

Representationalism and Linguistic Knowledge

Ronnie Cann, Ruth Kempson & Daniel Wedgwood

September 30, 2010

Abstract

1 Positing representations

In the analysis of natural language phenomena, linguistic theories typically have recourse to *representations* of one form or another.¹ Different types of representation are often posited as a means of generalising over aspects of form or interpretation as displayed in natural language constructions, and these are frequently invested with considerable theoretical significance. There are proposed representations of structure at all levels of linguistic systems: sounds, words, sentence strings, as well as representations of the meanings of words, sentences in the abstract and uttered sentences, and even representations of other people's intentions. Such a *representationalist* stance was firmly set in place by Chomsky (1965) as part of, indeed the central core of, cognitive science, with language defined as a system of principles for correlating phonological representations (on some abstraction from phonetics) with some representation of interpretation (on some abstraction from denotational contents), via mappings from a central syntactic system. In such an approach, more than one *level* of representation may be posited as interacting in different ways with other types of representation, for example deep structure and surface structure levels of syntax of Chomsky (1965) were taken to interact in different ways with other types of representation, in particular semantics and phonology.

Chomsky's move towards the explicit representation of linguistic properties as part of human cognition came to be assumed almost universally within theoretical linguistic frameworks, whether formally characterised or

¹This chapter in part reports work to which many people have contributed. We thank in particular, Eleni Gregoromichelaki, Wilfried Meyer-Viol, Matthew Purver for essential input to these developments. We also thank Andrew Gargett, Stergios Chatzikyriakidis, Peter Sutton, Graham White and many others for comments both over the years and in the preparation of this chapter. Research for this paper includes support from the Leverhulme Trust to the first author (XXX), and from the ESRC to the second author (YYY).

not. But all details remain controversial, as there are no *a priori* constraints on the number of levels or types of representations that may be posited to account for natural language phenomena nor on the modes of interaction between them. A multiplicity of types and levels of representation, however, threatens to result in a characterisation of language as merely a composite of different factors for which no integrated systematic underpinning is offered. Different types of representation (phonological, morphological, syntactic, semantic and so on) are typically associated with different types of data structures, with mapping relations needing to be defined between the different representation systems or between different levels within such systems and levels of representation in other systems. Such accounts may be descriptively successful but, their potential complexity aside, hardly provide an explanation of language as a system, for distinct types of representation may be posited *a priori* for all languages, suggesting that some central organisational principle is somehow being missed.

With this worry in mind, the question raised by the postulation of representations of one type or another is what such talk of representations amounts to. A *representationalist* account, according to linguists' use of that term, is one that involves essential attribution of structure intrinsic to the phenomenon under investigation. To say that the characterisation of the human capacity for language requires representations is not to establish just which systems of representation should be posited. The issue is then what the attribution of structure consists in – what underpins it and what it entails – and how many distinct types of structure have to be invoked.

First, and relatively uncontroversially, the characterisation of the sound system of a language as expressed through a phonological framework demands a system of representations that is distinct from that warranted by structural properties of language as expressed in their syntax/semantics (though issues of representation arise in phonology too: see Carr this volume). Phonology apart, the primary debate over the types of representation to be invoked takes place against two background assumptions. The first, derived principally from the structuralist practices of the early twentieth century (although with roots in the Graeco-Roman grammatical tradition), is that natural language strings exhibit syntactic structure as defined over words and the expressions that they make up within some string. This assumption is adopted almost without caveats in the major current linguistic formalisms: Categorical Grammar (e.g. Steedman 1996, 2000, Morrill 1994, 2010) Head-driven Phrase Structure Grammar (e.g. Sag, Wasow and Bender 2003), Minimalist versions of Transformational Grammar (e.g. Chomsky 1995, Hornstein et al 2005), Lexical Functional Grammar (e.g. Bresnan 2001). The second is that such syntactic structures have to be defined independently of the characterisation of meaning for natural language. This second assumption is far from universally adopted (in particular not within categorical grammar), but the presumed independence of syntax and seman-

tics is nonetheless widespread. Part of the debate about representations in linguistics is thus about whether that semantic characterisation itself has intrinsic structural properties that warrant a level of representation independent of syntactic structure.²

This aspect of the representationalism debate emerges out of an apparent incommensurability between *natural languages* and the *formal languages* used to articulate logics (the language of *predicate logic* in particular), whose grammatical properties constitute the starting point for defining the representations that are used to characterise the properties of natural languages. The problem lies in the fact that natural languages are endemically context-dependent in their construal. Each natural language has a stock of words and expressions which would seem to provide some stable input to the process of construal, but which nevertheless in all usages allow a vast array of interpretations, determined in some sense from the context in which an utterance occurs. This phenomenon does not arise with formal languages, an asymmetry which was initially analysed as somewhat peripheral, a matter to be stipulated but not of central importance. However, increasingly, theoreticians have been addressing the issues of context-dependence (see sections 4-5 below); and formal models of language are now seeking to address this core phenomenon of natural languages directly. As we shall see, it is in characterising this aspect of natural-language construal that structural aspects of interpretation seem to be needed that are distinct from what is made available by the conventional concepts of syntax.

This tension between the context-dependence of natural language construal and the context-independent interpretation of formal languages, although in principle logically distinct from any debate about the need for representations of linguistic phenomena in general, has nevertheless become intrinsically linked to that debate. Firstly, as noted above, this is because the general methodology has been to use some type of formal system to express the properties of syntactic relations which are interpreted through some decontextualised, denotational semantic machinery (Montague 1974) or through some equally decontextualised ‘Conceptual/Intentional’ interface

²We note in passing that this is not the only characterisation of *representationalism*. There are *inferentialist* approaches to natural language understanding, hence semantics, according to which what is taken as basic are the inferential processes associated with language use. The representational aspect of language, in the sense that expressions represent their attributed denotational content, is claimed to be wholly derivative on such practices (Brandom 1994, 2008). On this view, the term *representationalism* applies more broadly than just to those linguistic models which invoke a semantic level of representation over and above syntactic structure; for Brandom, the term applies to ALL linguistically based forms of explanation in so far as in all of these a primitive notion of representation is induced as an intrinsic property of natural language strings. On the Brandom view any such assumption of representations is seen as reducible to patterns of usage established through inferential practices involving assertion, justification, etc. Since the debate between Brandom and his adversaries takes place wholly within philosophy of language, we do not adopt this terminology here.

(Chomsky 1986, 1995; Jackendoff 2002). Secondly, it results from the move instigated in Chomsky (1965) that knowledge of language ('competence') should be isolated from all applications of this knowledge to parsing or production ('performance'). This separation imposes a methodology in which the primary data of linguistics are not that of language as used in context, but of judgements of the grammaticality of sentences without any reference to context or the dynamics of real-time language activities.³ This methodology of taking sentences rather than utterances as the core language data, with grammars constructed that bear no relation to performance considerations, has been adopted by linguists across different theoretical persuasions.⁴

In this chapter, we seek first to establish in a pre-theoretic way why representations of some sort have to be invoked to characterise the properties of natural languages, and, more particularly, why representations of meaning attributed to an uttered string have to be invoked. We then turn to the formal language concept of grammar in which syntax is by definition a set of principles for inducing all and only the set of sentence-strings of the language. In applying this methodology to natural language grammar construction, grammars are defined to be systems for inducing structure for sentence-strings, so that the natural-language grammar too could be seen as analogously inducing all and only the set of grammatical strings of the language. The making of the competence-performance distinction is essential here, in order to isolate a concept of linguistic knowledge commensurate with the formal language concept. Indeed, it has been assumed, broadly following the formal-language methodology, that grammars for languages comprise a set of principles that pair strings with structures inhabited by those strings, and that from those structures mappings are definable onto interpretations, thus capturing all and only the grammatical sentences of the language, paired appropriately with sentence-meanings. This, it is almost universally agreed, is the syntactic heart of the human capacity for language as expressed in a grammar, with performance factors totally excluded.

Once we have sketched this formal-language background, we will set out the arguments as to whether a distinct type of representation is warranted in addition to string-structure pairings, namely representations of such meanings as are expressible by that string. As we shall see, it is increasingly central to explanations of meaning for natural language expressions to characterise the way in which interpretation of a string relates to the context

³So-called 'intuitions' of native speakers whose precise nature is unclear and controversial, are problematic particularly as apparent categorical judgements of (un)grammaticality, the bedrock of this sort of linguistic investigation, may be manipulated in context to acceptability (see, for example, Cann, Kempson and Kaplan 2005 for a discussion of resumption in English relative clauses). See Baggio et al this volume for detailed discussion and critical evaluation.

⁴It is, as we shall see, a primary cause of the apparent need to posit multiple types of representation.

in which it is uttered. This has proved particularly important in attempts to articulate integrated explanations of anaphora and ellipsis, and it is here that representations of structure specific to interpretation are usually invoked. But in this exploration of context-dependence, the sentence remit of the grammar and the total exclusion of the dynamics of natural-language performance will be seen to be problematic.

