with psychology. On the one hand its proponents, sometimes explicitly, more often implicitly, took over Frege’s anti-psychologism, that made a principled distinction between logic as a normative science and the empirical study of actual reasoning, and they applied it to the study of natural language meaning, separating formal description of semantic structure from the study of the way in which language is produced and interpreted. But unlike logic, semantics never really was conceived as a purely formal discipline; after all, its aim is to describe and explain empirical facts, and it is therefore considered to be as much a branch of empirical linguistics as phonology or syntax.³⁹

From that perspective the classical model should have been quite compatible with the Chomskyan approach to grammar. But in fact the relationship turned out to be more complicated. For one thing, the Chomskyan model involved a close alliance with rationalistic thought and with the computational approach in cognitive psychology that developed from the 1960’s onwards. But not everybody felt comfortable with these particular philosophical presuppositions, and many semanticists working within the classical model preferred to keep their distance. In turn, many Chomskyans, in-

also Kamp & Reyle (1993)), explicitly aims to combine a declarative and a procedural view on natural language meaning. Other models of dynamic semantics include Heim’s file change semantics Heim (1983), Veltman’s update semantics Veltman (1996), and Groenendijk and Stokhof’s dynamic semantics Groenendijk & Stokhof (1990, 1991), cf., also Groenendijk et al. (1996).

38. This does not necessarily imply that the concept of truth is a completely epistemic notion. That depends on how states of complete information relate to states of the world. In fact, in the theories mentioned in footnote 37 the notion of truth is as objective as it is in formal theories that implement the classical conception.

39. Cf., Stokhof (2002) for some more discussion of this tension.

cluding Chomsky himself,⁴⁰ kept formal semantics at bay, arguing that the use of the concepts of truth and reference as central tools in the explication of meaning disqualified the classical model as far too externalistic to be compatible with the internalistic approach that they favoured. Semantics in the Chomskyan framework accordingly concentrated primarily on the way in which conceptual structure is expressed, mainly in lexical semantics.⁴¹ In this connection the approach of ‘cognitive semantics’⁴² should be mentioned as well. Though in many ways adverse to the generative framework as developed by Chomsky in his later work, it shares with that approach a focus on lexical semantics and an unwillingness to account for meaning in terms of reference, using the tools of logic in the way exemplified by the formal implementations of the classical conception, in particular in model-theoretic semantics. Characteristic for cognitive semantics is the emphasis on the fluidity of the distinction between semantic knowledge and encyclopedic knowledge and on the embodied nature of meaning.

With its focus on formal properties of natural language meanings, the classical model initially succeeded in maintaining something of a ‘splendid isolation’ from empirical work in psychology and biology. But as the counter current grew stronger, as more aspects of use were taken into account, context became more and more important, users and their conversational goals and strategies were incorporated as significant aspects of the context and as the emphasis accordingly shifted from formal structure to actual content and its use, these barriers began to crumble. For many it has become increasingly obvious that one of the tasks of semantics is a realistic modelling of language users and their interactions, for in the end natural language meaning can be properly understood only if we understand how it functions in real information exchanges and other linguistic interactions.

This has brought about a certain rapprochement between semantics and psychology, and to some extent also between formal semanticists and people working in the Chomskyan tradition. This rapprochement has also been helped by a growing interest in lexical semantics on the part of formal semanticists, who at long last have begun to respond to the charge that if all meanings are derived from lexical meanings, then explaining how they are derived is not good enough if one has nothing to say about what they are derived from. Nevertheless there remain substantial differences between the internalistic and the externalistic perspective (the former being preferred by those who take the Chomskyan approach). But as the computational model in cognitive psychology began to loose its grip, it became clear that the study of how language functions as one of the human cognitive faculties does not necessarily commit one to an internalistic view. There is room for a variety of perspectives, some working with a model

40. Thus early on, replying to a suggestion from Bar-Hillel that formal logic might contribute to linguistics, Chomsky stated that ‘the relevance of logical syntax and semantics [to the study of natural language] is very dubious’ Chomsky (1955). And throughout the years Chomsky expressed similar sentiments on a number of occasions. For example, in Chomsky (2005) he states, in keeping with his internalistic perspective, that ‘even the most elementary concepts of human language do not relate to mind-independent objects by means of some reference-like relation between symbols and identifiable physical features of the external world’.

41. Cf., Jackendoff (1990), Pustejovsky (1995).

42. Cf., Lakoff (1987), Talmy (2000).

that is individualistic and internalistic,⁴³ others favouring a more externalistic set up that emphasises the role of the linguistic community.⁴⁴

The rapid development of new techniques for studying brain processes and the consequent rise of cognitive neuroscience during the last decade has greatly contributed to a renewed interest in the underlying mechanisms of meaning. Language being one of the core cognitive functions of humans, it has always been an important object of study in cognitive psychology, as witnessed by a long tradition of studies in language acquisition, language pathologies, and language processing. For a long time such studies were generally based on computational, internalistic models of language, although some more empiristic and community oriented studies were undertaken as well.⁴⁵ The prospect of being able to study the brain almost ‘in vivo’ as it processes language, holds much promise. Particularly enticing is the possibility of experimentally testing different theoretical models that account for more or less the same linguistic data. The advent of more performance oriented models, such as dynamic semantics, optimality theory and game theoretical semantics have greatly facilitated this reorientation.⁴⁶ However, as our earlier discussions concerning externalism, internalism and individualism illustrate, we should be careful in our assessment of what exactly can be achieved in this fashion. The idea that research in cognitive neuroscience will be able to arbitrate between rival semantic frameworks all by itself is certainly not unproblematic: for one thing, the relationship between neurological correlates of semantic concepts and these concepts themselves cannot simply be taken for granted, and it seems that the relation between the two is much more symmetric than a reductionist approach would predict.⁴⁷ And the contributions of physical and social reality need to be taken into account as well.

Finally, it should be noted that the shift towards information exchange and other aspects of use that is embodied in these new approaches also has spurred a renewed interest in the biological and cultural origins of language, both phenotypically and genotypically.⁴⁸ Using techniques from evolutionary game theory and learning theory, semanticists have begun to study the way in which expressive systems can arise within a population of interacting agents, trying to isolate which factors are responsible for the characteristic features of human languages, notably recursive structure and semantic compositionality.⁴⁹

In conclusion, it seems fair to say that the current state of thinking in philosophy of language and natural language semantics about meaning is one of diversity. There seems to be no one, dominant conception of natural language meaning, and many,

43. Cf., the references given above.

44. Cf., Tomasello (2003).

45. Cf., work by Bruner and others Bruner (1983); Garton (1992), and early work in the connectionistic paradigm.

46. Cf., footnote 37 for references to work on dynamic semantics in relation to natural language; cf., van Eijck & Stokhof (2006) for a more general overview of various concepts from dynamic logic. For optimality theoretic semantics, cf., Hendriks & de Hoop (2001), for game theoretical approaches, cf., Hintikka (1983).

47. Cf., Baggio et al. (to appear) for an in-depth discussion.

48. Cf., Tomasello (1999), for an early, influential study.

49. Cf., Christiansen & Kirby (2003) for a collection of papers that gives an overview of current thinking about language evolution; for recursive structure and compositionality, cf., Kirby (2000).

sometimes quite divergent approaches to its analysis are being pursued concurrently. The resulting situation might strike some as somewhat paradoxical: on the one hand all these abstractions have led to success, yet what it is they purport to study, viz., natural language meaning, seems to fade from view, at least as one coherent, unifying concept.⁵⁰

Indeed, in some cases we appear to be dealing with incompatible underlying conceptions, as for example in the case of internalism and externalism. But more often it seems that differences arise because people focus on different aspects, and that, although it might not always look that way, the results could be unified in a single, more encompassing theory. The contours of such a theory are beginning to emerge, although no generally accepted format has been established as yet. It treats natural language meaning as a ‘thick’, i.e., substantial concept that gets its content and structure from a variety of sources (conversational goals, with a pivotal role for information exchange, the world, reflexive models of language users) and that ties in closely with other cognitive functions (perception, the emotional repertoire, everyday skills). Thus it reinstates the close relationship between meaning, information, judgement and the world that was characteristic for many of the earlier views on linguistic meaning that predate the classical model. But it does so based on a firm grasp of the underlying formal structure of the concepts involved, thus allowing for descriptions that have extended empirical scope and greater explanatory power.⁵¹

In the following sections we will illustrate a few important aspects of the present state of thinking about meaning and information in natural language by outlining in somewhat more detail the main elements of one particular way of describing and analysing how natural language expressions perform their information conveying roles. In section 3 we will discuss how the relational nature of linguistic meaning can be captured by means of representational techniques that are derived from model theory, allowing us to define the linguistic meaning of an expression in terms of the information carried by an utterance of it in various circumstances. The starting point of our exposition will be something akin to the classical model, which we will then subsequently modify and refine to capture more aspects of content and context. Next, section 4 will be devoted to an illustration of the way in which this particular conception can be used to capture how information is conveyed in larger units of linguistic material, such as texts and conversations. We will illustrate the way in which anaphoric relations and presuppositions establish discourse connections that enter into the specification of the informational content of utterances, and we will show how the model set up in section 3 can be enriched so that it accounts for this. In section 5 we analyse how the state of the recipient enters into the picture, again indicating how the model can be adapted

50. Which gives rise to difficult methodological questions as to what the nature of the success is: What is it that current semantics and philosophy of language are successful at? What are the measures of success here? Are these measures (relatively) theory independent? What do they apply to? And so on.

51. It should be noted that there is also a continuing tendency toward the use of notions of meaning and information that are at a greater distance from what we could call the qualitative, common sense notion, as witnessed by the rise of purely quantitative, statistical notions of information in combination with the use of ‘shallow’, non-rule based techniques in certain approaches in natural language processing, information retrieval, semantic web, and so on.

to account for this as well. In section 6 we come back to the question of what is characteristic of linguistic information. We conclude this chapter with a short overview of current further work in this area.

3 Modelling meaning in context

In the introduction to this chapter we observed that the notion of linguistic information is inseparable from that of linguistic meaning, that both are relational and that the richness of linguistic meaning is due in large part to the fact that the syntax and semantics of human languages involve recursion. In this section we discuss these issues in more detail.

First a few words on syntax. One of the oldest insights into language is that sentences have grammatical structure. For instance, the observation that the typical sentence of a language such as Latin, French, or English, contains a verb and that this verb has a subject can be found in the earliest grammars; and it is something that speakers of those languages will accept without demur when it is pointed out to them, and that they might find out without much trouble for themselves. It is also plain, and no doubt always was, that simple sentences can be used as building blocks for larger sentences, e.g., as conjuncts, or as relative clauses, or as subordinate clauses beginning with subordinate conjunctions such as ‘when’, ‘although’, or ‘because’. Speaking more generally, it was from the beginning a central aim of the ‘Art of Grammar’ to describe how grammatically correct sentences can be analysed into their grammatical constituents, as a way of proving that they are in accordance with what Grammar demands.

Modern generative grammar starts from a superficially different point of view, according to which sentences and other complex linguistic expressions are built from basic constituents (the words and morphemes of the language) according to rules that guarantee their grammaticality (or ‘syntactic well-formedness’, as terminology has it). And the way in which a grammatical expression is built from the words and morphemes occurring in it according to the rules of syntax shows its grammatical structure and is thus, once again, a demonstration of its grammatical correctness. What makes a generative grammar recursive is that some of its rules can be used repeatedly in the construction of a single sentence. More explicitly: the grammar is recursive if it has recursive rules — where a rule R is a recursive rule of a given grammar G if and only if for any number n there are sentences generated by G in which R is used at least n times. (In the generative grammars that have thus far been proposed for natural languages all or nearly all rules are recursive in this sense.)

In the end there is not much to choose between the generative and the analytical approach to grammar. In fact, on the basis of a generative grammar it is generally possible to construct parsers which compute syntactic analyses for those strings of words and morphemes that the grammar generates, by tracking how the string can be built using the grammar’s rules. So, when we proceed, as we do, from the assumption that grammaticality is defined in terms of generative grammars we do so without loss of generality.

The first formally explicit accounts of natural language meaning made use of generative grammars that fit our characterisation of such grammars perfectly in that they consisted exclusively of generative rules, which serve to build complex expressions out of simpler ones. The accounts assumed that for each such rule R that tells us how expressions $e_1, \dots, e_n$ can be combined into a complex expression e there is a corresponding semantic rule R' which states how the denotations $d_1, \dots, d_n$ of $e_1, \dots, e_n$ must be combined to obtain the denotation d of e .⁵² As a matter of fact, natural languages do not take well to the comparatively rigid regime that is imposed by generative grammars of this strict and simple generative form, and more complex rule systems are needed if their syntax is to be captured in intuitively plausible and theoretically convincing terms. But for our present purposes the way in which these more complex systems determine the meanings of complex expressions is the same as it is for the simpler generative grammars described above and the extra complications can safely be set aside.

We will therefore assume that the syntax of natural languages can be given as consisting of (i) a set of rules, determining how complex expressions can be built from smaller ones, and (ii) a lexicon, specifying words and morphemes.⁵³

All grammars make use of grammatical categories. This is true in particular of generative grammars: like other grammars they classify well-formed expressions into different categories. These categories are essential to generative grammars in as much as the rules refer to them. The perhaps best known illustration of this is the rule $S \rightarrow NP VP$, which, in some form or other, is part of most generative grammars that have been proposed for English. This rule says that an expression of the category ' S (entence)' can be formed by concatenating an expression of the category ' N (oun) P (hrase)' with an expression of the category ' V (erb) P (hrase)'. The members of a grammatical category can be either lexical items or complex expressions. Lexical categories are those which contain at least some words. (It is possible for a lexical category to consist of lexical items only, but in general this won't be so.) Examples of familiar lexical categories, which will be found in any grammar for a language such as English, are 'Noun', 'Verb', 'Adjective', 'Adverb' and 'Preposition'. In addition to lexical categories many grammars also postulate certain non-lexical categories, which contain no lexical items but only complex expressions.