Finally, we shall turn to claims that a characterisation of the dynamics of how utterance understanding is incrementally built up in context is not merely essential to the explication of natural-language interpretation, but turns out to be a possible grounding of syntactic explanations of natural language strings, reversing the generally presumed order of dependence. Moreover, according to these arguments, it is only by making this shift to a more dynamic perspective for natural-language grammars that an integrated account of context-dependent phenomena such as anaphora and ellipsis becomes possible. On this view, it is the concept of syntactic structure independent of semantic characterisation that turns out to be epiphenomenal. So finally we shall come to the conclusion that only a minimum duo of types of representations is required after all – a pairing of structured representations of sound with structured representations of meaning, as made available by a system of growth of such representations. But this conclusion is made available only by a shift in perspective to one in which the dynamics of language processing is directly reflected in the grammar itself. In this respect, the conventional notion of competence, and its associated methodology, has to be replaced by one in which the articulation of linguistic knowledge is much more closely connected to the uses to which it is put.

2 Representationalism in Linguistic Theory

Capacity for language is seen as the definitive mental activity of human beings. Humans manipulate sequences of sounds of a language that they understand with the clear presumption of these being interpretable relative to some rule-governed system, such that one agent can apparently succeed in passing on information to another, request clarification of another, induce activity on the part of another, and so on. This is a capacity that we all share, and one that children naturally acquire very early on in their development. The question at the core of the linguistic enterprise is: how should this ability be explained? It cannot be reduced merely to the observables of the speech signal and some induced behaviour, for there is no relation between any one speech signal and some correlated behaviour. But nor can it be seen as just knowledge about sentences in some abstract use-removed sense, since there is no deterministic relation between any one speech signal, simple or complex, and some fixed interpretation. Interpretations are highly dependent on context in what seem to be multifarious ways. This is very

well-known with regard to anaphoric expressions, as witness the interpretation of the pronoun *it* in (1-4):

- (1) It was rock-hard.
- (2) The cake had been cooked for so long it was rock-hard.
- (3) Each cake had been cooked so long that it was rock-hard
- (4) Each cake was inedible. It had been cooked for so long that it was rock-hard.

(1) involves a so-called indexical pronoun, with the pronoun construed as referring directly to some individual in context. (2) involves a so-called coreferential use of the pronoun, referring to some individual by virtue of the previous use of an expression that is taken to pick that individual out. (3) is a bound-variable use of a pronoun, its interpretation controlled by the quantifying expression, *each cake*, that “binds” it. In these cases, as in all others, it is the combination of an available antecedent and containing structure that, in some sense, determine the way that the pronoun gets to be understood.

Here problems immediately arise, for an interpretation may be decidable within the boundary of the uttered sentence, as in (2), but it may not be, as (4) shows: this example appears to involve bound-variable uses of a pronoun across a sentence boundary, unless the pronouns in (4) are given a wholly different explanation from that of (3). Then with (5), there appear to be two formulations of what is conveyed. Either there can be said to be a bound-variable use of a pronoun across a sentence boundary, or the sequence can be analysed as having a pronoun that picks out some arbitrary witness of what makes the first sentence true, viz. whatever cake was depicted to be inedible:

- (5) One cake was inedible. It had been cooked so long that it was rock-hard.

In uses of this type, the pronoun establishes its construal by a process of composing the full meaning of the first sentence and using that composed result as the basis for interpreting the pronoun in the subsequent sentence. These are the so-called *E-type* pronouns (Evans 1980). Given the assumption that quantifiers are a type of expression whose interpretation is not available across sentence boundaries, these E-type pronoun uses have been seen as constituting a distinct way in which expressions can display dependence on the context in which they are understood. So we appear to have at least four different types of interpretation for pronouns.

There is yet a further way in which pronouns can vary: a use in which they apparently do not pick up on the interpretation assigned to the antecedent expression, nor on its containing environment, but on some weaker

concept of sameness of interpretation, the pronoun being interpreted in the same MANNER as its antecedent. In (6), for example, the pronoun *it* is interpreted as though it somehow contains a genitive pronoun, just like its antecedent (*his* in *his bread*), and furthermore this “covert pronoun”, is likewise interpreted relative to its local, newly presented subject:

- (6) Sandro gets his bread to rise using a bread-machine; but Gianni gets it to rise in his airing cupboard.

So the resulting interpretation of the second conjunct (in most contexts) is that ‘Gianni gets his own bread to rise (not Sandro’s)’. But if this is an appropriate characterisation of this form of anaphoric interpretation, then the nature of the antecedent provided has to be that of some representation, since attributes of this replicated representation have to be interpreted relative to the new clausal context.

This conclusion is controversial, and, as we shall see, it remains a live issue whether representations of content are essential to capturing the nature of anaphoric dependence at an appropriate level of generality. In the meantime, what is certainly perplexing about the overall phenomenon of anaphora is that, despite the robust folk intuition that anaphoric expressions are transparently dependent on context for their construal, the detailed specifications of the apparently required meanings constitute little more than a heterogeneous list of types of interpretation, at least if meanings are to be given in denotational terms, following the formal language methodology. And if we are to address the challenge of modelling the phenomenon of context-dependence and the twinned flexibility of natural-language expressions themselves that makes this possible, then this heterogeneity has somehow to be explained as something other than such a mere list.

Tense and aspect give further indications of the context-dependent construal of language: their uses in placing described events in time involve many complexities that depend on specifics of the eventualities being described (Kamp 1980, and many others since). The issue is not simply that tense is a deictic category whose interpretation depends on the time of utterance; it also displays anaphoric properties, so one tensed clause may depend on another to identify the time at which it holds. For example, in a sequence of such clauses, it may be the case that the time at which one eventuality occurs is the same as that of some other eventuality (contrary to any naive assumption that discourse is made up of eventualities simply narrated in sequence):

- (7) When John came into the room, he was smiling.

Furthermore, the lexical semantics of a verb may determine how eventualities relate to each other temporally. So states and events differ with respect to how eventuality times and reference times relate, and other contextually

determined subtleties of interpretation abound (see, for example, Kamp and Reyle 1993).

The appearance of complex anaphoric phenomena against the background of a more general and intuitive sense of context-dependence is redolent of pronoun construal (and the pronoun/tense parallel has long been noted; Partee 1973). Indeed, construal of tense has also been used in the debate over whether representations of content are essential to the analysis of anaphora (Kamp 1980, Hamm et al 2006, Baggio et al this volume).

In ellipsis, another context-dependent phenomenon, the need for structural reconstruction is, in some pretheoretical sense, not in question. Indeed the motivation for positing structural representations of meaning for at least some cases of ellipsis is so clear that ellipsis has been taken by many not to be a matter of context-dependence at all, but a grammar-internal matter, with the syntactic characterisation of the string said to have some of its elements simply unpronounced:

(8) John was over-wrought about his results though Bill wasn't.

This analysis of ellipsis as invisible syntactic structure, preserving some concept of sentencehood for the elliptical fragment, might superficially seem to be the appropriate form of explanation. However, as with pronouns, such context-dependent effects cannot be reduced to mere reconstruction of the words themselves. We know this because of the many challenges which ellipsis construal poses for the analyst (Fiengo and May 1994). Amongst these is the way in which ellipsis construal can be ambiguous, even given a single interpretation of its antecedent phrase. For example, restricting the first predicate of (8) to the interpretation that John was over-wrought about his own results, the elliptical form *wasn't* in the second conjunct can be understood in two ways. Either Bill was not over-wrought about John's results (matching the predicate content of the antecedent clause, the so-called "strict" interpretation), or Bill was not overwrought about his own results (the so-called "sloppy" interpretation).

Moreover, despite the superficial distinctiveness of anaphora and ellipsis, this ambiguity echoes that of anaphoric dependence. On the "strict" interpretation, the interpretation of the predicate is exactly that of its antecedent, just as in instances of pronominal coreference. On the sloppy interpretation, the pronoun within the reconstructed predicate has to get reinterpreted relative to the new subject paralleling the MODE of interpretation assigned to its antecedent, rather than picking up on the denotational content of that antecedent. This more indirect form of interpretation parallels the mode of interpretation needed for (6) (itself labelled with the term "lazy pronoun" usage: Karttunen 1976). In both cases, in some sense to be made precise, it is the way in which interpretation is built up in the antecedent that is replicated in the construal of the ellipsis-site/lazy-pronoun. So there are clear parallelisms between types of construal in anaphora and ellipsis.

Like anaphora and tense construal, the strict/sloppy ambiguity associated with VP ellipsis may span more than one sentence:

(9) John was over-wrought about his results. Bill wasn't.

(9) is ambiguous in exactly the same way as (8). This phenomenon is not restricted to particular types of expression, as there is nominal as well as predicate ellipsis; for example, in (10) *most* is construed as 'most of my students', again proving to be both intra-sentential and super-sentential:

(10) I wrote round to all my students, though most were away.

(11) I wrote round to all my students. Most were away, but a few answered.

However, the issues are broader than some single sub-type of ellipsis, indeed broader than the objective of explaining variability in anaphora and ellipsis. The challenge for linguistic explanation posed by context-dependence in general is that of formulating what it is about expressions of language that enables them to make some single systematic contribution to the way information is conveyed while nevertheless giving rise to multiple interpretations. In the case of anaphora and, even more strikingly, ellipsis, this would seem to necessitate some form of representation other than that inhabited just by the linguistic string of words. As the "lazy"/"sloppy" forms of pronoun/ellipsis construal show, the input to construal may involve adopting some particular interpretation of an antecedent string and building up structural properties in parallel with that to yield some novel interpretation. In addition, it is far from obvious that the remit of characterising sentence-internal phenomena provides a natural basis for characterising such context-dependence, since all of these various context-dependence phenomena can be associated with an antecedent either sentence-internally or sentence-externally. Indeed, if we stick to identification of sentence-internal dependencies as the exclusive remit of grammar-internal characterisations, leaving sentence-external dependencies on one side, all context-dependent phenomena will fail to receive complete characterisation.

Further problems arise when the remit of data to be considered is extended to conversational dialogue, where ellipsis is rampant. Both speakers and hearers have to be able to manipulate decisions relative to the context in which the communication is taking place; indeed, the planning and linearization of some intended thought may make essential reference to the way choices made are relative to the context. In conversational dialogue, these decisions made severally by speaker and hearer appear to come together, giving rise to what seems to be yet a further type of ellipsis. Utterances can be split across more than one person, with the speaker who is interrupting her interlocutor transparently building upon the point which that interlocutor has reached:

- (12) Father: We're going to Granny's.
Mother: to help her clean out her cupboards.
Child: Can I stay at home?
Mother: By yourself? You wouldn't like that.

In (12), the mother seamlessly extends the sentence initiated by the father by providing a non-sentential addition, continuing the narrative for the child in ways that are dependent on what is previously uttered (here understanding the subject of the nonfinite *help* as the mother and father, and the pronoun *her* as Granny). The child, having heard both, then interacts with his parents by asking his own question, which the mother responds to with another fragment, *by yourself*, as an extension of his utterance. This is done in the full expectation that the child will understand her use of the demonstrative *that* as referring to the proposal that he should stay at home by himself, even though this proposal exists in the context only as a result of the extension of the child's suggestion by the mother's utterance; it has not been expressed by any one participant in the conversation. Such an interpretation essentially depends on the coordinated accumulation of information by both parties, each building on the structure to date, whether as speaker or as hearer.

Seen as a phenomenon of ellipsis, this is no more than an illustration of the context-relativity of the construal of elliptical fragments; but in these cases such relativity strikingly includes dependence on partial structure provided by the context, which is essential to determine the wellformedness of the whole as built up from the severally uttered parts. In this connection, the last exchange between mother and child has special significance. It shows that the concept of structure presumed upon by both participants in the building up of interpretation is not that of the string each utters, but is rather some representation of the content that their strings give rise to. What the two participants - mother and child - utter is *Can I stay at home all by yourself?* It is clear that neither child nor mother takes what they have jointly been building to be a structure decorated by those words: as a sentence, this sequence is plainly ungrammatical, but what the interlocutors understand each other to have said is taken by both to be fully wellformed (as evidenced by the unproblematic subsequent use of anaphoric *that* discussed above). Looked at in more abstract linguistic terms, the issue here is what kind of structure crucial notions like 'locality' are defined over. Reflexive pronouns like *yourself* are conventionally defined as requiring identification of an antecedent determining their interpretation within a locally identifiable structure. This example provides evidence that the structure over which such locality has to be checked is not that of the words themselves. If it were, the result of the exchange between child and mother would not be wellformed: there is mismatch in person specification between the reflexive pronoun and the form of wording in which its antecedent has been presented.

Since this kind of exchange is wholly acceptable, it must instead be the case that the relevant notion of locality is defined over some representation of interpretation successfully constructed by the participants involved. Hence the necessity for structures representing meaning established by the use of linguistic expressions.

Note that this is general phenomenon: such speaker-hearer role switches can take place across any syntactic dependency whatsoever:

- (13) A: Has John [negative polarity dependency]
 B: read any books on the reading list?
- (14) A: I need the.. [determiner-noun dependency]
 B: mattock.
- (15) A: I think she needs [PRO control dependency]
 B: to see a doctor.

As these examples illustrate, it is characteristic of conversational dialogue that the construal of fragments may involve both recovery of content and recovery of structure, each participant adding a subpart to what becomes a shared whole. This phenomenon is so widespread that it is arguably diagnostic of human conversational dialogue. The question is: what does all this amount to in the search for what constitutes the human capacity for language, or knowledge of language?

Until very recently, such dialogue data have been totally ignored by linguists as no more than a demonstration of dysfluencies observable in language use. Being designated a performance matter, such data lie quite outside the normal canon of data familiar to linguistic argumentation (though see Purver et al 2006, Cann et al 2007, Gargett al 2009, Kempson et al 2009). Yet this phenomenon is wholly systematic, forming a large proportion of the sole type of data to which young children are exposed in learning language. If we take seriously the challenge that grammars of natural language must be definable as licensing systematicity of structure in the data the language-acquiring child is exposed to, then in these dialogue data we are brought up against evidence that models of linguistic knowledge must include representations of content as part of the vehicle of explanation. For cases such as (12)–(15) show that these, not representations inhabited by the string of words, are the representations over which syntactic restrictions have to be defined. So the conclusion here buttresses that reached through more conventionally recognised forms of ellipsis: structural representations of content and context are essential to the characterisation of such fragments as part of a general explanation of the context-dependence that natural languages invariably display.

Of course, there is always the option of ignoring these data, simply deeming them to be a dysfluency of performance, and in principle not within the

remit of a conventional competence-based form of explanation. In particular, the phenomenon of split utterances is at odds with a grammar formalism in which the concept of sentence is the central notion. Standardly, all dependencies to be characterised within the grammar are defined relative to sentence-sized units independent of any application of them, and without any concepts of underspecification and growth.

It is important to recognise that such a decision comes with a cost. The effect will be that no single linguistic phenomenon will receive a complete characterisation. On the one hand, accounts of context-dependent phenomena such as anaphora and ellipsis will be incomplete, bifurcated into those phenomena which can be characterised sentence-internally, and those which cannot. On the other hand, the structural dependencies themselves will suffer the same fate: not a single one of them will be fully characterised by the grammar, for such a characterisation will not include the use of fragments in discourse in which participants shift roles mid-sentence. By the same token, we cannot look to performance mechanisms to unify such phenomena: as long as we stick to the sentence-based view of competence, the grammar will explain what it can, and stop short of the rest. Note that, on conventional assumptions, fragmentary utterances like those in (12)–(15) do not even provide a grammatical trigger for performance mechanisms. Not being wellformed sentences, these fragments simply will not form part of the output of the grammar. There is therefore no possibility of even a degree of continuity between such examples and the more widely recognised kinds of ellipsis that they resemble in significant ways.

This failure of the grammar to provide any coherent basis for an account of the general phenomenon of context-dependence is a serious defect, since it constitutes a failure to address the very phenomenon which critically distinguishes natural languages from formal languages. It is thus this conception of the competence-performance distinction, and the divide it imposes on the data to be characterised by taking the remit of grammar to be sentence-internal phenomena, to which we shall ultimately have to return in seeking to include some explanation of the phenomenon of context-dependence as an essential objective for any model of natural language. But this is to leap ahead. To see how and why certain representationalist views of language emerged as a result of articulating formal grammar specifications of natural language, we have to go back to the point of departure for core notions in syntax and semantics – the grammars of formal languages.

3 Syntactic and semantic representations

3.1 Formal languages and the form of grammar

The classical formal languages of propositional and predicate logic were defined not for the study of language but for the formal study of mathematical

inference, though predicate logic incorporated a partial reflection of natural-language structure in its articulation of subpropositional predicate-argument structure. Logic is the formal modelling of inference, which involves truth-dependence between wellformed formulae of the defined language. The objective in defining such a formal language is to capture all and only the valid inferences expressible in that language, via some concept of truth with respect to a model. The task is to posit a minimal number of appropriate units and structure-inducing processes that together yield all and only the appropriate outputs of the grammar, viz. the infinite set of wellformed strings, over which the inferences under study can be defined. The objective is to derive the infinite complexity of valid inference patterns from the interaction of a small number of primitive notions.