For the theory of meaning grammatical categories are important in that expressions of the same category will have denotations of the same logical type. For instance, the denotations of expressions of category Noun generally are properties — or, in another formulation which we will favour in what follows, the extensions of those prop-

52. Cf. Montague (1970a).

53. Among the building rules for a language like English there are those which state how full words can be built out of their stems by addition of certain morphemes. For instance, the past tense form 'called' of the verb 'to call' is formed by concatenating the stem 'call' with the past tense morpheme '?ed'. In what follows we will ignore the distinction between words, stems and morphemes. For our purposes morphemes and stems can both be thought of as 'lexical items', i.e., as elements of the vocabulary of the language, and full forms like 'called' can be thought of as complex expressions. (The interaction between syntax and morphology is one of the aspects of natural languages that make it awkward to press natural language grammars into the strict format of sets of construction rules.)

erties, i.e., the sets of entities of which a given property is true.⁵⁴ Another example: the denotations of elements of the category *S*, i.e., of well-formed sentences, are always propositions — the denotation of a sentence *s* is the proposition expressed by *s* — or, in the formulation favoured, the truth values of those propositions. It should be noted that being of the same category is a sufficient but not in general a necessary condition for having denotations of the same type. For instance, the denotations of verb phrases (i.e., members of the category *VP*) in many semantic theories are properties (or, alternatively, property extensions) just like the denotations of nouns.

So much for syntax. The standard method to account in a formally precise way for denotation and meaning is that of model theory. The method consists in (i) defining structures — the so-called ‘models’ — in which expressions of the different grammatical categories can be assigned suitable denotations; (ii) a specification in each model *M* of denotations for each of the lexical items of the language or language fragment *L* under consideration; and (iii), in order to account for the denotations of complex expressions, a general definition of how the denotation of any complex expression is determined by the denotations of its constituents. (Cf., the remarks made earlier about the semantics of generative rules.) Together (ii) and (iii) will assign in each model *M* a denotation to each well-formed expression. In particular we obtain a denotation in *M* for each of the sentences of *L* (which, as noted, will, depending on how the theory is set up, either be a proposition or a truth value (‘true’ in case the sentence is true on the interpretation provided by *M* or ‘false’ in case the sentence is false on that interpretation).⁵⁵

The model-theoretic concept of meaning is relational in that it connects expressions and models. This can be seen most clearly for the case of sentences, assuming that their denotations are construed as truth values. On this assumption a given sentence *s* is positively related to those models in which its denotation is ‘true’ and negatively to those in which its denotation is ‘false’. For expressions of other categories the matter is somewhat different insofar as their denotations aren’t simply truth values. But here too the denotation is the product of the interaction between expression and model, and can be seen as the manifestation of the way in which the two are related. We can isolate the contribution that the expressions make to these manifestations by associating with each expression *e* the function which maps each model *M* to the denotation of *e* in *M*. It has been suggested that the meaning of an expression can be identified with this function, in particular, that the meanings of sentences can be regarded as functions from models to truth values. We will see in sections 3.1 and 3.2, however, that such an identification isn’t possible in general.

In order that a model-theoretic account for a language *L* does justice to our intuitions about what words and complex expressions mean, great care must be taken with the definition of its models. Of particular importance is that only such models

54. ‘Generally’ because, e.g., so-called ‘natural kind terms’ (nouns such as ‘water’ and ‘gold’) may be taken to denote, not properties, but abstract essences.

55. In what follows we will, as indicated above, assume that the denotations of sentences are truth values and the denotations of nouns and other property expressions extensions; but we will briefly return to the other option, according to which sentences denote propositions and nouns properties, in section 3.2, footnote 60.

be admitted in which the denotations of words represent realistic possibilities. To give just one example, assume that our account identifies denotations of nouns as sets. Let $n_1, \dots, n_k$ be nouns. Then as a rule not any combination $S_1, \dots, S_k$ of sets of entities will constitute a conceptually possible combination of denotations for these words. Suppose for instance that n_1 is the noun ‘woman’ and n_2 the noun ‘man’. Then in any model M the denotations of n_1 and n_2 should be disjoint. This is a rather simple case of a semantic connection between two words that imposes a restriction on the models that should count as admissible in a satisfactory account of meaning. In general the connections are much more complex and more difficult to identify. And at the present time semantics is nowhere near a comprehensive inventory of the constraints that such connections impose.⁵⁶

The reason why this issue is relevant for the specific concerns of this chapter is that what information a sentence carries depends on the models it excludes — i.e., those models which are incompatible with what the sentence says, and which we are entitled to ignore when we take the sentence as giving us true information. But evidently, what the set of those models is, and how large a proportion it represents of the totality of all admissible models, depends on which models are admissible to start with.

In view of the importance that the question of constraints on models has for the central topic of this chapter, it is appropriate to dwell on it a little more. First something that has been implicit in the assumption that the denotations of nouns are sets. On that assumption the elements of those sets must be in some way part of M . The way in which this requirement is met is to assume that each model M comes with a domain of entities, or ‘individuals’, which includes all denotations of nouns as subsets. The domain of individuals can be used as a foundation on which a hierarchy of further domains of other, higher logical types can be built, using certain elementary set-theoretical operations. Some of these higher type domains correspond to the logical types that are associated with grammatical categories of L , with the understanding that the denotations in M of expressions of a category C will be members of the domain of the logical type associated with C . For example, the denotations of nouns are members of the domain of the logical type associated with the category Noun. So, if noun denotations are sets of individuals of M , then this domain must be the power set of the domain of individuals, i.e., the set of all subsets of that domain.

A model M must thus have at a minimum the structure of a range of domains of different logical types, held together by the relations which are entailed by the way in which higher type domains are constructed from the individual domain. But this is only one, comparatively simple part of the structure that is presupposed by the constraints that singles out the conceptually admissible models. For one thing, we need more structure within the different domains. This is true in particular of the domain of individuals itself. We already noted that the denotations of ‘woman’ and ‘man’ should always be disjoint. The same constraint applies to the nouns ‘wife’ and ‘husband’. But connected

⁵⁶ It should be noted that such conceptual restrictions derive from language, and one may well argue, as Quine has done in his attack on the analytic–synthetic distinction Quine (1953b), that they may not hold as such.

with these last two nouns there is a further constraint. Their denotations are delimited by the restriction that they can be felicitously applied only to human beings; or — to put this in the current terminology used by many linguists — both ‘wife’ and ‘husband’ come with what is called a ‘selection restriction’ to the set of human beings. When we look more closely at the use of nouns (in English or other languages), we see that pretty much all of them come with selection restrictions of some kind. Furthermore, the sorts that form the selection restrictions of the different nouns of the language form a complex hierarchical structure, with some sorts being proper sub-sorts of others. A simple example: the noun ‘bachelor’ is, in its most prominent use, restricted to men of a certain age and social position (excluding for instance those who have made a formal vow of celibacy). So its selection restriction is a sub-sort of the selection restriction of ‘husband’ and ‘wife’. A more thorough exploration of this phenomenon also makes clear that part of what is needed is an articulation of a sort hierarchy that provides, among other things, the selection restrictions of the different nouns of the language.

Another source of complexity is that most denotations change over time. In fact, this is so for two reasons, as can be seen plainly for nouns such as ‘wife’ and ‘husband’. First, the set of human beings changes over time, as new people are born and other people die. So death affects the denotations of these nouns directly in that they lose members because these disappear from the scene altogether. But people also enter and leave the denotations of ‘wife’ and ‘husband’ while alive — viz., by getting married or divorced.⁵⁷ To do justice to this temporal dependence of the denotations of nouns and expressions of other categories, models must include a time structure. Since this is an aspect of model-theoretic meaning accounts that is especially important in connection with what will be said below, we do well to be a little more explicit about it. We keep things as simple as possible, assuming that each model M includes a time structure $\langle T, < \rangle$, where T is a set of temporal instants and $<$, the ‘earlier-later’ relation, is a linear ordering of T . Furthermore, the domain of individuals of M may now vary as a function of time — that is, as a function of T — and the same goes for the higher type domains that are constructed from domains of individuals, for the sortal hierarchies that subdivide these various domains, and for the denotations in M of the words of L . Thus each expression e of L no longer has a single denotation in M , but a possibly different denotation for each $t \in T$.

These are just some of the complications that model-theoretic accounts of meaning must address. This is not the place to do more than indicate that these issues need to be dealt with, and give a rough idea of what they are. But one further remark, of a more general tenor, is in order. Both the structure of sort hierarchies and the nature and structure of time are matters of ontology, the science of ‘what there is’⁵⁸ This endeavour, of determining the kinds of entities that must be assumed to exist and their logical properties and relations, was for centuries the exclusive province of philosophy. In more recent times it has become a major concern in artificial intelligence and cognitive psy-

57. As the tabloids keep reminding us, weaving your way in and out of these denotations can become a form of life in its own right.

58. Quine’s happy phrase, cf., Quine (1948).

chology and this is where now much of the kind of work on ontology is being done that is relevant to the theory of meaning. That is indicative of an important aspect of the meanings of linguistic expressions and the information they carry: ontology is not just a part of a theory of the meanings of words (although, as we have argued, it is an indispensable part of such a theory too), but rather a general theory of the structure of the world that presents itself to, and is projected by our cognitive faculties — of the different kinds of entities of which that structure is composed and of the principles that hold this multiplicity of kinds together. Up to a point the languages we speak presuppose and mimic this structure, would be there even if we didn't speak a language, or didn't speak the particular languages that we do speak. But as pointed out earlier, languages also contribute to this ontology by projecting certain kinds of entities and structures on it. Thus, the 'ontology of language' is a complex affair, the result of external, causal influences from reality, the structuring principles underlying general cognitive abilities, such as perception, and the contributions made by linguistic structure.

Assuming that this assessment of the nature of ontology is correct, the claim that an account of natural language meaning must include parts of it amounts to the acknowledgement that the meanings of words (and, by implication, also those of larger expressions) are not 'autonomous', but are also constrained by general conditions that relate to the ways in which we perceive the world and structure our expectations about its regularities. A similar conclusion follows for the information that is carried by linguistic utterances: it too depends on the structure that cognition imposes on what it receives as input. Note that this implication is two-sided. On the one hand utterances could be said to succeed in carrying as much information as they do because their meaning implicitly relies on, and thus implicitly incorporates, so much of the cognitively based, though not specifically linguistic structures that our languages presuppose and exploit. On the other hand, the new information that an utterance conveys is limited by the fact that one must already be in possession of much of this implicit information in order to be able to interpret the utterance in the first place.

3.1 Utterance dependence of content

The next aspect we must consider of the way in which the denotations of natural language expressions are determined is very different from the one we have just discussed and it requires a shift of perspective, from expressions as such to their uses on particular occasions — that is, to utterances. For an illustration of the point at issue consider the following sentences:

- I. (a) That is a man.
- (b) He is a widower.

First (1a). One of the things one must understand in order to understand an utterance of (1a) is what is denoted by the word 'that'. What is its denotation? Well, that depends on whom the speaker of the utterance intends to denote by her use of 'that'. An interpreter will be able to determine what that entity is only insofar as the speaker provides him with some clue, for instance by pointing at the individual that she intends as denotation, or by gazing pointedly in its direction. This is a general property of 'that' and other so-called 'demonstrative' expressions: they can be used to denote pretty much anything,

and what they denote on a particular occasion is determined by what the speaker wants them to denote, as long as she conveys this to her audience by providing the right clues. Much the same goes for personal pronouns like the ‘he’ occurring in (1b). The denotations of ‘he’ are more restricted in that they always must be male (and usually human). But which male is again a matter of the speaker’s current intentions and her ability to get her intention across.

Because the denotations of the words ‘that’ and ‘he’ may vary from utterance to utterance, this is also true for the denotations of the sentences (1a) and (1b) themselves, since these depend on the denotations of the words they contain. There is however also another reason why the denotations of (1a) and (1b) vary, and this is a very general one. Time determines which denotations of the nouns occurring in (1a) and (1b), ‘man’ and ‘widower’, are to be combined with the denotations of ‘that’ and ‘he’, respectively. For instance, an utterance of (1b) at time t is a statement to the effect that the denotation of ‘he’ belongs to the denotation of ‘widower’ at t (rather than to the denotation of ‘widower’ at some other time). It is clear that this temporal dependence of (1a) and (1b) has to do with their tense, viz., that it is the present, rather than a past or future tense. For example, had the tense of (1b) been the simple past, as in (1c) below, then an utterance made at t would not have expressed that the denotation of ‘he’ belongs to the denotation of ‘widower’ at t , but to its denotation at some time before t :

1. (c) He was a widower.

Note well, however, that the utterance time is as indispensable to the interpretation of (1c) as it is to that of (1b). For although in the case of (1c) it is not the denotation of ‘widower’ at the utterance time t itself that is involved, the denotation times that are relevant are those which stand to t in a certain relation. They are times that precede t , and not t itself or times following it.