The perspective of mathematical inference imposes an important restrictive capacity: it is global and static. There is no modelling of context external to the structure being defined – mathematical truths are by definition independent of context. There is no modelling of growth of information or of its corollaries, underspecification of content and the concept of update. In fact, the flow of information is in exactly the opposite direction: rather than building information, inference involves only what follows from some information that is already given, the premises. There are therefore fundamental reasons to doubt that the methodology of describing these formal languages could ever be sufficient for modelling natural languages. If the interpretation of expressions of natural language necessarily involves the building up of information relative to context, then a formal explication of this process is required. Models based in mathematical inference will not provide this, even though insights as to the nature of inference undoubtedly form a sub-part of a full characterisation of natural language interpretation.

Despite its restrictiveness, the methodology for articulating formal languages has transformed the landscape within which formal frameworks for natural language and natural language grammars have developed; and the assumption of a truth-conditional basis to semantics for natural language is very widely adopted. In predicate logic, the grammar defines a system for inducing an infinite set of propositional strings which are taken to be truth-value denoting; and sentence-sized units are defined as having predicate-argument structure made up of predicate forms and individual constants, with naming and quantificational devices. Syntactic rules involve mappings from (sub)-formulae to (sub)-formulae making essential reference to structural properties: these rules constitute a finite set of principles inducing an infinite set of strings. Semantics is the provision of an algorithm for computing denotations for arbitrarily complex strings: the result is a pairing of strings and the objects they represent. This pairing is determined on the basis of some notion of content assigned to elementary parts, plus rules that determine how such contents are to be composed, through stepwise correspondence with syntax, yielding the values true and false as output.

The pattern provided by such formal languages was famously extended to natural language semantics by Montague (1974), who argued that natural languages could be characterised as formal languages, with semantics defined in terms of reference/denotation/truth with respect to a model. To achieve this, Montague developed a program of logical types and formulations of content for expressions of the language which are defined and articulated in the lambda calculus. These were defined to enable the preservation of predicate logic insights into the meanings to be assigned to these expressions even while sustaining the view that composition of such meanings is determined from some conventional analysis of the syntax of the language. Consequently, the natural language grammar, like a formal language, is conceived of as a system that induces an infinite set of strings paired with denotations, where these denotations are determined by semantic rules which directly match the combinatorial operations that produce the strings themselves. For example, a functional abstract can be defined using a functional operator, the λ operator, which binds an open variable in some propositional formula to yield a function from individuals to propositional contents, as in $\lambda x.x \text{ smiled}$. If we take this to be the contribution of the intransitive verb *smiled*, and we take a constant of type $\langle e \rangle$, *john*, to be the contribution of the noun phrase *John*, then it is clear that this allows semantic composition to mirror the combination of the words yielding the string *John smiled*. At the same time, a further functional abstract can be defined in which a predicate variable is bound in some open propositional formula to yield a function from properties to propositional contents $\lambda P.P(\textit{john})$ (semantically, the set of properties true of John, or the set of classes that include John of type $\langle \langle e, t \rangle, t \rangle$). This is equally able to combine with the predicate $\lambda x.x \text{ smiled}$ to yield the proposition expressible by *John smiled*, only in this case the contribution of the subject is the functor and that of the verb is the argument.

As even these simple examples show, there are in principle a number of different modes of combination available, involving different functor/argument relations – with potentially many ways of deriving the same string-meaning pair as more complex sentences are considered. If this approach is applied with no independent constraints on syntactic analysis, the syntactic structure assigned to strings is effectively epiphenomenal, being no more than a vehicle for the semantic operations: this account is notably espoused in categorial grammar formalisms (Moortgat 1988; Morrill 1994, 2010, this volume; Steedman 2000). These grammars are non-representationalist in that, on the one hand, the semantics of strings is defined in denotational terms (in terms of individuals, sets of individuals, and functions on those which ultimately map on to concepts of truth and inference); and, on the other, the rules of syntax constitute nothing more than mappings from strings onto denotationally interpreted strings (that is to say, mappings from strings to strings suitably paired with mappings from

denotational contents to denotational contents). Any invocation of structure is then no more than a convenient way of talking about such pairwise mappings of strings and assignable denotational contents.

Even without adopting this strict variant of the Montague claim about natural languages as formal languages, the influence of the formal-language methodology holds sway very generally. On a more pluralistic model of natural-language grammar – influenced by Montague solely with respect to semantics – a natural language grammar is a finite set of rules which assigns structure to sentences, and it is these syntactic structures to which denotational interpretations are assigned, defined in terms of truth with respect to a model. Semantic operations are thus defined in tandem with syntactic ones, most transparently applied in the so-called ‘Rule-to-Rule Hypothesis’ whereby each structure-defining syntactic rule is paired with an appropriate semantic one. This is the dominant model in work within theoretical linguistics that labels its object of study ‘the syntax-semantics interface’ (e.g. Heim and Kratzer 1998) – which is to say most work that purports to be formally explicit about both syntactic structure and its interpretation. In terms of representationalism as a whole, this constitutes a mixed approach. The view of syntax is representationalist, in that there are assumed to be fixed structures defined over strings of words. But the semantics is not representational, at least if conceived of in terms of conventional logic formalisms, because the semantic characterisation assigned to each syntactic structure is given in terms of denotation with respect to a model (or some equivalent).

It remains a matter of controversy in linguistic theory whether syntactic and semantic operations can be directly matched in this way. While for some analysts the direct assignment of content to syntactic structures remains an ideal worth striving for, others work on the basis that this is demonstrably impossible for natural languages. Broadly, speaking, there are two common kinds of claim for the necessary divergence of syntactic and semantic structures, necessitating multiple levels of representation. One is that the interpretation of natural languages requires representations of meaning that are not directly interpretable in terms of denotations. This relates to issues of context-dependence, and we return to it in section 4. The other kind of claim is that natural language syntax has distinctive properties that are neither reducible to, nor reflected in, its semantics; this is our next topic of discussion.

3.2 The syntax-semantics split

As we saw in a preliminary way in section 1, the phenomena of anaphora and ellipsis and their interaction in conversational exchanges provide informal evidence that in order to interpret natural language expressions, structure may have to be built up over and above that presented by the expressions

themselves. The minimal stance, as we set it out there, was that (contrary to the Montagovian position outlined in section 3.1) representations of content must indeed be posited, in addition to representations of the words themselves and their phonological properties. We return to this claim in more detail in section 4. Assuming for now that it is correct, an obvious question arises with respect to the economy of the overall system (though it is rarely asked): do sequences of words (i.e. sentences) need to be assigned representations of structure over and above whatever representations of content might be attributable to them?

At the level of the observable data, the answer is very widely taken to be yes. This assumption is generally held to follow from the incontrovertible fact that there is no one-to-one correspondence between a string and an assignable interpretation. In the one direction, there are clearcut cases where expressions have more than one interpretation, and these are taken to warrant the invocation of discrete tokens; i.e. so-called structural ambiguities. More tellingly, in the other direction, there are strings which systematically differ in some structural properties but have identical interpretation, at least at the level of denotational content. These are the cases of so-called discontinuity made famous by Chomsky (1957), which feature a pair of sentences that express the same semantic content, but are asymmetrically related to each other in a particular way: one appears to have structure that can be mapped directly onto the associated semantic content, while the other seems to map onto this content only by making reference to the structure of the first sentence (at least so Chomsky 1957, 1965 and others argue). There are both local and non-local variants of this phenomenon. The first of these is displayed by so-called *expletive pronouns*:

(16) That Eliot will bring food for the party is likely.

(17) It is likely that Eliot will bring food for the party.

(18) A man is singing outside in the garden.

(19) There is a man singing outside in the garden.

These examples show relatively local discontinuity between the expression in question, here the expletive pronoun, and the linguistic material providing its interpretation (in (17), for example, the interpretation assigned to the pronoun is provided by the end-placed clause *that Eliot will bring food for the party*).

Non-local cases, like (20) and (22), have famously been said to involve movement from one structural position to another. Here, the discontinuity between a certain expression and the site at which it is assigned an interpretation may span an indefinitely long intervening sequence: *what* in (20) and *the new book by Sue* in (22) must each be interpreted as the internal argument of the predicate given by the verb *read*, and so, by hypothesis,

must be related to the normal, postverbal position of the syntactic object of *read* (which is shown in the position of these expressions in (21) and (23)):

(20) What did John say that we should all read by tomorrow?

(21) John said that we should all read what by tomorrow?

(22) The new book by Sue, John said we should all read by tomorrow.

(23) John said we should all read the new book by Sue by tomorrow.

Arguably, both (20) and (21) express a question about what John said should be read by tomorrow, yet in (20) this involves front-placement of the expression *what*. In like manner, (22) and (23) do not essentially differ in interpretation. There thus appear to be structural properties of strings that have no necessary counterpart in their semantic characterisation, this being presumed to be a characterisation of conditions for the truth of such strings.

Moreover, such ‘displacement’ effects have other properties that have been claimed to necessitate separate representations for syntax and semantics. Since the work of Ross (1967), linguists have investigated apparent constraints on the locations of certain expressions, relative to the positions in which they seem to be interpreted (often conceived of as constraints on syntactic movement). These constraints appear to be inexpressible in semantic terms, and so are taken to warrant an independent level of representation for syntax. For example, according to one of the most well known of these constraints, the Complex NP Constraint, expressions cannot be construed across a relative clause boundary. Though (20) is well-formed, (24) is not:

(24) *What did John introduce us to the American lecturer who said that we should all read by tomorrow?

To understand why this is taken to demand the separation of syntactic and semantic representations, it is important to note that the conventional tools of formal semantic analysis fail to predict the illformedness of (24). Standard methods definable using the lambda calculus allow the material following *what* to be treated as a form of predicate abstract, with the lambda operator binding a variable (the internal argument of *read*) at an arbitrary depth of embedding. This means that *what* can be related to the object of *read* in (24), just as in (20). More specifically, it should be possible to treat *what* as a functor of type $\langle\langle e, t \rangle, t\rangle$, which will apply the relevant lambda-abstract to yield an overall interpretation in which *what* queries the object of *read*. There is nothing on the semantic side that blocks this in the case of (24). In essence, the problem is that the lambda calculus is blind to syntactic details such as the presence or absence of a relative clause boundary. It would seem to follow that semantics and syntax require distinct forms of

generalisation, expressed in different vocabularies. Indeed the sensitivity of some phenomenon to the Complex NP Constraint is commonly taken to be diagnostic of its syntactic underpinnings (see e.g. Partee 1976, Merchant 2009, Stainton 2006).

We have thus arrived at a point where phonology, syntax and semantics are widely accepted to require discrete forms of generalisation.⁵ As such, the overall tendency towards representationalism is solidly established. Given this, it has seemed to many linguists a relatively small step to posit additional types of representation in a grammar, any one of which by definition requires a different form of vocabulary. This is manifested in a framework like Lexical Functional Grammar (LFG, Bresnan 2001), which posits separate levels of representation for a range of putatively distinct types grammatical information, each of which is expressed with wholly distinct vocabulary, and furthermore with transition rules in yet a further vocabulary mapping representations at one level onto those at another (Dalrymple, M. 1999). So we have c(onstituent)-structure representations with enriched ‘functional equations’ to derive the mapping from this type of representation onto f(unctional)-structure, which encodes the grammatical functions played by the various constituents and is expressed in an entirely distinct vocabulary. Further rules relate f-structures to semantic structures and other types of representation have been variously posited as necessary to correctly characterise the grammatical properties of natural languages. LFG is by no means unique in this tendency, however, nor the most extreme at positing distinct modes of grammatical representation (cf. for example Role and Reference Grammar, van Valin and LaPolla 1997), but it illustrates one end of a spectrum of representationalist approaches.

There is a distinction at this point which is important to bear in mind. The issue of how many levels a grammar defines is independent of the issue of types of representation advocated. In one direction, discrete levels may be posited relative to a single type of representation, as in the transformational grammar of the nineteen-seventies with its deep and surface structure levels, both of which were expressed as syntactic forms of representation, articulated in a common vocabulary. In the inverse direction, a grammar may be mono-level (or ‘mono-stratal’), in which a single format of presentation is adopted. In Head-driven Phrase Structure Grammar (HPSG, Pollard and Sag 1994), for example, all grammatical information is encoded in the Attribute Value Matrix (AVM) notation and linguistic signs are considered to be fully defined by such matrices, but this apparent single level of representation nonetheless masks the fact that wholly discrete vocabularies are used within various parts of an AVM, with different modes of unifying information as the signs for individual words and phrases are combined. These

⁵It is not standardly argued in phonology that structures in phonology should be reducible to syntax.

different vocabularies correspond to disjoint categories, with distinct forms of generalisation expressible within them, effectively giving rise to different types of representation for different categories of grammatical information. So the issue of how many types of representation to advocate within grammar formulation cannot be reduced to the decision as to how many levels of representation are considered to be required. Decisions over how many levels of analysis to posit are driven by the objective of capturing generalisations in the most revealing manner. But the issue of how many types of representation to invoke is a claim as to the richness of the ontological base of natural language as a cognitive phenomenon, hence is intrinsically a foundational issue.

The issue of the relative independence of syntax and semantics has constituted the core of the representationalism debate as argued within linguistic theorising, if simply because of the live issue of whether generalisations about natural language interpretability requires the articulation of representations of content at all. As we shall see, as conventionally expressed, this turns out to be the wrong question to ask. But let us first examine the issue as it is usually formulated.⁶

4 Representations in semantics?

The argumentation about distinguishing syntax and semantics that we considered in the previous section is usually taken to show that syntax requires a securely distinct type of representation, one that, following Chomsky (1965), acts as the core of the grammar. Its uniquely defining vocabulary, involving tree structures, category specifications, and whatever other features are taken to be needed to express appropriate syntactic generalisations, articulates a system of rules that is intended to license all and only the grammatical strings of some language independently of context. Such a stance is broadly in line with the formal language conception of syntax.

Leaving on one side more general problems associated with the concept of defining grammaticality in terms of the well-formedness of sentences (Pullum and Scholz 2001), there are significant problems with this view of language when we turn to a fuller consideration of semantics. Right from the outset

⁶Morphology, and the status of morphological sub-word generalisations, is another case in point. We do not address these here, but note the debate within minimalism as between those who take the set of morphological phenomena to require wholly disjoint vocabulary specific to morphology, e.g. the positing of *morphological templates*, and those who take the set of phenomena constitutive of morphology to be broadly syntactic, with explanations involving movement relative to appropriate feature-specified triggers. There is also the stance more pertinent to the issues raised in this chapter in which it is argued that the major morphological generalisations thought to provide evidence for such morphology-specific phenomena such as morphological templates can be expressed solely as a mapping from phonological sequences onto representations of content (see Chatzikyriakidis and Kempson 2011).

of the formal semantics programme, the need for some characterisation of context-dependence was recognised, but at least initially the problem was taken to be essentially peripheral to the core insight that truth-conditions can be presented as a basis for articulating natural-language semantics relative to orthodox assumptions of syntax (Lewis 1972). However, as research on the problems posed by context-dependence have deepened (from Kamp 1981 and Kaplan *** onwards), there has been increasing emphasis on the on-line dynamics of language processing and update; and such moves are almost diametrically opposed to the perspective in which syntax and semantics are defined in tandem, with reference only to sentences. Yet, as we saw informally in section one, there is evidence that in conversation interlocutors rely on some notion of representation that allows them to interpret an elliptical utterance, extend or complete an utterance, clarify some aspect of what has been said, and so on. Moreover this notion of representation must involve attribution of structure in order, for example, for reflexives to be used coherently in split utterances.

In coming to grips with the challenges posed by such context-dependence of content, linguists have begun to ask the further question: what is it in the nature of linguistic expressions that enables their apparent intrinsic specification of content to interact with aspects of context to determine their particular construal in a given linguistic situation (see in particular Cooper 2005, this volume, Larsson 2008)? In other words, what sort of underspecified content can be attributed to lexical items and what sorts of updates can be provided by linguistic and non-linguistic context to derive the meanings of words and expressions that interlocutors understand in particular conversational exchanges? This provides another source of tension with the formal-language specification of the relation between syntax and semantics, as the assumptions of that methodology are no more set up to express a concept of underspecified meaning or its update than they are to express dependence on a structured concept of context. Moreover, the static and context-independent conception of linguistic competence is destabilised by any attempt to accommodate underspecification and update, since these notions require that the model of competence takes account of processes which apply when language is used in context.

The first moves in the direction of this form of analysis were taken by semanticists modelling the nature of information update as the basis for language understanding. These moves have led to ever-increasing richness of structure in the formal specifications of both content and context, in particular in order to explain phenomena such as anaphora and ellipsis. As we shall see, it is this trend which enables the question of relative dependence between syntax and semantics to be re-opened, with the potential for new answers.

As the initiator of this movement, Kamp (1981) set out two points of view of what meaning is, each of which has been assumed in different parts

of the linguistic literature:

- (25) (a) meaning is what determines truth conditions; the view taken by truth-theoretic and model-theoretic semantics
- (b) meaning is what language users GRASP when they understand the words they hear. This representationalist view of meaning is in, principle, what the majority of psychologists, computer scientists, linguists, and others working in cognitive science, aim to characterise.