The temporal dependence exemplified in (1c), with its one verb in the simple past, is relatively simple. But a closer look at the full range of tenses, as well as at other expressions with which the tenses interact reveals a very complex field of temporal relations in which the events described may be linked to the utterance time complicated ways.⁵⁹ Here, however, it is not the complexity of these relations that matters, but the mere fact that what is needed to determine the denotations of sentences containing past or future tenses are not just the denotations of words and morphemes at the utterance time t , but also their denotations at other times. This is important because it entails that the denotations in a model M of sentences uttered at time t will in general require not just the denotations of their lexical constituents in M at t , but the entire ‘temporal history’ of M , providing denotations for all instants of its temporal structure.

59. Examples of such expressions can be found among adjectives (e.g., ‘former’, ‘repeated’), conjunctions (e.g., ‘while’, ‘after’, ‘before’) and prepositions (e.g., ‘after’, ‘before’, ‘during’, ‘ago’). Within the class of temporal adverbials we find representatives of a whole spectrum of distinct functions, as the following examples illustrate: ‘Monday’, ‘the twentieth of March’, ‘last week’, ‘often’, ‘every other Sunday’, ‘still’, ‘again’, ‘the second time’, ‘for the second time’. The semantics of tenses and other temporal expressions in English and a few other languages is one of the most assiduously researched areas of the theory of meaning, and much in this area is by now quite well understood. For a recent study, cf., Rothstein (2004).

Echoes of what we have just observed in connection with time can be found in the realm of modality, that part of the theory of meaning that has to do with the difference between the actual and the possible, the difference between what is true and what isn't but could have been. In fact, in the early days of the model-theoretic approach to meaning time and modality were treated as two dimensions of a simple ontological structure. Since then the general perspective has changed. According to more recent views the differences between time and modality outweigh the similarities, and most current formal treatments reflect this. But there is one similarity between the temporal and the modal that is as prominent in recent treatments as it is in older ones. This similarity can perhaps be brought home most forcefully by a look at subjunctive conditionals. Consider for instance an utterance of the conditional in (1d):

1. (d) If he had been a widower, she would have married him.

The sentence in (1d) relates two constituent sentences, the 'if'-clause and the main clause, and it is this relation which determines whether the conditional claim as a whole is to be counted as true. Moreover, whether the conditional is true does not just depend on what is the case in the world as it is. The conditional implies that both 'if'-clause and main clause are false in the actual world. But that is not enough; what is required in addition concerns other worlds than the actual one. Roughly, the additional requirement is that in any relevant possible world in which the 'if'-clause is true, the main clause should be true as well.

There is a large literature on conditionals (different in spirit from that on temporal reference and tense, but comparable in size). A large part of this literature is concerned with the difficult question how to define the concept of a 'relevant' possible world, as it occurs in the truth requirement we just stated for (1d). But once again, it is not such details that matter here, but only the general fact that to determine the denotations of utterances of conditionals (and other sentences containing modal terms or constructions) in one world we must have recourse to denotations in other worlds.

If we want our model-theoretic approach to deal with modality along the same lines that we have outlined for dealing with tense, then we must extend our models with yet another layer of complexity. What we need are not simply models that provide the development of denotations through time, but whole bundles of such models which cover not only the actual world but also other worlds that are relevant to modality-involving sentences of L . We will call such bundles 'intensional models' and refer to the models considered up to this point as 'extensional models'. (In other words, an intensional model is a bundle of extensional models.) Since intensional models will play an important part in all that follows, it will be useful to stipulate a specific form for them. The following definition is simple, but suits our needs.

By an intensional model $\mathcal{M}$ for a given language or language fragment L we understand a pair $\langle W, M \rangle$, where W is some non-empty set (of 'possible worlds') and M is a function which maps each $w \in W$ to an extensional model for L , i.e., to a model for L of the kind considered so far. (We write M_w (rather than $M(w)$) for the extensional model that M associates with each $w \in W$.)

This definition gives what you might call a 'bare bones' characterisation of intensional models. For many purposes the models it specifies won't be enough. For instance,

in order that a model yield a satisfactory analysis of various kinds of conditionals, it must provide, apart from what is specified by our definition, also certain relations between worlds (which tell us which worlds are relevant to the denotations of various modal sentences in which other worlds). But as we have said, the exact analysis of the denotations of particular sentences is not what concerns us here. And as we will see, for the purposes of this chapter our present definition gives us just what we want.

There is one aspect of intensional models, however, that does require our attention. This is the structure of time. So far we assumed that each extensional model has its own tense structure. But what can we say about the time structures of the different extensional models that make up a single intensional model? Are all these time structures the same, or may we expect them to vary from one extensional model to the next?

Behind this question lurks an age-old debate about the nature of time and the formal properties that follow from it. It is a debate that started out within philosophy, but that spread to several other disciplines once these had taken on their own topical and methodological identity, most notably to physics and psychology. The various positions that have been argued in the course of this debate can be divided into two main groups. On the one side there have been those who see time as a given absolute, either along the lines of Newton's *Principia*, or, more in the spirit of psychology or cognitive science, as some sort of Kantian category. On the other hand time has been seen as an immanent feature of an unfolding world of successive events, as an abstraction from its flow of events. (Well-known representatives of this second view are Leibniz and Russell.) On this view it cannot be excluded a priori that time structures inherit also some of the contingent properties of the event flows from which they are derived and thus that they vary from one world to the next. For proponents of a view of time of the first type it will go without saying that all extensional models come with the same time structure; for proponents of a view of the second kind this will not be self-evident, and some at least will want to deny it.

This is not the place to take sides in this debate. In general we should allow for intensional models that are consistent with either position, thus including those in which time structures may vary between their component extensional models. Variability of tense structure within intensional models, however, leads to certain conceptual and technical complications that it is better to side-step here. We will therefore, in the interest of presentational perspicuity, restrict our attention to intensional models that each have a single time structure.

3.2 Content and meaning

In section 3.1 we drew attention to two complications that model-theoretic accounts of meaning must deal with: (i) the dependence of denotations on utterance features other than the linguistic form of expression uttered, and (ii) the power of an utterance to make a statement not just about its here-and-now, but also about what lies beyond — in the past, in the future or even in other possible worlds. These are by no means the only complications that theories of meaning have to deal with. But we have singled them out because it is they which affect the general form of a theory of meaning most deeply and therefore it is they also which have the greatest impact on answers to the

questions that are the principal business of this chapter.

The main questions that will occupy us in the remainder of the chapter are: What is the propositional content of an utterance? What, if anything, are the contents and meanings of sentences? And what is the information carried by a natural language utterance? We will deal with the first two of these in the remainder of this section. The last question — which is the central question of this chapter — will be discussed in section 5.

As a preamble to answering the first question recall that earlier in this section, when we first spoke of denotations, we mentioned that the denotations of sentences could be either construed as propositions or as truth values, and that meaning theories vary on this point. There is a close relation between those two concepts of sentence denotation, just as there exists a close relationship between properties and property extensions: a proposition can be either true or false, depending on whether the situation to which it is applied is compatible with what the proposition says or not. It has been argued that these manifestations of the proposition, its being true in some situations and false in others, are all there is to its identity, i.e., that a proposition is nothing but the truth values it takes in different possible situations. On this view a proposition can be identified with the function which returns for each situation the truth value that it has in that situation. We use this as our leading idea in formulating our answer to the first question.⁶⁰

We have seen in section 3.1 that in general it is only utterances of natural language sentences that can be said to have definite denotations, but not those sentences by themselves. By the same token it is only to utterances that we can attribute definite content, and not to sentences per se: it is utterances, not sentences, that express propositions. So, if we want to stick to the spirit of our leading idea, it is to utterances, and not to sentences as linguistic expressions, that we should apply it. That is, utterance content should be defined as ‘propositional content’ — viz., as the range of truth values that an

60. This is the point to return to the question how one might choose between model-theoretic accounts which construe sentence and noun denotations as truth values and sets, respectively, and those which construe them as propositions and properties. At the level at which the answer to this question is at all relevant to the issues of this chapter, it is quite simple, and also quite uninteresting. In model-theoretic accounts which exclusively make use of extensional models only the former denotations (truth values and sets) are well-defined, so it is only in that way that sentence and noun denotations can be understood. In accounts that use intensional models, both construals are possible, but there is little to choose between them. First, everything that can be done with truth values and sets as denotations can also be done when the denotations of sentences and nouns are taken to be propositions and properties, since we can always pass from propositions to the truth values they have in particular worlds or models, and from properties to their various extensions. Conversely, when propositions are defined, as suggested above, as functions from possible worlds to truth values, then in an intensional model $\mathcal{M}$ it is in principle possible to recover propositions from the corresponding truth values in the different extensional models that are part of $\mathcal{M}$ (and a similar reconstruction is possible when properties are construed as functions from worlds to extensions). So in theories that make use of intensional models the two ways of construing denotations are equivalent so long as the technical machinery is in place for going from propositions and properties to truth values and sets, and back. In all model-theoretic accounts of which we are aware, however, this machinery is available. As noted earlier, there appears to be a preference for theories in which denotations are construed, like we have been doing here, as truth values and sets. But as far as we can see, there are no compelling reasons for this preference.

utterance determines in different possible situations. Or, stated in terms of intensional models: the content of an utterance relative to an intensional model $\mathcal{M} = \langle W, M \rangle$ should be defined as the range of truth values that the utterance determines in the extensional models M_w associated with the different worlds $w \in W$.

At first sight it may look as if defining utterance content along these lines runs into a snag. Consider once more a sentential utterance u , for instance one of sentence (e) (which is like (b), except that it doesn't have the pronoun 'he' so that the only utterance feature that its interpretation depends on is the utterance time):

1. (e) Helmut is a widower.

According to what we have just suggested, the content of u relative to $\mathcal{M}$ should be identifiable as the function which maps each world $w \in W$ to the truth value of u in the model M_w . But what is this truth value of u in models associated with worlds in which u has not actually been made? Since the truth value determined by u depends not only on the sentence uttered, but also on some further properties of u , notably the utterance time, it cannot, one might think, be taken for granted that truth evaluation is possible also in relation to other worlds.

Fortunately this worry can easily be put to rest. Intuitively it seems clear that the world in which the utterance u is made could have been different from what it is, but that that would not have made any difference to the possibility of enquiring whether or not it makes the statement that u expresses true. The only difference might have been that the enquiry might have led to a different outcome. The intuitive reason why this should be so is that once the utterance time t of our utterance u has been fixed, as the time at which the utterance act is performed in its world w , the denotation of u at that time t can be computed just as easily in models M'_w that are associated with worlds w' different from w as it can be in the model M_w associated with w itself. All we need to assume for this is that t can be identified as a time of those other worlds too.⁶¹

Now that we have resolved the apparent snag, nothing stands in the way to the intended characterisation of utterance content:

The *propositional content* of an utterance u of a sentence s , relative to an intensional model $\mathcal{M} = \langle W, M \rangle$, made at a time t (of the time structure of $\mathcal{M}$), is the function which maps each world $w \in W$ to the truth value of s at t in M_w .

We now turn to the second question: What, if anything, is the content or meaning of a sentence? As regards sentence content we can be brief and simply repeat what we have noted already: given that there can be no definite sentence content without definite sentence denotations — that is, definite truth values —, there can't be a defi-

61. It is here that our assumption that all extensional models belonging to a given intensional model have the same time structure is being used. Without this assumption arguing for the present conclusion becomes more complicated, since it will involve the question how the possibility of identifying t in other worlds than w correlates with the relevance of those worlds for determining the denotation of the utterance in w . Other complications arise when additional utterance features besides the utterance time play a part in the content of u . All in all there are many non-trivial details that an elaboration of the argument we have sketched here must deal with. We refer the reader in particular to the locus classicus for these issues Kaplan (1989). Further discussion can be found, e.g., in Almog et al. (1989).

nite content for any sentence of which the interpretation depends on additional features of its utterances. It is still possible, however, to make sense of the notion of sentence meaning, viz., as that which enables the different possible utterances of a sentence to express their respective propositional contents. Understood in this way the meaning of a sentence s can be identified with the function that maps each possible utterance of s to its propositional content. This brings us to the following formal characterisation:

The *meaning* of a sentence s relative to an intensional model $\mathcal{M}$ is that function which maps each possible utterance u of s in some world w from $\mathcal{M}$ at some time t of its time structure to the propositional content of u .

We have pointed out several reasons in this section why the information state of a recipient who is in a position to interpret an utterance, and thereby profit from the information it contains, cannot be a tabula rasa. But there also is a further reason, which is connected with the fact that natural language utterances show a strong tendency to build upon those that precede it in discourse. In fact, human languages are rich in devices that serve this very purpose — devices for linking the sentences in which they occur to the sentences that precede them in the texts or dialogues of which they are part. These devices enable the recipient to interpret the sentences that contain them in the way the speaker intends — viz., as integral pieces of a larger discourse. But of course this can work only if the recipient has already interpreted those preceding sentences and has thereby acquired the information which they carry. In this sense too the information he will get from the new sentence takes the form of an increment to the information he already had. In the next section we will have a closer look at this incremental dimension of interpretation, and of the acquisition of linguistic information that goes with it.

4 Modelling discourse connections

Much of what we want to say we say in several sentences. Single sentence utterances suit only the simplest of messages, as soon as the message becomes a little more complex, a single sentence won't do. Strictly speaking, of course, conveying a complex message in a single sentence isn't impossible in principle. But the sentence that one would have to use would be so long and convoluted that others would have the greatest difficulty in unscrambling the message; and even the speaker himself would be likely to get tied up in knots and lose track of what he was saying. This humdrum fact about the use of language points at an aspect of our language handling capacities which is also quite obvious. Our ability to parse sentences, i.e., to ascertain their syntactic form, is not commensurate with our capacity for grasping and retaining content. For whatever reason parsing is, apparently, something that we humans find hard as soon as we are confronted with strings that exceed a certain length or structural complexity. Such sentences should therefore be avoided, and instead the story one has to tell must be broken up into a sequence of sentences that are each of manageable size.

But breaking up a message into a sequence of sentences each of which covers some part of it comes at a price. It requires that each sentence can be recognised as

making a particular contribution to the larger content. That is, the recipient must be able to see how and where the contribution of each new sentence fits within the part of the message that he has already reconstructed from preceding sentences. Sometimes it is clear from the nature of the message that is conveyed in a sequence of sentences and from the sequential order in which the sentences are arranged how the contributions of the successive sentences fit together. But this is by no means always so. In such cases it will be helpful, or even imperative, that the new sentence contains certain elements that make its connections with the preceding sentences clear. Given how important it is to get these connections right, it should not be surprising that natural languages include various types of such 'discourse linking' elements. In fact, there are many such elements and many of them are in constant use. (We know this to be the case at least for English and the range of other languages for which the question has been investigated, and we suspect that it is universal.)

Among the classical examples of sentence constituents which are capable of linking the sentences in which they occur to the preceding discourse are anaphoric pronouns. Anaphoric pronouns can have antecedents which occur at some earlier point in the same sentence, but as often as not their antecedents are not sentence internal. Often, but not necessarily, they occur in the immediately preceding sentence. In such cases the link between pronoun and antecedent also establishes a link between the content of the sentence containing the pronoun and that of the sentence that the antecedent belongs to. By way of illustration consider (2):

2. All these years Bill has kept in touch with one of the girls from his class in his final year in high school. He met her again last summer.

Here the pronoun 'her' can (in the absence of further context) only be construed as referring to the 'one of the girls from his class in his final year in high school' who is spoken of in the first sentence. This construal links the content of the second sentence to that of the first: the woman that, according to the second sentence, Bill met last summer is the same person as the girl that he has kept up with since his high school days. And in so linking the new content to the preceding one it also makes the former dependent on the other. It is a kind of 'add-on', i.e., an additional specification of the relation between Bill and the girl that the first sentence has already put on the interpretational map. This incrementality of discourse meaning, with the contributions by later parts building on those of earlier parts, is an aspect of linguistic meaning that substantially alters and complicates the picture of sentence meaning and utterance meaning sketched in section 3. Yet, as a feature of how natural languages work it is pervasive, and it comes in many different forms, the range of which is being uncovered only gradually.⁶²

This is not the place to explore this range in depth, and we present just a few more examples that may give some flavour of what forms discourse linking can take. In example (3), the subject phrase of the second sentence, 'the other two', establishes a number of connected links with the subject of the first sentence. Because of the

62. The systematic study of the effect of pronouns and other expressions with discourse linking effects is of comparatively recent date. It is one aspect of the approach to the study of meaning now widely known as 'dynamic semantics'. Cf., the references above, in footnote 37.

constraints that accompany these links (3) will be acceptable only if the number of students in the speaker's logic class was three. As a consequence, a recipient of (3), who assumes that the speaker has expressed herself in a way that is in keeping with what she is trying to convey, will conclude (in case he didn't already know that) that there were three students:

3. One of the students in my logic class flunked the final. The other two didn't turn up.

The constraints just spoken of are typical of expressions that establish this kind of 'anaphoric' discourse links.⁶³ Constraints of this kind, which utterances impose on the contexts in which they are made, are very common. In many (and presumably in all) languages there is a large variety of words and grammatical constructions which encode such constraints. The cover term that has come to be used for them is that of '(linguistic) presuppositions', or 'presuppositional constraints'.⁶⁴

Presuppositional constraints do not always affect the content of the utterances which generate them in the way illustrated by (2) and (3). Two examples where this is not the case are the presuppositions triggered by the words 'too' and 'again' in (4a) and (4b), respectively.

4. (a) Yesterday John came too.
(b) Yesterday John came again.
(c) Yesterday John came.

An utterance of (4a) carries the presupposition that there was somebody else who came yesterday, and one of (4b) the presupposition that there was an occasion before yesterday when John came. Here the content that is asserted is in both cases the same as would have been conveyed by an utterance of (4c). Utterances of (4a) and (4b) differ from utterances of (4c) only with regard to the contexts in which they 'sound right', but not in the content they contribute when they do. But even so they, too, tend to produce discourse-linking effects. For instance, the 'again' of (4b) draws attention to the fact that the described event — that of John coming (presumably his coming to an occasion of a certain, repeatable kind) was a repetition of something that had happened before. Such

63. This is true also for the pronoun 'her', which requires that its referent be a female person (if we ignore special uses such as making reference to a ship or to a female animal to which the speaker feels or wants to imply a person-like relationship). This constraint is confirmed by the anaphoric link between the occurrence of 'her' in the second sentence of example (2) and the argument of 'with' in the first sentence. Had the 'with'-argument been 'classmate from his final year in high school' instead of 'one of the girls from his class in Bill's final year in high school', then the interpreter, seeing the 'with'-argument as the only possible antecedent for 'her', would have concluded that the classmate in question was a girl.

64. The current tendency to subsume a large variety of context constraints under the term 'presupposition' is justified insofar as there is much that such constraints have in common, both in the ways in which they limit the contexts in which the expressions that generate them can be felicitously used and in the way they establish links to utterance contexts and thereby help to shape the content of connected discourse. But the term has the drawback that it tends to conceal some real differences which nevertheless exist between the various constraints that are subsumed under it. For the first clear recognition and articulation of the insight that anaphora and presupposition are closely related phenomena, and in fact that the terms 'anaphora' and 'presupposition' can be seen as each addressing one side of what is the same coin, cf., van der Sandt (1992). A somewhat different, though also essentially dynamic perspective on presupposition and anaphora can be found in Beaver (2001).

connections are often crucial to proper understanding. They belong to a dimension of discourse interpretation that lies beyond the reach of the notions of meaning and content outlined here, and which is only slowly becoming accessible to systematic, formally precise investigation.⁶⁵

Presuppositions are usually described as ‘constraints on the context’. That is a good way of describing their role and status, but it requires a certain understanding of the notion of context. As should be clear from our two anaphoric examples (2) and (3), the contexts that are the targets of anaphoric presuppositions are due to the preceding discourse. In fact, it is the content of the preceding discourse, as established by the interpretation that the reader or listener has made of it, that plays this role, and it is because of this double role — as content of what has been interpreted already and as context for what is being interpreted currently — that anaphoric constraints can engage with it in the way they do, relying on it for the satisfaction of the constraints they express and at the same time augmenting it (in its role as discourse content) with the content contribution derived from the current utterance. Thus discourse contexts — as contexts deriving from a discourse or discourse segment are usually called — function the way they do because they are content and context all in one. This unity of content and context is a direct reflection of the fact that the process of discourse interpretation is incremental, in that it modifies the discourse content/context step by step, adding each time the content contributed by the utterance or sentence that it has reached, after checking that the discourse context meets the current context constraints.

The incremental picture of interpretation throws an important new light on the nature of utterance content. The notion of content that we arrived at towards the end of section 3 was that of a set of possible worlds — those in which the current utterance is true. But for sentences whose interpretation requires linking one or more constituents to the discourse context this notion is no longer viable. Rather than determining a set of possible worlds in its own right all that an utterance of such a sentence can be said to identify by way of content is what it contributes to the discourse context established by the antecedent part of the discourse to which it belongs. In abstract terms this contribution can be characterised as a pair $\langle C, C' \rangle$ of discourse contexts, where C is provided by the antecedent discourse and C' is the discourse context that results from updating C with the contribution that is made by u , assuming that updating C with u is possible, i.e., that C is a discourse context which provides all that is needed for a proper interpretation of u .

In section 3.2 we have defined the meaning of a sentence S relative to an intensional model $\mathcal{M}$ as the function which maps each utterance u of S at a time t of $\mathcal{M}$ ⁶⁶

65. The most ambitious current approach to this dimension of discourse interpretation that is familiar to us is the ‘segmented discourse representation theory’, developed by Asher and others. Cf., Asher & Lascarides (2003).

66. In the preceding section we emphasised the dependence of the propositional content of an utterance u on the time t at which u is uttered. The central topic of the present section is the dependence of utterance content on the discourse context. But of course, the second dependence does not abrogate the dependence on utterance time and other features of the utterance context, such as the identity of the speaker and that of her addressee(s). In other words, in general we find dependence both on the discourse context and on features of the utterance context. However, in the remainder of this chapter we will suppress explicit

to the propositional content of u in $\mathcal{M}$. In analogy to the relational characterisation of utterance content just given, we revise the notion of sentence meaning as follows. The meaning of a sentence S relative to an intensional model $\mathcal{M}$ is the partial function f_S which (i) maps each utterance u of S made in a discourse context C onto the pair $\langle C, C' \rangle$ in case C can be updated with the propositional content of u in $\mathcal{M}$ and the result of that update is C' ; and (ii) is undefined otherwise. Such functions f_S are known as ‘context change potentials’ (*CCPs* for short), or ‘update potentials’.

What *CCPs* are like depends first and foremost on how we identify discourse contexts. The first proposal that might come to mind is that discourse contexts can take over the role of our earlier utterance contents, and thus can be identified with sets of possible worlds. That is to say, one might think that, although the content of a sentence utterance can no longer be identified in those terms, it should still be possible to identify the content of a discourse in this way (and by the same token the content of any initial segment of it). On this assumption sentence meanings become functions from sets of possible worlds to such sets.

But this proposal won’t work. It fails as soon as the question what are the available interpretation options for anaphoric pronouns is taken seriously. What issued may be involved in settling this issue is illustrated by the following example:

5. (a) One of the ten balls is missing from the bag. It has probably rolled behind the sofa.
- (b) Only nine of the ten balls are in the bag. *It has probably rolled behind the sofa.

The point of this example is this: The first sentence of (5a) and the first sentence of (5b) are true in precisely the same circumstances; an utterance of the one is true in precisely the same worlds as an utterance of the second.⁶⁷ So if the contents of initial discourse segments are identified with utterance contents old style, then the discourse context determined by an utterance of the first sentence of (5a) will be indistinguishable from the discourse context established by an utterance of the first sentence of (5b). But then it becomes inexplicable why the anaphoric interpretation of the occurrence of ‘it’ in the second sentence of (5a) is possible but a similar interpretation of its occurrence in the second sentence of (5b) is not.

reference to features of the utterance context, including the utterance time, assuming that these are given with each individual utterance and could be recovered from these when necessary. So, when from now on we speak of the content of an utterance u in an intensional model $\mathcal{M}$ we mean the content of u at the time t of $\mathcal{M}$ at which u is made.

A few remarks about utterance context and discourse context as parts of a more inclusive notion of context can be found later in this section.

67. Or, to put the point more pedantically, suppose that u is an utterance of the first sentence of (5a) that is made at some time t in some world w of a given intensional model $\mathcal{M}$, that u' is an utterance of the first sentence of (5b) that is made at the same time t in some other world w' and that in all other respects w and w' are exactly the same (so that in particular they contain exactly the same situation pertaining to the balls, bag and sofa that the utterances of (5a) in w and of (5b) in w' are targeted on. (It is a reasonable assumption that for all or most worlds w in which there is an utterance of the first sentence of (5a) at t there is such a world w' with a corresponding utterance u' of the first sentence of (5b), and conversely.) Then the set of possible worlds of $\mathcal{M}$ in which u is true will be the same as the set of worlds of in which u' is true.

Examples like (5) indicate that pronoun interpretation is sensitive not just to what the described world must be like given what has been said about it — not just, to put the matter somewhat more abstractly, to what possible worlds are left as candidates for the described world — but also to certain aspects of how the discourse has described it. The first sentence of (5a) provides a description that permits interpreting it as referring to the missing ball, while the description provided by the first sentence of (5b) does not. A notion of discourse content that is to provide the basis for an account of these facts of pronoun interpretation must differentiate between the discourse contexts established by utterances of these two sentences; and utterance contents old style just don't do that.⁶⁸

How can we modify the notion of discourse content so that this difference is captured? Not much reflection on examples like those in (5) is needed to see that the decisive difference between (5a) and (5b) is that the first sentence of (5a) 'introduces the missing ball into the discourse', as a kind of discourse entity, whereas the first sentence of (5b) does not do this. In (5b) the existence of the missing ball can only be inferred from the (old style) content, and apparently that is not good enough for the purpose of providing an antecedent for a singular pronoun. It may not be immediately clear how this idea can be turned into a formal definition of discourse content. To some extent this is brought out by the existing dynamic semantics literature, where a non-trivial number of different definitions can be found. We mention just one of these proposals, which seems to us a comparatively simple and natural realisation of the basic idea. According to this proposal each discourse content (relative to an intensional model $\mathcal{M}$) is based on a certain set X of so-called 'discourse referents' and consists of a set I of pairs $\langle w, f \rangle$ where w is a world from $\mathcal{M}$ and f is a function from X to parts of w . (The common domain X of all the functions which occur as second components of pairs in I is called the referential base of I .) The discourse referents that make up the referential base X of a discourse content I should be thought of as the entities that are explicitly introduced by the discourse whose content is I . (Thus the referential base of the discourse content I_1 determined by an utterance of the first sentence of (5a) will consist of three discourse referents, one for the missing ball, one for the set of 10 balls of which the missing ball is a member and one for the bag. The referential base for the discourse content I_2 determined by an utterance of the first sentence of (5b) will also contain three discourse referents, one for the set of nine balls in the bag, and further, as in the case of I_1 , one for the set of ten balls and one for the bag. The base of I_1 contains an element that can serve as antecedent for it, the base for I_2 does not.)