The first view of meaning is what Montague Grammar promulgated; and many philosophical accounts concerned with reference and truth can be taken as espousing this view. It articulates a concept of a language as a set of interpreted symbols wholly independent of any agent's knowledge of it. On the other hand, the representationalist point of view which Kamp advocates within semantics involves defining formal constructs which are assumed to model the mental representations humans employ in response to linguistic processing. This is the view taken by various cognitive science and linguistic approaches, with perhaps the most prevalent being the so-called *computational theory of mind*. On this view, the mind is a system for manipulating symbols according to syntactic rules which determine the recursive complexity of thoughts a so-called 'language of thought' (Fodor 1975, 1983). On this view, human cognition operates systems of internal cognitive representations (possibly more than one) enabling individuation of objects for the purposes of mind-internal reasoning (Fodor 1983 and subsequently, Hamm et al 2006, and Baggio et al this volume).

Discourse Representation Theory (DRT) was the first theory which aimed to combine the two approaches to meaning listed in (25); DRT is motivated by the need to give an account of how discourse processing leads to the generation of a representation of the semantic content of a discourse.

4.1 Discourse Representation Theory

The immediate objective of Kamp (1981) was to articulate a response to the challenge of modelling anaphoric dependence in a way that enables its various uses to be integrated, contrary to the simple postulation of ambiguity to account for different modes of interpretation (as in the indexical, bound-variable, or E-type interpretations of pronouns exemplified section 2). Sentences of natural language were said to be interpreted by a construction algorithm for interpretation which takes the syntactic structure of a string as input and maps this onto a structured representation called a *Discourse Representation Structure* (DRS). This constitutes a partial model for the interpretation of the natural language string which contains named entities (*discourse referents*) introduced from natural language expressions, and

predicates taking these as arguments (*conditions* on referents). The sentence relative to which such a partial model is defined is said to be true as long as there is at least one embedding of the model so constructed into an overall model. So, for example, for a simple sentence-sequence such as (26), involving an E-type use of a pronoun,

(26) John loves a woman. She is French.

the construction algorithm induces a DRS for the interpretation of the first sentence containing discourse referents corresponding to the name and the quantifying expression, together with a set of predicates corresponding to the verb and noun:

x,y
John =x loves (x,y) woman (y)

Such a DRS can then be extended by applying the construction algorithm to the second sentence, extending the initial DRS to an expanded DRS:

x,y,z
John =x loves (x,y) woman (y) z = y French (z)

On this account, indefinite NPs are defined as introducing a new discourse referent into a DRS, while definite NPs and pronouns require that the referent entered into a DRS be identical to some discourse referent already introduced. Names, on the other hand, require a direct embedding into the model providing the interpretation. As noted above, once constructed, a DRS is evaluated by its embeddability into a larger model, being true if and only if there is at least one embedding of that partial model within the overall model. Notice that the basis for semantics is relatively conservative, in that the grounding model-theoretic assumptions are preserved unaltered.

Even without investigating further complexities that license the embeddability of one DRS within another (and the famous characterisation of *If a man owns a donkey, he beats it*), an immediate pay-off of this approach is apparent. E-type pronouns fall under exactly the same characterisation as more obvious cases of co-reference: all that is revised is the domain across

which some associated quantifying expression can be seen to bind. It is notable in this account that there is no structural reflex of the syntactic properties of the individual quantifying determiner; indeed, this formalism was among the first to come to grips with the name-like properties of quantified formulae (see also Fine 1986). It might seem that this approach obliterates the difference between names, quantifying expressions, and anaphoric expressions, since all lead to the construction of discourse referents in a DRS. But these expressions are distinguished by differences in the construction process. The burden of explanation for natural language expressions is thus split: some aspects of their content are characterised in essentially dynamic terms, in the mode of CONSTRUCTION of a representation (the intervening DRS), while some aspects are characterised in more traditional semantic terms, through the embeddability conditions of that structure into the overall model.

The particular significance of DRT lies in the Janus-faced properties of the DRS's thus defined. On the one hand, this intervening level is a partial model – or, more weakly, a set of constraints on a model – defined as true if and only if it is embeddable in the overall model, and hence essentially the same type of construct. On the other hand, it is argued (Kamp & Reyle 1993) that the specific structural properties of the DRS are needed in defining the appropriate antecedent-pronoun relation, hence such a level constitutes an essential intermediary between the natural language string and the denotations to be assigned to the expressions it contains. Nonetheless, this intermediate level has a fully defined semantics, via its embeddability into an overall model.

Despite the explicit model-theoretic characterisation provided by Kamp and Reyle, the DRT account of anaphoric resolution sparked an immediate response from proponents of the model-theoretic tradition. Groenendijk and Stokhof (1991) argued that the intervening construct of a DRS was not only unnecessary but illicit, in making compositionality of natural language expressions definable not directly over the natural language string but only via this intermediate structure (see also Muskens 1996). Part of their riposte to Kamp involved positing a new Dynamic Predicate Logic (DPL) with two variables for each quantifier and a new attendant semantics: one of these variables is closed off in ways familiar from predicate-logic binding, but the second remains open, bindable by a quantifying mechanism introduced as part of the semantic combinatorics associated with some preceding string. This was argued to obtain cross-sentential anaphoric binding without any ancillary level of representation as invoked in DRT (see chapter on lexical meaning, this volume). Such a view, despite its novel logic and attendant semantics, sustains a stringently model-theoretic view of context-dependent interpretation for natural language sentences commensurate with e.g. Stalnaker: 1974, 1999): the progressive accumulation of interpretation across sequences of sentences in a discourse is seen exclusively in terms of inter-

sections of sets of possible worlds progressively established, or rather, to reflect the additional complexity of formulae containing unbound variables, intersection of sets of pairs of worlds and assignments of values to variables (see Heim 1982, where this type of program is set out in detail). However, the DPL account fails to provide an explanatory account of anaphora, in that they merely articulate a semantics for the outcome of the interpretation process, and they do not address Kamp’s objective of modelling what it is that pronouns contribute to interpretation that makes such diversity of resulting interpretations possible.⁷ It was this assigned objective which led Kamp to postulate a “DRS construction algorithm” with subsequent evaluation rules determining the embedding of the resulting structure within the overall model. It was only this way, he argued, that one could capture both the denotational diversity of anaphoric expressions in context and the intrinsic uniformity of such expressions within the language system.

Over and above the departure from orthodox semantic theories in shifting to representationalist assumptions, there is a further sense in which DRT constitutes a radical departure from previous theories. In defining a formal articulation of the incremental process of building an interpretation of discourse, relative to some previously established context, there is an implicit rejection of the severe methodology whereby no reflex of performance should be included in any specification of aspects of natural language competence. Indeed the construction algorithm for building DRS’s yields a formal reflection of the sentence-by-sentence accumulation of content in a discourse (hence the name *Discourse Representation Theory*). So DRT not only offers a representationalist account of natural language meaning, but one that reflects the incrementality of utterance processing, albeit one that takes clauses as the basic units of information and organisation in a text.

There remains a lurking problem in this move, however, of direct relevance to the representationalism issue. In assigning meaning to a sentence string, the DRS construction algorithm was defined to take as its input a syntactic structure for that string as articulated in some independently defined syntactic framework (for example some form of Phrase-Structure Grammar as in Kamp and Reyle 1993), and progressively replace that structure with a DRS. Hence the explicit advocacy of more than one type of representation, semantic as well as syntactic. Yet, looked at as a basis for a uniform account of structure in language, DRT lacks only some characterisation of those phenomena which linguists take to be within the purview of syntax; and many of these involve just the kind of interaction with anaphora which one might expect should fall within the remit of DRT-based explanations, given the assigned aim to provide a unitary account of anaphora. There is, for example, no characterisation of long-distance dependency, expletives

⁷The DPL account explicitly relies on some presumed independent indexing which the DPL characterisation is defined to reflect (Groenendijk and Stokhof 1991).

or quantifier scoping; and so, equally, there is no account of their systematic interaction with anaphora. If, however, such phenomena could be reconstructed in terms that build on the intermediate level of representation which DRT does indeed provide, the situation would be very different; for in that hypothetical scenario it would not be the level of content representation that would be redundant, but the level of structure as inhabited by the string. It is this scenario we shall shortly seek to actualise.

In the meantime, the debate between advocates of DPL and DRT is far from over; there has been continuing debate as to whether any intervening level of representation is justified over and above whatever syntactic levels are posited to explain structural properties of natural language expressions, preserving orthodox assumptions of syntax. Examples such as (27)-(28) have been central to the debate (Kamp 1996, Dekker 2000):

(27) ?Nine of the ten marbles are in the bag. It is under the sofa.

(28) One of the ten marbles isn't in the bag. It is under the sofa.