Discourse contents of this type are usually referred to as 'information states'. Information states of this kind are about the simplest and most conservative refinement of our earlier notion of content as a set of possible worlds. They reflect only one aspect of the form of the discourse, viz., the set of discourse entities it introduces, which the information state captures through its referential base. Notwithstanding this conservativeness, information states thus defined can account for a remarkable variety of discourse linking mechanisms, some of which look at first glance quite different from the constraint

68. Alternatively, one might attempt to analyse these examples in terms of descriptions, as proposed, e.g., by Neale Neale (1990); but cf. Peregrin & von Heusinger (2004) for arguments why this will not do.

on pronominal anaphora illustrated in (5). This is not to say that all such linking phenomena can be accounted for on the basis of discourse contents of this particular form. In fact, other, richer notions of discourse content have been proposed in order to deal with certain phenomena for which information states of the sort just defined do not seem adequate. At present the question what notions of discourse content are optimal for dealing with which aspects of discourse linking is far from settled.

We note for good measure that, obviously, the notion of a *CCP* co-varies with that of an information state, *CCPs* are always partial functions from information states to information states. Refinement of the notion of information state is automatically reflected in a similar refinement of the corresponding notion of a *CCP*.

There are two other fairly obvious points, which we record for further reference. The first is that an information state always determines a content in the sense of section 3. For an information state I of the kind defined above this content is the set $\text{prop}_{\mathcal{M}}(u)$ consisting of all worlds w such that for some $f: \langle w, f \rangle \in I$. In case the notion of information state is refined (in order to adapt it to the explanation of more complicated cases of discourse linking), the reduction to propositional contents may take a different form and be somewhat more involved. However, it is an essential ingredient to the notion of information state that each information state determines as propositional content, and thus that such reductions are always possible.

The second point is that among the sentences of a language such as English there are many whose content does not depend on the discourse context. Or, more accurately, there are many sentences such that any utterance of them has a content that is independent from the discourse context in which it is made. This is so whenever the sentence is free from anaphoric requirements and other presuppositional constraints. The content of an utterance of such a sentence S can still be identified in the manner of section 3 with the set of possible worlds in which the uttered sentence is true. Or, putting things more formally, given an intensional model $\mathcal{M}$ an utterance u of a sentence S (relative to a time t and a world w from $\mathcal{M}$) will have propositional content $\text{prop}_{\mathcal{M}}(u)$ that is independent of the discourse content C in which u is made. In such cases updating C with u will always be possible and will lead to a new discourse context C' of the propositional content is the intersection of the propositional content of C with $\text{prop}_{\mathcal{M}}(u)$. In particular, when C is identified with information state I , then the result of the update with u update is an information state I' such that $\text{prop}_{\mathcal{M}}(I') = \text{prop}_{\mathcal{M}}(I) \cap \text{prop}_{\mathcal{M}}(u)$. Of course this does not determine I' completely. In fact it is part of the point of example (5) that the first sentence of (5a) and the first sentence of (5b) can update the same information states I , that the resulting updated information states I_a and $I - b$ have the same propositional content, but that they nevertheless differ in such a way that the same second sentence in (5a) and (5b) can serve as an update of I_a but not of I_b .

The question how much structure must be given to discourse contents so that they are suitably equipped for accounts of the various forms that discourse linking can take must be sharply distinguished from another question: should discourse contents be defined, as we have done so far, as model-theoretic constructs, i.e., as set-theoretic constructs built out of models and their components)? Or should they be characterised as semantic representations ('logical forms'), i.e., as structures that have their own syntax

as well as a semantics determined by their syntactic structure (in the same sense in which, say, formulae of the predicate calculus have a semantics fully determined by their syntax)? Questions of the first type arise irrespective of how the second question is resolved, no less for those who opt for a representational approach than for those who prefer the ‘non-representational’ mode of analysis we have been following.⁶⁹

Reasons for choosing between a representational and a non-representational approach should be looked for elsewhere. First of all, for the computational linguist, whose task is the design and implementation of algorithms for processing language on a computer, the representational approach is the only option. Only finite objects like representations can be computed and manipulated to further computational ends. The typically infinite structures which non-representational approaches use to model content and information are, as such, fundamentally non-computable, and even a computational linguist whose theoretical inclinations lean towards the non-representation perspective will be obliged to work with syntactic expressions which provide finitary descriptions of the infinitary objects he favours.

A second observation that might be seen as pointing towards the representational approach has to do more specifically with the cognitive dimension of language. Languages are used by people, whose processing capacities are, just like those of the computers they manufacture, finite; and when language is used, it is the finite minds of people which do the processing that is involved in both production and interpretation. Cognitive science is still at a stage where the most fundamental questions about how the mind works are a matter of debate, and this is true in particular with regard to the question how the mind processes language and represents the results of this processing. We still aren’t in a position to assert with full confidence that there is any sense in which the mind can be said to process language ‘symbolically’, in the sense of forming and manipulating representations similar to those assumed in theoretical linguistics. This goes for all levels of linguistic representation, including those of syntax and of semantics. Nevertheless, many of those concerned with the cognitive dimension of language work on the assumption that language processing is largely symbolic in a non-question-begging sense, and that this is so in particular for processing at the syntactic and semantic level. When this general assumption is combined with the observations about discourse interpretation to which this section has been devoted, the inevitable conclusion seems to be that interpreting discourse involves representations of discourse content that can be incremented along the lines our discussion has indicated.⁷⁰

Alternatively, one might regard the choice between a non-representational account that assumes infinite entities as semantic values and a representational one that offers finite ‘reductions’ of those entities as a false dilemma that in fact is a mere artifact

69. This is not to say that the two approaches offer the same ranges of options for dealing with such questions. The question exactly how the representational and the non-representational approaches compare on this issue is still largely unanswered.

70. One of the long-term goals of the most representational version of dynamic semantics, discourse representation theory, is to uncover aspects of the semantic representation of content derived from linguistic input (i.e., the content that the recipients of linguistic input in spoken or written form extract from what they hear or read. Cf., Kamp (1984–1985, 1990).

of the specific model-theoretic assumptions that derive from the classical model, and ultimately that classical logical approach to the treatment of formal languages. After all, there seems to be no principled reason why one would assume the various ingredients in a non-representational account (such as domains, sets of worlds, models, and even the language itself) to be infinite. In effect, it can be argued that the very notion of a language as an infinite object, which, when combined with semantic compositionality, brings along the conception of an infinite set of meanings, itself is a theoretical construct, the result of a set of assumptions we may make in order to facilitate the study of certain linguistic phenomena (such as studying syntactic productivity without worrying about actual performance limitations, or defining information update as elimination of possibilities), but that we may also discard again when they get in the way.⁷¹

In section 3 we drew attention to the role of the utterance context in the determination of the utterance content. In the present section we have focused on the role of the discourse context, and in doing so we have kept the dependence of content on the utterance context out of sight. This compartmentalisation is consistent with a practice that up to the present time has been quasi-universal: dependence on utterance context and dependence on discourse context are hardly ever discussed in the same breath. But it is a practice which has little to speak for itself. What we want is an integrated account of information content, which deals with dependence on utterance context and dependence on discourse context in tandem. As far as we can see there are no fundamental obstacles that stand in the way of such an integrated account. But to our knowledge none has yet been fully worked out.

In an account of this kind utterance context and discourse context could be treated as completely separate and distinct, much as we have been presenting them in their respective sections 3 and 4. But from the perspective which provides the motivation for such an account it is tempting to see utterance context and discourse context as two components or aspects of a single comprehensive context, much as the terminology, which describes both as ‘contexts’, suggests. This second perspective seems especially compelling when we see utterance interpretation as something that takes place in the mind of the recipient. The recipient has to work with whatever information is available to him, and that information consists, apart from the expression uttered, of contextual information of various kinds, including information about the circumstances of the utterance event and the discourse context as the recipient has thus far constructed it.

As a matter of fact, the contextual information that the interpreter of an utterance relies on typically includes more than just his discourse context and information about the utterance context. For instance, world knowledge (including knowledge about both necessary as well as defeasible regularities that govern the events of the world in which we live) and the encyclopedic information that competent speakers of a language associate with most of its words are notorious for being indispensable to interpretation, and thus it is natural to take them to be parts of the context too. Along these lines we are led to a notion of an interpretation context as a complex structure,

71. For an early, philosophical argument along these lines, cf., Davidson (1986); Groenendijk & Stokhof (2005) contains some thoughts in this direction regarding compositionality.

in which utterance context and discourse context are just two of the components. One way in which these components differ from each other is that some of them change in the course of a multisentence discourse while others, such as for instance the world knowledge component, normally remain fixed. But the dynamics of such integrated contexts will also involve interactions between different context components, leading for instance to information being transferred from the utterance context to the discourse context. Describing this dynamics correctly will be one of the major challenges for such an integrated context theory.

Let us take stock once more and see what the present section has taught us about content and information. We found that since discourse interpretation is incremental in nature, utterance content cannot in general be identified with sets of possible worlds, but rather has to be accounted for in terms of the updating effects that utterances produce and the context change potentials of sentences of which those updates are the manifestations. This led to a pair of two related conceptions of information, that of an information state as embodying the information content of a discourse, and that of a context change potential — a partial function from information states to information states — embodying the contributions that the utterances of sentences make to the discourses of which they are part.

It should be emphasised, however, that while these new notions cast a significantly different light on the nature of linguistic information, they are, just like the notions presented in section 3 they are meant to supplant, notions of information of a user-neutral sort, which abstracts away from the needs, interests and antecedent beliefs and convictions of those to whom information is imparted. They address the question what the information that is conveyed by an utterance or discourse actually means for the one that it is conveyed to only insofar as they identify what information he could get from the linguistic input if he interprets it in accordance with the rules of the language, including those rules that govern the incremental interpretation of multisentence discourse. But they have nothing to say about, for instance, the different effects that the same linguistic communication will have on two different recipients, either because of differences in prior information with which the newly acquired information can be inferentially combined, or because of differences in their respective needs and interests that the same new information may address to different extents or in different ways. In the next section we will have a brief look at some aspects of the cognitive perspective according to which information should be assessed in terms of what difference it makes to the one who receives it.

5 Modelling the recipient

Information content as we have explicated it so far is, as we noted at the end of section 4, independent of who is receiving it. When a speaker A makes a statement, using some discourse-context-independent sentence S , and B and C are among her audience, then, given how we have defined the notion, the information content of her utterance will be the same for B as it is for C . In an important sense this is right: on a natural

understanding of the term ‘information’ *B* and *C* have obtained the same information. But how informative an utterance is, is not determined solely by the information it conveys, it also depends on what is in and on the recipient’s mind at the moment he processes what the utterance has to say.

For one thing — this is the simplest but also the most telling distinction that can be made here — what is communicated to the recipient may be either new to him or it may be something that he knew already. For example, suppose that Anna tells Bernhard and Carl over dinner that Dorothea has gone to Paris. And suppose that this was news to Carl, but not to Bernhard, who had been informed about Dorothea’s trip the day before. Then there is a sense in which the information which Anna’s statement imparts to Bernhard is nil, while for Carl it may be significant news.

But this is not the only way in which one and the same communication may carry information of different significance to different recipients. It may be that the information is new to both Carl and Bernhard, but that Carl knew that Dorothea had applied for a job in Paris and accordingly infers from what Anna says that she must have been given the job, whereas to Bernhard, who knows nothing about Anna’s job application, the question what she will be doing in Paris may not even occur. Or, yet another scenario, suppose that neither Bernhard nor Carl knew that Dorothea had left for Paris, but that their attitudes towards her are different. To Bernhard she is just someone he vaguely knows and that he has never paid much attention to, but to Carl she has been the object of deep and unrequited love. Again Anna’s words will provide new information to both Bernhard and Carl. But in Carl they are likely to trigger heated speculation about what reasons Dorothea may have had for going where she did, and about what she will be doing now that she is there. Bernhard, on the other hand, may be expected to file the new bit of information without giving it a second thought, and perhaps he will have forgotten it as soon as the dinner party is over.

The obvious fact these examples illustrate is that the significance that a piece of information has for the one who gets it is a function of his antecedent information,⁷² and often also of his other mental attitudes — his concerns, convictions, affections, desires, goals and plans. This means that an account of the impact of information presupposes a systematic way of representing mental states, as composed of attitudes of these different kinds. At present no such theory of mental structure exists that has found general approval. But even if there were such a theory, this would not be the place to expound it. So we shall limit ourselves to a look at the first type of dependence mentioned above, viz., the dependence of the impact of newly conveyed information on prior information.⁷³

The information that a person has at a given time must be represented in his mind in some way. For our present purposes it won’t really matter in what way it is represented. It may be that the representation of at least some of the information takes the form of representations with a specific ‘syntactic’ structure — of formulae or terms

72. Here and henceforth we use ‘information’ as an epistemically neutral term, covering both what a speech participant actually knows and what he only thinks he knows.

73. Cf., section 7 for some references to work that takes dependencies on other factors, such as the action goals of speech participants, into account.

from some ‘language of thought’ — but this assumption won’t be essential for most of what we will have to say. What does matter is that the representations determine content. In view of what we have observed in section 4 this means that each such representation must determine an information state, where information states have at least the complexity of sets of world-assignment pairs. (For most of what we will say, however, it will suffice to assume that such representations determine propositional content.)