According to the DRT account, the reason why the pronoun *it* cannot successfully be used to refer to the one marble not in the bag in (27) is because such an entity is only inferable from information given by expressions in the previous sentence. No REPRESENTATION of any term denoting such an entity in (27) has been made available by the construction process projecting a discourse representation structure, on the basis of which the truth conditions of the previous sentence are compiled. So although in all models validating the truth of (27) there must be a marble not in the bag described, there cannot be a successful act of reference to such an individual, using the pronoun. By way of contrast, the anaphoric resolution is successful in (28), despite its being true in all the same models in which (27) is true, because the term denoting the marble not in the bag is specifically introduced and hence represented in the DRS. Hence, it is argued, not only is the presence of an intermediate level of representation essential, but the concept of context from which expressions can pick up their interpretation must be representational in kind also.

4.2 Ellipsis

Even outside DRT debates, semanticists continue to argue over whether there is definitive justification of the representational nature of content and context (see von Stechow and Egli 2000, Ginzburg and Cooper 2004). In particular, ellipsis provides a case of particularly rich and structured concept of context-dependence, as we have already hinted in section 2; and there the dynamically evolving nature of both content and context is amply demonstrated. Indeed, ellipsis arguably provides a window on context, relying as it does on information that is in some sense manifestly available in the immediate context within which the fragment is produced/interpreted.

Consider again (8), repeated here:

- (8) John was over-wrought about his results though Bill wasn't.

As we saw, ellipsis is like anaphora in allowing multiple interpretations, determined by what information is accessed from context and how. This means that there no reason to anticipate an algorithmic relation between ellipsis site and assignable interpretation, nor between the fragment expression at that site and the assigned interpretation. So interpretation is not algorithmically determinable at either the level of the fragment itself or at the level of the propositional construal derived from it. Without some formal characterisation of how the context and content of sequences of utterances can evolve in tandem, one possible response is that the best that can be done is to analyze the relation between ellipsis site and the string from which its interpretation is built up, if there is one, as a grammar-internal phenomenon, leaving the nondeterminism of contextual factors aside. In principle, indeed, this has been the prevailing methodology. Until research into formal dialogue was developed (as, recently, by Ginzburg forthcoming, Ginzburg and Cooper 2004, Purver 2004, Cann et al 2005), the only forms of ellipsis that were addressed were indeed those where the ellipsis site can in some sense be analyzed sententially – either as a second conjunct of a compound conjunctive form or as an answer to a question. Relative to this restriction, the grammar-internal dispute has been whether a model-theoretic alternative is available for all such instances of ellipsis, with that semantic characterisation defined on the surface form of the string. Such an account would be non-representationalist, in the sense of invoking no additional semantics-internal concept of representation.

For some cases in particular, a model-theoretic account of this kind has seemed competitive. The debate is thus analogous to that between DRT and Dynamic Predicate Logic, but in this case played out over predicate meanings. For example, following Dalrymple et al (1991), ellipsis construal has been defined as an operation of (higher-order) abstraction on the model-theoretical content defined for an immediately antecedent conjunct. The resulting lambda-abstracted predicate can then be applied to the newly provided subject of the elliptical second conjunct.

By this means, as is minimally needed, a semantic explanation is available for cases such as (8) where, for a single antecedent form and assigned interpretation, ambiguity nonetheless arises. The process of abstraction over the content provided by the antecedent ('John is overwrought about John's results') yields two types of abstract: abstraction over solely the subject of the antecedent, as in (29), or abstracting over the subject and all other references to the individual it denotes, as in (30):

- (29) $\lambda x.x$ overwrought about Bill's results

- (30) $\lambda x.x$ overwrought about x 's results

The two resulting predicates apply to the lambda term assigned to the elliptical fragment to yield the so-called strict/sloppy ambiguity. Hence, it is claimed, the semantic basis for explanation does not require any intermediate DRT or other representational construct, contrary to the conclusion informally evidenced in section 1.⁸ Indeed, on this view, there is no parallelism with anaphora: the process is simply an operation on the denotational content of one conjunct to provide a fully specified construal for the second, elliptical conjunct.

Yet there is competing evidence that sensitivity to structure is essential to the way in which elliptical fragments are reconstructed. In particular, the very restrictions taken to be diagnostic of syntactic phenomena constrain some cases of ellipsis. So-called *antecedent-contained ellipsis* displays sensitivity to the presence of a relative clause boundary, as expressed in the complex NP constraint of Ross (1965), for no such boundary can intervene between the relative pronoun of the containing structure and the ellipsis site (Fiengo and May 1994, Merchant forthcoming):

- (31) John interviewed every student who Bill already had.
- (32) *John interviewed every student who Bill ignored the teacher who already had.

This is exactly as though the structure in question were present, not elided:

- (33) John interviewed every student who Bill already had interviewed.
- (34) *John interviewed every student who Bill ignored the teacher who already had interviewed.

On the syntactic account that such data are said to indicate, such structure is indeed present at some level of abstraction (only to be deleted through some low-level “PF Deletion” operation). Hence, so the argument goes, at least some types of ellipsis require syntactic explanation, involving full projection of clausal structure at the ellipsis site, with subsequent deletion of phonological material – hence an ineliminably representationalist account for at least some types of ellipsis. This brings its own problems, the first of which is that the VP apparently provided as the antecedent of the ellipsis site contains that site, threatening circularity. This immediately requires increased complexity in the syntactic analysis. A more general problem, however, is that once ellipsis is recast in structural terms, with full determinism of structure underpinning some required interpretation, the postulation of multiple ambiguities becomes inevitable. For example, the existence of sloppy and strict interpretations of ellipsis, as in (8), not only imposes distinct analyses of the ellipsis site, but it forces the use of different characterisations of coreference within the antecedent conjunct from which the construal is built up: one

⁸The identification of the subject of predication is stipulated within the semantic vocabulary, and the restriction to only two such bases for abstraction is controversial.

in terms of a binding mechanism, the other as free indexing, yielding the postulation of ambiguity in the source string even on a single interpretation. The matching of structure at ellipsis site and antecedent then forces the need to posit many different types of ellipsis (Merchant forthcoming).

We thus appear to have arrived at the division of ellipsis into semantic types and an array of syntactic types, all supposedly displaying sentential structure/content even though in fact realised in a fragmentary form, with the parallelism of anaphora and ellipsis construal totally lost. Over and above this, there are yet further cases where, arguably, there is no determinate linguistic basis for assigning interpretation to the fragment (Stainton 2006):

- (35) A: [coming out of lift] McWhirter's?
B: Second left.

Stainton argues that, contrary to both syntactic and semantic analyses, such cases do not allow any analysis as sentential reconstructions but have to be seen as a speech act that is performed without recourse to a sentential structure.⁹ These cases are redolent of the very broad use of fragments in dialogue (section 1), where a rich array of aspects of context may determine the way the fragment is understood:

- (36) Father: We went to parents' evening. The teacher
Child: Mrs Smith?
Father: No, your other teacher. Mr Jones. He said you were doing fine.

These fragmentary cases pose problems not merely for grammar-internal syntactic accounts of ellipsis, which effectively presuppose a single non-split utterance phenomenon, but also for the denotational account of ellipsis. On that analysis, the interpretation of the fragment would be expected to follow the general rule that its denotational content is provided by some lambda-abstract constructed from the content of the preceding utterance. This would require that the fragment *Mrs Smith?* in (36) must be interpreted as asking whether Mrs Smith went to parents' evening; yet this is not even a possible interpretation. The clarification requested has to be about the identity of *the teacher* just introduced into the conversation. This illustrates the special – and, for conventional approaches to ellipsis, highly problematic – properties of cases where the fragment occurs very early in the interpretation process that is under way: there is an unexpected and unexplained asymmetry between early and late use of fragments relative to

⁹Nevertheless, Stainton's argument also rests on the assumption that ellipsis mainly requires a sentential basis: his argument for the distinctiveness of what he calls *pragmatic ellipsis* is an argument to the effect that only these fragments have no such sententially based construal.

the context from which they build up their interpretations. Worse than this, whatever the fragment provides may have to ELIMINATE what has just been processed in context, violating monotonicity, rather than constituting a re-usage or development of that context. This is so even in cases where the interpretation is accumulative, and not a correction:

(37) Every doctor – John, Bill AND Mary – signed the certificate.

Again, such cases are problematic for denotational accounts of ellipsis. Denotationally, the extension of *every doctor* has to be eliminated before any interpretation can proceed, since the denotation of that phrase cannot be taken, simpliciter, to be equivalent to the set made up of John, Bill and Mary. Consequently, this kind of approach cannot capture the sense in which the addition of *John, Bill and Mary* is an elaboration on, or extension of, the original, quantified expression.