Let us consider, then, a person B who is the recipient of a statement u of some context independent sentence S . To make sense of u , B will have to process it in the light of what he knows. As was pointed out in section 3, this requires in the first place that B recognises the subject matter of u , i.e., the part of the world that u is about. And recognition means activating those elements of his information that pertain to this subject matter. It is this part that will be directly relevant to B ’s interpretation of u and that will be augmented with the information which his interpretation of u will produce.⁷⁴

Evidently, where there is selection of elements there must be elements to be selected. So we must assume that B ’s information can be subdivided into elements that pertain to different subject matters. We leave open whether such divisions of a person’s information into subject-related elements is always possible, and we doubt that the dividing lines can ever be entirely non-arbitrary and sharp. However, for our present purposes a rough idealisation will suffice. Let us assume that these elements, bits of information, can be identified in the simplest and most abstract way possible, viz., as sets of possible worlds. Second, let a subject matter SM be the following equivalence relation between worlds: two worlds w and w' stand in the relation if they contain exactly the same facts pertaining to the subject matter, but otherwise they may be as different as you like.⁷⁵ Example: Suppose that the subject matter is what happens in Paris on the first of January 2006. Then w and w' stand in the corresponding relation if and only if what is the case in Paris in w coincides exactly with what happens in Paris on that day in w' . Furthermore, if SM is a subject matter and C a bit of information (i.e., a set of possible worlds), then we say that C has nothing to say about SM if and only if C has a non-empty intersection with each equivalence class of SM .⁷⁶ And finally, two distinct subject matters SM and SM' , are said to be mutually independent if and only if each equivalence class of SM has a non-empty intersection with each equivalence class of SM' .⁷⁷

About the selection of that part C_{SM} of his information that the recipient B of an utterance u activates as pertaining to the subject matter of u we make two assump-

74. The importance of this selection as part of the process of interpreting utterances has been underlined in particular in ‘relevance theory’; cf., Sperber & Wilson (1995).

75. This means that the subject matter can be identified with a question in the so-called ‘partition semantics of questions’, with the equivalence classes representing the exhaustive answers; cf., Groenendijk & Stokhof (1984), (Groenendijk & Stokhof, 1997, section 4).

76. In the terminology of the partition semantics of questions: if C is uninformative with regard to the question, i.e., if it does not exclude any of the possible exhaustive answers to the question.

77. In terms of the analogy with questions once more: if no answer to either question implies or excludes an answer to the other.

tions. First, that C_{SM} is about a subject matter SM in the sense that it is the union of some set of equivalence classes of SM . And second, that B 's total information can be decomposed into C_{SM} and some other part C_{RM} (the 'remainder' of B ' information) in the sense that $C = C_{SM} \cap C_{RM}$, while at the same time C_{RM} has nothing to say about SM .

This puts us in a position to say something about the epistemic effect that u will have on B . Interpretation of u will lead to an augmentation of C_{SM} with the result of that interpretation. We denote this augmentation as $C_{SM} \otimes u$. It seems intuitively clear that the set-theoretic difference:

$$C_{SM} \setminus (C_{SM} \otimes u)$$

between C_{SM} and $C_{SM} \otimes u$ is a measure of the epistemic impact that u has on B .

On the assumption that B can represent thoughts about SM in some representation language L , there is also another, inference-related way of assessing the epistemic impact of u , viz., as the set:

$$\{\phi \in L : (C_{SM} \otimes u) \models \phi \ \& \ C_{SM} \not\models \phi\}$$

consisting of those representations belonging to L which B is in a position to deduce after he has obtained the information that is conveyed by u , but would not have been able to before. It is easily verified that the correlation between this characterisation of epistemic impact and the set-theoretic one is monotone: suppose that u and u' are two utterances such that:

$$(C_{SM} \setminus (C_{SM} \otimes u)) \subseteq (C_{SM} \setminus (C_{SM} \otimes u'))$$

Then also:

$$\{\phi \in L : (C_{SM} \otimes u) \models \phi \ \& \ C_{SM} \not\models \phi\} \subseteq \{\phi \in L : (C_{SM} \otimes u') \models \phi \ \& \ C_{SM} \not\models \phi\}$$

(The converse implication doesn't hold without further assumptions about the expressive power of L .)

The information C_{SM} pertaining to SM can be seen as one component of the recipient's total information at the time of interpretation, but it is not the only one on which interpretation depends. In fact, the most important component, which plays a central part in every act of utterance interpretation, is his linguistic knowledge — his knowledge of the grammar of the language and of its lexicon — and, presumably, of a host of so-called 'encyclopedic knowledge'.⁷⁸ Let us represent this conglomerate of the recipient's linguistic and extra-linguistic knowledge as C_{LK} .

78. Encyclopedic knowledge is knowledge that isn't purely linguistic, but that nevertheless is important to interpretation, partly because it includes many of the preconditions of individual words — you cannot properly understand the meaning of 'levitate' without having some knowledge of the 'common sense physics' of gravity and its practical effects, or of the financial term 'futures' without knowing something about the stock market, or of the noun 'quark' without a substantial portion of knowledge about quantum physics (which is even harder), and so forth —, and partly because it guides us in distinguishing plausible from implausible interpretations, and thereby helps us to deal with ambiguity and vagueness. To be sure, the boundary between what is linguistic knowledge *sensu strictu* and what counts as extra-linguistic, encyclopedic or world knowledge is itself rather vague.

The way in which interpretation depends on C_{LK} is of course very different from the way in which it can depend on C_{SM} , and it may seem odd to mention these two almost in the same breath. But it is important to emphasise this second dependency as well, for it is the knowledge that goes into C_{LK} which is responsible for the very special character of information in natural language: linguistic expressions have the capacity to carry the information they do because of this very large package of knowledge that is shared (with close to total overlap) between the members of a speech community and thus in particular between any two members who use their language in an act of communication: the speaker uses her knowledge to encode a thought in words and her interlocutor makes use of the same knowledge to decode the verbal message to reconstruct the encoded thought. This is what makes linguistic information into the special thing it is and language into the uniquely powerful communication tool that it is.

Note that as a rule C_{SM} and C_{LK} will be quite different, and in fact can be assumed to be mutually independent in the sense defined above: any way that the given subject matter could have been is compatible with any way that the language could have been. For instance, suppose once more that the subject matter of an utterance u is the current whereabouts of Dorothea. Presumably that subject matter has nothing to do with the conventions of the language; any possible fact about where Dorothea is, has moved from or is going to is compatible with any way the language of u could have been. Only in the special case when it is language itself that is the topic of discussion will C_{SM} and C_{LK} overlap, or even coincide.

Even in the special case where language itself is the subject matter, but certainly in all those where C_{SM} and C_{LK} are independent, is the way in which interpretation depends on C_{LK} clearly very different from the way it depends on C_{SM} — so different in fact that casting linguistic knowledge as part of the over-all context in which utterances are interpreted might seem rather artificial. But linguistic knowledge clearly is knowledge without which normal interpretation would be impossible, and, to repeat, for a proper appreciation of what makes linguistic information special the dependence of interpretation on this part of the interpreter's knowledge is crucial: that linguistic expressions have the capacity to carry the kind of information they do — and thus to carry as much information as they do — is due to the very large package of linguistic and encyclopedic knowledge on which the interpreter can and must rely, and which is common (largely, if not totally) to those who share knowledge of a given language.

This fact about human languages — that linguistic and paralinguistic knowledge is very extensive and that it is wholly or largely shared by those who can be said to speak them — is of particular importance for understanding the special nature of information as linguistically expressed and communicated. Part of the point here is not specific to natural language: In many contexts the information that is carried by a code belonging to a coding system of any kind is understood in terms of the coding and decoding algorithm that makes the system a coding system; and transmission of information using the system will function only if this algorithm is known to both sender and receiver. But of course, this is a notion of information that is derivative insofar as it presupposes some other language or medium in which information can be represented and with which the coding system is connected via its coding algorithm. It is a notion which passes, one

might say, the question ‘What is information?’ on to that language or medium.

Human languages are special on the one hand because of the sheer quantity of knowledge that must be shared by those who use them to communicate. It is this which explains the possibility of packing as much information into an expression of modest size as we often manage to do and yet getting it across to our audience. But it isn’t just the quantity of linguistic knowledge which makes the case of natural language special; it is also its quality. Our knowledge of our language isn’t just a coding system that enables us to translate into and from it messages that are given in some other representation system (some ‘language of thought’). It is in part genuinely semantic knowledge which links the expressions of our language directly to the world. It is because of these two properties of linguistic knowledge — that it is truly semantic and that it is shared between all speakers — which explain why linguistic expressions can be said to carry information in a non-derivative sense and at the same time be such remarkably effective information transmitters.

6 Information in natural language

In his chapter in this handbook Dretske reminds us of two distinct pretheoretical uses of the term ‘information’. When we say that the puff of smoke we see in the distance means that there is a fire there, or that puffs of smoke normally carry the information that there is a fire, we use the term in a sense in which information by definition is true information. But this is not the sense we intend when we speak of ‘linguistic information’. An utterance of the sentence ‘There is a fire over there’ can be described as meaning, or as carrying the information, that at the time of the utterance there is a fire somewhere in the direction that the speaker indicates. By itself this assessment does not entail that there is indeed a fire, in the relevant direction and at the relevant time, whenever someone makes a statement by uttering this sentence. Thus in cases such as this the term ‘information’ is used in such a way that its factual correctness is not assumed, i.e., in a way that allows for information that is true, but also for information that is false. It is this second sense in which we have been using the word ‘information’ throughout the present chapter. For it is this sense of ‘information’ that reflects the most fundamental characteristic of natural language meaning, viz., its ability to be about non-existent objects and non-obtaining situations.

Nevertheless, statements do carry a commitment to truth. It is constitutive of the practice of making statements that they are intended to convey true information, even if on occasion speakers fail to do so by mistake, or abuse the trust of their audience by lying. One situation in which this commitment to truth makes itself felt in the context of verbal communication is when the speaker makes statements which contradict the recipient’s beliefs — in other words, statements u such that the propositional content of $C_{SM} \otimes u$ is the ‘contradictory proposition’ (the empty set of worlds). In such a situation the recipient may react in one of several different ways. He may conclude that if the speaker felt confident enough to make her statement, she must be right, and revise his beliefs to fit the opinion she has expressed. But he may also feel certain enough about

his own beliefs to conclude that the speaker must be wrong; and in that case he may either keep his disbelief to himself or try to convince the speaker that she is wrong. And of course there are many shades between these two extremes. The recipient may feel strong enough about his own views not to accept the speaker's opinion without further ado, but not strong enough to dismiss her opinion out of hand. In such cases a discussion may ensue, perhaps ending in a joint view of whose opinion should be considered most likely.⁷⁹

Even if we ignore the range of possibilities between the extremes of unquestioned acceptance and unconditional rejection, just the two extremes themselves show that our earlier binary distinction between old and new information is too simple. There are three basic relations that the speaker's statement u can stand in to the recipient's assumptions C_{SM} about the subject matter of u (or what he takes to be its subject matter): (i) the interpretation he assigns to u may be entailed by C_{SM} — in this case the information carried by u is old for him; (ii) the interpretation may contradict C_{SM} and (iii) it may be that the interpretation neither contradicts C_{SM} nor is entailed by it. So far we just counted cases (ii) and (iii) both as cases of new information, but as the remarks above should have made clear, from the recipient's point of view they are really very different.

This tripartite distinction is important not only in connection with the content of the statements speakers make, but also with the anaphoric and other presuppositions their statements carry. Suppose that A has just made the statement that 'Dorothea has just gone to Paris again.' This choice of words introduces the presupposition that Dorothea has been to Paris before. What is the interpreter to do with this presupposition? Again this will depend on how the presupposition is related to C_{SM} , and once again we have to distinguish between three possibilities — (i) C_{SM} entails the presupposition, (ii) the presupposition contradicts it, and (iii) neither. The first case, (i), is usually treated as the normal one: the antecedent information about the subject matter entails the presupposition; that is as it should be, and the interpreter can, detaching the presupposition, move straightaway to the non-presuppositional component of his interpretation. Cases of type (iii) are the ones that we believe presupposition theorists usually think of when they discuss accommodation. If a presupposition cannot be verified as following from the 'context', i.e., from what the interpreter currently knows or assumes, then he will adjust the context — accommodate it, as the technical vocabulary has it — so that it does entail the presupposition.⁸⁰

79. Assuming the discourse is a cooperative one. For a different perspective, cf., Merin (1999); van Rooij (2004a).

80. By and large accommodation comes easily. In fact, it is one of the classical observations of presupposition theory that speakers will often exploit the readiness with which interpreters accommodate presuppositions by choosing wordings for their statements which trigger presuppositions of which they do not think that their interlocutors believe them already, but that they want them to adopt. And usually the ploy works: the interpreter will take the content of the presupposition on board much as he would have done if it had been asserted. In such cases, where the interlocutor reacts in accordance with the speaker's expectations, the effect on his beliefs will be the same as it would have been if the speaker had made the presupposition into a separate statement followed by the statement he actually made. The term 'accommodation', in the present technical use of it, can be traced back to Lewis (1979). Some of Lewis' remarks can

And then there still is the second case, in which the presupposition contradicts C_{SM} . Once again there are several ways in which the interpreter could react: he could conclude that since the speaker was making a statement with this presupposition, she must have known the presupposition to be true, and adjust his own beliefs to fit; or he may conclude that the speaker's apparent belief in the truth of the presupposition is wrong; and in that case there are once again several options; he may try to point her mistake out to her or he may let the matter rest and take the non-presuppositional content of her statement as if no presupposition had been attached to it.

We note that these same observations also apply to presuppositions that are generated in speech acts other than assertions, such as questions or directives. Such non-assertive speech acts were mentioned in passing in the introduction, but since then nothing was said about them. This seems a suitable point to return to them. Our lack of attention to non-assertive speech acts throughout the chapter should not be construed as showing our lack of awareness of the crucial part they play in the normal use of language. (In any theory of the semantics and pragmatics of conversation their analysis is absolutely indispensable.) Rather it is a reflection of our conviction that in an account of linguistic information there is no need to make them the topic of a separate discussion. For by and large the information content of non-assertive utterances is determined according to the same principles as it is for assertions. The only difference is that non-assertive speech acts put their information content to different uses than assertions do. For the question what the information content of an expression is and how it gets transmitted that difference seems immaterial.

But the situation is different with regard to presuppositions. The presuppositions connected with non-assertive utterances have the same status as those connected with assertions. In either case their content must be verifiable in the context before the utterance can be accepted as a legitimate transmitter of its message. One consequence of this is that presuppositions are more markedly set aside from the non-presuppositional part of utterance content in the case of non-assertive speech acts than they are in the case of assertions. It is for this reason that non-assertive speech acts are particularly useful as presupposition tests: For instance, a presupposition triggered by a constituent of an interrogative sentence used in a yes-no question will often manifest itself more clearly as a presupposition there than it does in relation to an assertion involving the corresponding indicative sentence. Especially striking are those cases where you, the addressee of the question, think that the presupposition is false. It won't feel right to you to answer the question with either 'yes' or 'no'. For instance, suppose I ask you the question in (6a) and you know (i) that Fred did come to the session last night, but (ii) that he wasn't at any of the previous sessions. It wouldn't seem right for you in such a case to simply reply with 'yes', since that would imply that for all you know the contribution made by 'again' in (6a) — that Fred came to one or more earlier sessions — is true. Rather, you would feel it incumbent upon you to point out (in the words of

be read as suggesting that accommodation is always possible, but in the meantime we have learnt that there are presuppositions for which accommodation is subject to certain constraints (although it appears that the question which accommodation constraints apply to which presuppositions is still largely unanswered). Cf., Beaver & Zeevat (2007).

(6b), say) that the presupposition was false, perhaps adding, after having got this matter out of the way, that as a matter of fact Fred was present at last night's session.

6. (a) Was Fred at the reading group again last night?
(b) Well as a matter of fact he had never been there before. But yes, he was there yesterday.

In the above we have emphasised the importance of true information. The importance, we saw, shows up both in connection with presuppositional and with non-presuppositional information. Failure of a presupposition puts the whole communication process in jeopardy, something that can be observed with particular clarity in the case of non-assertive speech acts. But, as we also noted, truth is just as important for non-presuppositional content, in particular the content of assertions. (It is part of the conventions associated with speech acts of that kind, we observed, that the speaker commits herself to their content being true.) None of this is really surprising. For what people need and want first and foremost is true information about their world — information that makes it possible for them to plan their actions, by enabling them to make predictions about the consequences that the different lines of action open to them might have.

But none of this should blind us to the fact that it is nevertheless linguistic information as we have defined it — information about how the world might be, rather than information about how it actually is — that is the central notion in relation to human language; it is this kind of information that is language's principal commodity, not the kind of information that has truth built into it. One indication of this is that all we have said about the interpreter's handling of both non-presuppositional and presuppositional content that is motivated by the concern for truth is ultimately not about the truth as such but about what the interpreter thinks is true. It is because the interpreter can represent the world as being of a certain kind, and thus imagine it to be of that kind, that he is also capable of thinking that it is of that kind. But in the case of thought, as in that of language, the commitment to the world actually being of a certain kind is distinct and detachable from the conception of a world of such a kind as such. This distinction — between truth and mere possibility, or, if you prefer, between belief and imagination — is at the core of information both as a cognitive and a linguistic commodity.

The detachability of truth from linguistic information content comes into particularly clear view when we compare discourse about the real world with fiction. When we read a novel or listen to a story we assign information content to the words we hear or see in much the same way as we do when interpreting utterances that we take to be about the real world. By and large the same principles of interpretation apply, including those which regulate the resolution of anaphora or the contextual justification of presuppositions. But there is nevertheless one crucial difference: since the world that the fiction unfolds presents itself as one that is at the author's disposition, the interpreter has no basis for objecting to any bit of non-presuppositional or presuppositional content. For that would require detection of a conflict with what he knows (or thinks he knows) to be true on independent grounds, and in this case it is true by definition that there can be no independent grounds. (At best the interpreter could detect internal inconsistencies in the story or violations of the basic laws and regularities which any world, real

or imagined, should obey.) In fictional discourse we see language at work as a means for providing pure information content, untrammelled by the concern that the world described might prove different from the world whose description is intended.

7 Prospects and challenges

In this chapter the focus has been on the information conveying function of natural languages. We have seen that the concepts of meaning and information have been related throughout much of the history of modern linguistics, though not always to the same extent and in the same sense. We have also described in some detail what kinds of concepts and ideas are needed to build a descriptively satisfactory and theoretically sound theory that models the information conveying capabilities of natural languages. To be sure, such a theory is still ‘in the works’, and no doubt other concepts and ideas will be needed for it to be developed further, but the basic contours of what natural language information is and how it is shaped by syntactic structure, the structure of the utterance context, the discourse context and the participants’ doxastic states and strategic goals, seem reasonably clear. In what follows we mention a few current trends, and then end this chapter with some thoughts on the information exchange function of natural languages.

One aspect that currently gets a lot of attention relates to the strategic goals that language users have when they enter into a conversation, read a text, or communicate linguistically in some other form. Simple information exchange as such is hardly ever the ultimate goal: information is needed for certain purposes, e.g., in order to decide which action to undertake oneself, or to predict or influence actions of others, to explore possible courses of events, and so on. Language use then becomes part of a solution of a decision problem, and understanding the nature of the problem is essential since it determines what information (in terms of content and/or amount) is relevant in a given situation.⁸¹

Situations of information exchange become even more complex once one acknowledges that not only providing information, but also withholding information, or divulging information selectively (part of the information, to part of the other participants), may be a crucial element of an overall strategy. This type of information exchange is often analysed using tools from game theory.⁸² So-called ‘higher order effects’ of information disclosure are a central topic here: if *A* tells *B* that *p*, then *B* potentially learns a lot of other things beside *p*: that *A* believes (knows) that *p*, that *A* wants *B* to believe (know) that *p*, that *C*, who happened to overhear the utterance, now also believes (knows) that *p* but not as a result of an intentional action of *A*, and so on.

A particular aspect of this problem set that has been studied quite extensively is how language users choose means of expressions, as speakers, and determine interpretations, as hearers. Given the fact that the relation between expression and content, form and meaning, is not one–one, but in general many–many, the problem how to express

81. Cf., Ginzburg (1995), van Rooij (2003).

82. Cf., van Benthem (2006).

certain information, and how to decide what content a certain expression is used to convey, is a substantial one. In order to solve these problems, language users need general principles that they can use themselves and that they can assume the other users employ as well. Gricean pragmatics provides a first and partial indication. So-called ‘bi-directional optimality theory’ aims to generalise and systematise these ideas.⁸³ Although it greatly enhances the explanatory power of the classical Gricean framework, it stays within that framework in that it relies on an independent characterisation of the space of possible forms and the space of possible meanings. A framework in which meanings are not a precondition but a result of linguistic exchanges is provided by so-called ‘signalling games’,⁸⁴ which in that respect constitute a major step away from the classical model.

The integration of the various theories and paradigms that are currently being explored and that we can only mention, is still very much an open matter. No unified framework exists as yet, and developments are rapid. However, most approaches somehow build on the general principles that we have outlined in this chapter, which suggests that the current phase of diversification could be followed by one in which more comprehensive theories can be developed.

A question that we haven’t addressed so far is to what extent information conveying is really natural language’s ‘core business’, as many would claim. The reason we have not gone into that discussion is minimally this, that although not everyone agrees that information conveying is the function of natural language ‘par excellence’, no-one really wants to deny that it is one of the things natural languages are used for, and the how and why of that is really what this chapter is all about. Nevertheless, as a final reminder it may not be superfluous to indicate, albeit only very briefly, the possible limitations of this view on natural language.

First of all, in as much as the concepts and formal machinery that are put to use in theories of natural language semantics and pragmatics are taken from logic and theoretical computer science, where they were developed with the explicit purpose of providing tools for the description and analysis of processes of information exchange, it is hardly surprising that the resulting theories make natural language, too, appear as primarily concerned with that specific goal. Anything that doesn’t fit simply disappears from sight, by being ‘abstracted away’ from. That by itself doesn’t mean, of course, that there actually is something that doesn’t fit, or that, if there is, it is of importance. But the fact remains that the tools used in the study of natural language are derived from the domain of formal systems and that whereas the latter are straightforwardly designed with a specific purpose in mind, the former can hardly be said to answer to such a description. So at least we need to allow the possibility of a certain one-sidedness, and concomitant distortion.

A second consideration pertains to the status of the use of natural language for information exchange. Of course nobody would deny that natural languages serve other

83. Cf., Dekker & van Rooij (2000), Blutner et al. (2006). For a game-theoretical approach to the issue of ambiguity resolution, cf. Parikh (2002).

84. Cf., van Rooij (2004b).

purposes as well: we flatter and comfort (each other and ourselves), we sing songs and write poetry, we congratulate and curse. There is no denying that such utterances, too, convey information, if only of the ‘higher order’ type just mentioned, but according to many it would be an unjustified generalisation to state that their purpose is that of exchanging information. The real issue, they feel, is how these other uses and the information exchange use are weighed relative to each other. Is information exchange the core function of language, with other functions being somehow dependent on it? Or are the various uses we make of language relatively independent from each other? Or is there some other function than information exchange that is the primary one? These are difficult questions, and it would take us too far afield to discuss the various options that have been proposed and defended in the literature. Some seek the answer in an (often speculative) account of language evolution: did language evolve from the use of signals, to indicate foods, predators, and such?⁸⁵ Or are its origins rather to be found in a need for social cohesion, and did it start out as a way of maintaining social bonds within a group?⁸⁶ Others address these questions from a more systematic point of view. As a matter of fact, the history of western thinking about language displays quite a variety of opinions about what constitutes its inner nature, and the logical one, with its emphasis on reference, description and information exchange is but one of them. In modern times, the ‘information oriented’ views of, e.g., Locke and Leibniz, were balanced by, e.g., those of Rousseau, who saw the essence of language in the expressions of the passions, and of Herder and Humboldt, who emphasised its expressive role with regard to the spirit of a culture. In the twentieth century Wittgenstein explored the variety of the uses we make of language and emphasised their ‘co-originality’, and Austin and Searle further systematised certain elements of this view. And in other philosophical traditions, too, people have expressed yet other views on the nature of language, such as the hermeneutic perspective of Heidegger and Gadamer, the moral-political view of Habermas,⁸⁷ or the phenomenological one of Merleau-Ponty.⁸⁸

What these alternative views have in common is that they reject, to some extent at least, the instrumentalism that seems inherent in the information exchange perspective. There language essentially is an instrument, a tool that is put to a use, viz., that of asking for and providing information. Again, the origins of the concepts and tools of modern semantics and pragmatics are conducive to such a view: in formal systems the languages and their semantics are defined according to predetermined specifications, which makes them instrumental through and through. In many of the alternative views, the distinction between the instrument and the use to which it is put is less clear, more difficult to make. One might say that the opposition really is a matter of whether ‘language has use’ or ‘language is use’.

These observations, of course, merely scratch the surface of a very complicated, and still ongoing debate. We draw attention to them merely in order to balance the

85. Cf., Pinker & Bloom (1990).

86. Cf., Dunbar (1998). And there are many other views, cf., various contributions in the already referred to Christiansen & Kirby (2003).

87. Cf., Lafont (1999).

88. Cf., Edie (1987).

perspective, not to throw serious doubts on the view that information exchange is a vital function of natural languages. The latter point is uncontroversial, as is the contention that theories exploring this perspective provide a real insight into the nature of natural languages, and have furthered our understanding of its structure, its meaning and its use immensely. The point to bear in mind is just that other perspectives, in their own way, contribute to such an understanding as well.

References

- Almog, Joe, Perry, John, & Wettstein, Howard K., eds. 1989. *Themes from Kaplan*. Oxford University Press.
- Asher, Nicholas & Lascarides, Alex. 2003. *Logics of Conversation*. Cambridge University Press, Cambridge.
- Austin, John Longshaw. 1962. *How To Do Things With Words*. Oxford University Press, Oxford.
- Baggio, Giosuè, van Lambalgen, Michiel, & Hagoort, Peter. to appear. Language, linguistics and cognition. In: Groenendijk, Jeroen & Stokhof, Martin, eds., *Handbook of Philosophy of Linguistics*. Elsevier, Amsterdam.
- Bartsch, Renate. 1996. The myth of literal meaning. In: Weigand, Edda & Hundsnurscher, Franz, eds., *Lexical Structures and Language Use*, pp. 3–16. Niemeyer, Tübingen.
- Barwise, Jon & Perry, John. 1983. *Situations and Attitudes*. MIT Press, Cambridge, Mass.
- Barwise, Jon & Seligman, Jerry. 1997. *Information Flow. The Logic of Distributed Systems*, *Cambridge Tracts in Theoretical Computer Science*, vol. 44. Cambridge University Press, Cambridge.
- Beaver, David I. 2001. *Presupposition and Assertion in Dynamic Semantics*. CSLI, Stanford.
- Beaver, David I. & Zeevat, Henk. 2007. Accommodation. In: Ramchand, Gillian & Reiss, Charles, eds., *The Oxford Handbook of Linguistic Interfaces*. Oxford University Press, Oxford.
- van Benthem, Johan F.A.K. 2006. One is a lonely number. In: Chatzidakis, Z., Koepke, P., & Pohlers, W., eds., *Logic Colloquium '02, Lecture Notes in Logic*, vol. 27. Association for Symbolic Logic.
- Biletzki, Anat & Matar, Anat, eds. 1998. *The Story of Analytic Philosophy: Plots and Heroes*. Routledge, London.
- Blutner, Reinhard, de Hoop, Helen, & Hendriks, Petra. 2006. *Optimal Communication*. CSLI, Stanford.
- Bruner, J. 1983. *Child's Talk. Learning to Use Language*. Norton, London.
- Burge, Tyler. 1990. Wherein is language social? In: Anderson, C. Anthony & Owens, Joseph, eds., *Propositional Attitudes*, pp. 113–30. CSLI, Stanford.
- Cardona, George. 1988. 2nd ed 1997. *Pāṇini. His Work and Its Tradition*. Motilal Banarsidass, Delhi.
- Chomsky, Noam. 1955. Logical syntax and semantics: Their linguistic relevance. *Language*, 31(1), 36–45.

- . 2005. Three factors in language design. *Linguistic Inquiry*, 36(1), 1–22.
- Christiansen, Morton H. & Kirby, Simon, eds. 2003. *Language Evolution*. Oxford University Press, Oxford.
- Cresswell, Max J. 1973. *Logics and Languages*. Methuen, London.
- Davidson, Donald. 1967. Truth and meaning. *Synthese*, 17, 304–23.
- . 1974. On the very idea of a conceptual scheme. *Proceedings and Addresses of The American Philosophical Association*, 47.
- . 1984. *Inquiries into Truth and Interpretation*. Oxford University Press, Oxford.
- . 1986. A nice derangement of epitaphs. In: LePore, Ernest, ed., *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson*, pp. 433–46. Blackwell, Oxford. Reprinted in Davidson (2005).
- . 2005. *Truth, Language and History*. Clarendon Press, Oxford.
- Dekker, Paul & van Rooij, Robert. 2000. Bi-directional optimality theory: An application of game theory. *Journal of Semantics*, 17, 217–42.
- Dummett, Michael. 1996. *The Origins of Analytical Philosophy*. Duckworth, London.
- . 2004. *Truth and the Past*. Columbia University Press.
- Dunbar, Robin I.M. 1998. *Grooming, Gossip and the Evolution of Language*. Harvard University Press, Cambridge, Mass.
- Edie, James M. 1987. *Merleau-Ponty's Philosophy of Language*. University Press of America, Washington.
- van Eijck, Jan & Stokhof, Martin. 2006. The gamut of dynamic logic. In: Gabbay, Dov & Woods, John, eds., *Handbook of the History of Logic, Volume 6 – Logic and the Modalities in the Twentieth Century*, pp. 499–600. Elsevier, Amsterdam.
- Farkas, Katalin. 2006. Semantic internalism and externalism. In: LePore, Ernest & Smith, Barry, eds., *The Oxford Handbook of Philosophy of Language*. Oxford University Press, Oxford.
- Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press, Cambridge, Mass.
- Frege, Gottlob. 1879. *Begriffsschrift. Eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Louis Nebert, Halle a.S. English translation in van Heijenoort 1970.
- . 1882. Über die wissenschaftliche Berechtigung einer Begriffsschrift. *Zeitschrift für Philosophie und philosophische Kritik*, pp. 48–56. English translation in Frege (1964).
- . 1918–19. Der Gedanke: eine logische Untersuchung. *Beiträge zur Philosophie des deutschen Idealismus*, 2, 58–77. English translation in Frege (1977).
- . 1964. On the scientific justification of a concept-script. *Mind*, 73, 155–60. Translated by J.M. Bartlett.
- . 1977. *Logical Investigations*. Blackwell, Oxford. Translated by P.T. Geach.
- Garton, A.F. 1992. *Social Interaction and the Development of Language and Cognition*. Erlbaum, Hillsdale.
- Ginzburg, Jonathan. 1995. Resolving questions, I & II. *Linguistics and Philosophy*, 18(5 & 6), 459–527 & 567–609.
- Grice, H.P. 1957. Meaning. *The Philosophical Review*, 66, 377–88.
- . 1975. Logic and conversation. In: Cole, Peter & Morgan, Jerry, eds., *Syntax and Semantics. Volume 3: Speech Acts*, pp. 41–58. Academic Press, New York.

- Groenendijk, Jeroen & Stokhof, Martin. 1984. *On the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, Department of Philosophy, Universiteit van Amsterdam, Amsterdam. (Available from: <http://dare.uva.nl/en/record/123669>).
- . 1990. Dynamic Montague grammar. In: Kálmán, László & Pólos, László, eds., *Papers from The Second Symposium on Logic and Language*, pp. 3–48. Akadémiai Kiadó, Budapest.
- . 1991. Dynamic predicate logic. *Linguistics and Philosophy*, 14(1), 39–100.
- . 1997. Questions. In: van Benthem, Johan & ter Meulen, Alice, eds., *Handbook of Logic and Linguistics*, pp. 1055–1124. Elsevier/MIT Press, Amsterdam/Cambridge Mass.
- . 2005. Why compositionality? In: Carlson, Greg & Pelletier, Jeff, eds., *Reference and Quantification: The Partee Effect*, pp. 83–106. CSLI, Stanford.
- Groenendijk, Jeroen, Stokhof, Martin, & Veltman, Frank. 1996. Coreference and modality. In: Lappin, Shalom, ed., *Handbook of Contemporary Semantic Theory*, pp. 179–213. Blackwell, Oxford.
- Hacking, Ian. 1975. *Why Does Language Matter to Philosophy?* Cambridge. Cambridge University Press.
- van Heijenoort, Jean, ed. 1970. *Frege and Gödel. Two Fundamental Texts in Mathematical Logic*. Harvard University Press, Cambridge, Mass.
- Heim, Irene. 1983. File change semantics and the familiarity theory of definiteness. In: Bäuerle, Rainer, Schwarze, Christoph, & von Stechow, Armin, eds., *Meaning, Use, and Interpretation of Language*. De Gruyter, Berlin.
- Hendriks, Petra & de Hoop, Helen. 2001. Optimality theoretic semantics. *Linguistics and Philosophy*, 24(1), 1–32.
- Hintikka, Jaakko. 1983. *The Game of Language*. Reidel, Dordrecht.
- Jackendoff, Ray. 1990. *Semantic Structures*. MIT Press, Cambridge, Mass.
- Kamp, Hans. 1981. A theory of truth and semantic representation. In: Groenendijk, Jeroen, et al., eds., *Formal Methods in the Study of Language*. Mathematical Centre, Amsterdam.
- . 1984–1985. Context, thought and communication. *Proceedings of the Aristotelian Society*, pp. 239–61.
- . 1990. Prolegomena to a structural account of belief and other attitudes. In: Anderson, C.A. & Owens, J., eds., *Propositional Attitudes: The Role of Content in Logic, Language, and Mind*. CSLI, Stanford.
- Kamp, Hans & Reyle, Uwe. 1993. *From Discourse to Logic*. Kluwer, Dordrecht.
- Kaplan, David. 1989. Demonstratives. In: Almog, Joe, Perry, John, & Wettstein, Howard K., eds., *Themes from Kaplan*, pp. 481–563. Oxford University Press, Oxford.
- Kirby, Simon. 2000. Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners. In: Knight, C., ed., *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, pp. 303–23. Cambridge University Press, Cambridge.
- Lafont, Christina. 1999. *The Linguistic Turn in Hermeneutic Philosophy*. MIT Press, Cambridge, Mass.

- Lakoff, George. 1987. *Women, Fire and Dangerous Things*. The University of Chicago Press, Cambridge.
- Levinson, Stephen C. 1983. *Pragmatics*. Cambridge University Press, Cambridge.
- Lewis, David K. 1970. General semantics. *Synthese*, 22, 18–67.
- . 1979. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–59.
- McGinn, Colin. 1989. *Mental Content*. Blackwell, Oxford.
- Merin, Arthur. 1999. Information, relevance and social decisionmaking. In: Moss, Larry, Ginzburg, Jonathan, & de Rijke, Maarten, eds., *Logic, Language and Computation*. Vol. 2. CSLI, Stanford.
- Montague, Richard. 1970a. English as a formal language. In: Visentini, Bruno, ed., *Linguaggi nella Società e nella Tecnica*, pp. 189–224. Edizioni di Comunità, Milano.
- . 1970b. Universal grammar. *Theoria*, 36, 373–98.
- . 1973. The proper treatment of quantification in ordinary English. In: Hintikka, Jaakko, Moravcsik, Julius, & Suppes, Patrick, eds., *Approches to Natural Language*, pp. 221–42. Reidel, Dordrecht.
- . 1974. *Formal Philosophy. Selected papers of Richard Montague. Edited and with an Introduction by Richmond H. Thomason*. Yale University Press, New Haven and London.
- Neale, Stephen. 1990. *Descriptions*. MIT Press, Cambridge, Mass.
- Parikh, Prashant. 2002. *The Use of Language*. CSLI, Stanford.
- Partee, Barbara H. 1973. Some transformational extensions of Montague grammar. *Journal of Philosophical Logic*, 2, 509–34.
- Peregrin, Jaroslav & von Heusinger. 2004. Dynamic semantics with choice functions. In: Kamp, Hans & Partee, Barbara H., eds., *Context Dependence in the Analysis of Linguistic Meaning*, pp. 255–74. Elsevier, Amsterdam.
- Pinker, Steven & Bloom, Paul. 1990. Natural language and natural selection. *Behavioral and Brain Sciences*, 13(4), 707–84.
- Plato. 1921. *Sophist*. Loeb Classical Library. Harvard University Press, Cambridge, Mass. Translated by H.N. Fowler.
- Pustejovsky, James. 1995. *The Generative Lexicon*. MIT Press, Cambridge, Mass.
- Putnam, Hilary. 1975. The meaning of ‘meaning’. In: *Mind, Language, and Reality*, pp. 215–71. Cambridge University Press, Cambridge.
- Quine, Willard Van Orman. 1948. On what there is. *Review of Metaphysics*, 2, 21–38. Reprinted in Quine (1953a).
- . 1953a. *From a Logical Point of View*. Harper & Row, New York.
- . 1953b. Two dogmas of empiricism. In: *From a Logical Point of View*, pp. 20–46. Harvard University Press, Cambridge, Mass.
- Recanati, Francois. 2004. *Literal Meaning*. Cambridge University Press, Cambridge.
- Robins, R.H. 1990. *A Short History of Linguistics*. Longman, London, 3rd edn.
- van Rooij, Robert. 2003. Asserting to resolve decision problems. *Journal of Pragmatics*, 35.
- . 2004a. Cooperative versus argumentative communication. In: Rebuschi, Manuel & Tulenheimo, Tero, eds., *Logique & Théorie des Jeux, Philosophia Scientiae*, vol. 8. Kimé, Paris.

- . 2004b. Signalling games select Horn strategies. *Linguistics and Philosophy*, 27, 492–527.
- Rothstein, Susan. 2004. *Structuring Events: A Study in the Semantics of Lexical Aspects, Explorations in Semantics*, vol. 2. Blackwell, Oxford.
- van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics*, 9(4), 333–77.
- Searle, John R. 1969. *Speech Acts*. Cambridge University Press, Cambridge.
- Seuren, Pieter A.M. 1998. *Western Linguistics: An Historical Introduction*. Blackwell.
- Soames, Scott. 2003. *Philosophical Analysis in the Twentieth Century, Volumes 1 and 2*. Princeton University Press, Princeton.
- Sperber, Dan & Wilson, Deirdre. 1995. *Relevance: Communication and Cognition*. Blackwell, Oxford, 2nd edn.
- Stalnaker, Robert. 1974. Pragmatic presuppositions. In: Munitz, M. & Unger, P., eds., *Semantics and Philosophy*. New York University Press, New York.
- . 1979. Assertion. In: Cole, Peter, ed., *Syntax and Semantics 9 – Pragmatics*. Academic Press, New York.
- Stanley, Jason. 2005. Semantics in context. In: Preyer, Gerhard & Peter, Georg, eds., *Contextualism in Philosophy*. Oxford University Press, Oxford.
- Steiner, George. 1975. *After Babel. Aspects of Language and Translation*. Oxford University Press, Oxford.
- Stenius, Eric. 1967. Mood and language-game. *Synthese*, 17, 254–74.
- Stokhof, Martin. 2002. Meaning, interpretation and semantics. In: Barker-Plummer, Dave, et al., eds., *Words, Proofs and Diagrams*, pp. 217–40. CSLI, Stanford.
- Sundholm, Göran. to appear. A century of judgement and inference: 1837–1936. In: Haaparanta, L., ed., *The Development of Logic*. Oxford University Press, Oxford.
- Szabó, Zoltán Gendler, ed. 2005. *Semantics versus Pragmatics*. Oxford University Press.
- Talmy, Leonard. 2000. *Towards a Cognitive Semantics*. MIT Press, Cambridge, Mass.
- Tomasello, Michael. 1999. *The Cultural Origins of Human Cognition*. Harvard University Press, Cambridge, Mass.
- . 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Harvard University Press, Cambridge, Mass.
- Veltman, Frank. 1996. Defaults in update semantics. *Journal of Philosophical Logic*, 25, 221–61.
- Wittgenstein, Ludwig. 1958. *Philosophical Investigations*. Blackwell, Oxford, 2nd edn.
- . 1960. *Tractatus Logico-Philosophicus*. Suhrkamp, Frankfurt a/M.