At this point, it might seem that there is an inevitable drift towards recognizing that ellipsis simply is a complex set of phenomena. Indeed, ellipsis has been characterised as displaying “fractal heterogeneity” (Ginzburg and Cooper 2004), calling on arbitrary types of information across the grammar-internal spectrum for reconstructing the requisite construal. The challenge of providing a uniform account of ellipsis understanding is thus signally not met: all that is achieved is an itemisation of disjunctive sets of ambiguities, some of them involving representations ineliminably, others said to be reconstructable from denotational contents directly. This amounts to a counsel of despair, admitting that no explanation can be found of what it is about elliptical fragments that enables them to yield such diverse interpretations, and to do so relative to context. Nor is any basis provided for why there is such widespread parallelism between anaphoric and elliptical forms of construal. If we are to find a solution to the puzzle of ellipsis in a way that reflects this parallelism, ellipsis has to be taken as illustrative of the much broader challenge of how to articulate the basis for context-dependence of natural language construal within formal linguistic theory.

5 A dynamic solution: from representation to construction

To address this challenge, we take up the underlying methodology of DRT and pursue it further. To achieve a fully integrated account of anaphora and ellipsis resolution, we seek to incorporate the dynamics of how interpretation is built up relative to context, while articulating a sufficiently structurally rich concept of such representations of content that structural constraints are expressible as necessary. We contend that the overarching problem lies in the assumption that representations are statically and independently defined. Another approach however is conceivable, one that

is equally representationalist, but adds another dimension - that of update from one representation to another, hence a concept of CONSTRUCTION of representations. The development of DRT, we suggest, pointed in the right direction with its modelling of anaphora resolution as reflected in the dynamics of how users build up interpretations in context; and to capture ellipsis, we generalise this to all structural dependencies, whether sentence-internal or supra-sentential.

The problem for DRT was the limit it imposed on the remit of its account. This reached a ceiling with forms of anaphora that are subject to syntactic explication, as these were, as sub-sentential structural dependencies, to be subject to grammar-internal characterisation, hence falling outside the DRT account. With ellipsis, the very same problem is replicated: the data are split into grammar-external and grammar-internal characterisations and then the latter into syntactic and semantic characterisations. In order to cross this hurdle, we need an account that is not restricted to the remit of the sentence-boundary, not restricted to arbitrary sub-types of structural dependency, and is essentially representational in its perspective on meaning specifications.

One way of extending the DRT form of analysis promises to remove these obstacles, thereby making possible an integrated account of context-dependent phenomena. This is to define a concept of time-linear incrementality within the grammar formalism itself, reflecting that of processing, so that both subsentential and supra-sentential dependencies are expressed as a process of update. The core syntactic notion becomes incremental growth of semantic representation and the concept of structure intrinsic to the string is replaced by the dynamics of building a progressively enriched semantic representation from a linear sequence of words, relative to a context that preserves a record of this growth process. With this shift, as we shall see, there is no intrinsic syntax-semantics divergence, no arbitrarily distinguished types of representation separating the two, and a structurally rich concept of context from we can express an integrated characterisation of ellipsis. One framework, in particular, provides the requisite formal tools to effect this perspectival and conceptual shift: Dynamic Syntax (Cann et al 2005, Kempson et al 2001, Gargett et al 2009, Kempson and Kiaer 2010).

5.1 Dynamic Syntax

Dynamic Syntax (DS), like conventional syntactic frameworks, provides a theory of language form and its interpretation. Unlike in conventional frameworks, however, this theory of form emerges from a model of interpretation, and more particularly from the growth of interpretation. DS is avowedly representationalist, in the sense that it depends upon explicit representations of both semantic structure (prior to model-theoretic interpretation) and the ways in which this structure is progressively induced from the incremental

processing of strings of words.

In this model, interpretations are represented as binary tree-structures of functor-argument form, and these are built up relative to context. Individual steps in this building process reflect the incrementality with which hearers (and speakers) progressively build up interpretations for strings, using information from context as it becomes available. It is like DRT in spirit, in that local predicate-argument structures are induced in a way that reflects the time line of processing, and such structured representations are taken to be essential to the expression of anaphora and ellipsis resolution. But it goes considerably further than DRT in a number of ways.

First, the mechanisms for building up such structures are presumed to apply in a strictly incremental way. Following the dynamics of on-line processing, representations of meaning are built up on a word by word basis – as opposed to the sentence by sentence approach of DRT. Second, this process of building up structure is taken to be what constitutes the syntactic component of the grammar: with the dynamics of structural growth built into the core grammar formalism, natural-language syntax is a set of principles for articulating the growth of structured semantic representations. Syntactic mechanisms are thus procedures that define how parts of interpretation-trees can be incrementally introduced and updated; they are therefore causally related to, but not constitutive of, the representations themselves. Third, reflecting the systemic context-dependence of natural language construal, all procedures for structural growth are defined relative to context; and context is defined to be just as structural and just as dynamic as the concept of content with which it is twinned. Context, by definition, constitutes a record not merely of the (partial) structures built up, with the typed formulae that decorate them, but also the procedures used in constructing them (Cann et al 2007). In short, the general methodology is a representationalist stance *vis-à-vis* natural-language construal (Fodor 1983), with the further assumption that concepts of underspecification and update should be extended from semantics into syntax. But the bonus of such explicit adoption of representationalist assumptions with respect to content is the avoidance of any further levels or types of representation – a clear application of Occam’s razor.

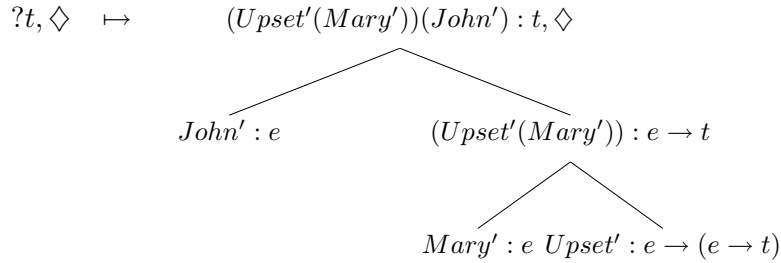
5.1.1 The tree logic and tree-growth processes

The general process of parsing is taken to involve building as output a tree whose nodes reflect the content of some uttered formula – in the simple case of a sentence uttered in isolation, this is a complete propositional formula. The input to this task, in such a simple case, is a tree that does nothing more than state at the root node the goal of the interpretation process to be achieved, namely, to establish some propositional formula.

For example, in the parse of the string *John upset Mary*, the output tree

to the right of the \mapsto in (38) constitutes some final end result: it is a tree in which the propositional formula itself annotates the topnode, and its various subterms appear on the dominated nodes in that tree, rather like a proof tree in which all the nodes are labelled with a formula and a type. The input to that process is an initial one-node tree (as in the tree representation to the left of the \mapsto in (38)) which simply states the goal as the requirement (shown by $?Ty(t)$) to establish some content, which is to be a formula of appropriate propositional type t (there is also a *pointer*, \diamond , which indicates the part of the tree that is under development at any given time):¹⁰

(38) John upset Mary.



Parsing *John upset Mary*

The parsing task, using both lexical input and information from context, is to progressively enrich the input tree to yield an output of appropriate type using general tree-growth actions and the sequence of words of the string. In order to talk explicitly about how such structures grow, trees need to be defined as formal objects; and DS adopts a (modal) logic of finite trees (LOFT: Blackburn and Meyer-Viol 1994).¹¹ DS trees are binary with the argument always appearing by convention on the left and the functor on the right. These are defined over mother and daughter relations, indicating a possible sequence of mother relations, or a possible sequence of daughter relations. The LOFT language makes available a basis for structural underspecification, using Kleene star (*) operators, making concepts

¹⁰This exegesis omits all indication of tense construal and quantification. In brief, the language of the formula representations is that of the epsilon calculus, with all quantified terms of type $\langle e \rangle$, matching the arbitrary names of predicate logic natural-deduction (of which the epsilon calculus is the formal study). The essential property of arbitrary names is that their structural representation is simple, but their semantics complex, by definition reflecting their containing environment. The advantage of this, in the linguistic application, is that this makes available a natural basis for name growth, following the general pattern of the DS framework, so that initial underspecification of a name under construction and its subsequent extension is naturally expressible. Tense construal then projects an epsilon event-term, with tense, adjuncts and aspect all taken to add to the restrictor specification of such terms, again a concept of term extension (see Cann fcmg).

¹¹There are two basic modalities, ways of describing node relations: $\langle \downarrow \rangle$ and $\langle \uparrow \rangle$. $\langle \downarrow \rangle \alpha$ holds at a node if α holds at its daughter, and the inverse, $\langle \uparrow \rangle \alpha$, holds at a node if α holds at its mother.

of *dominate* and *be dominated by* expressible for a tree relation even before the fixed number of such mother or daughter relations is fixed. For example, $\langle \uparrow_* \rangle Tn(a)$ is a decoration on a node indicating that somewhere dominating it is the node $Tn(a)$.¹² All that is determined is that the node in question must always be dominated by the $Tn(a)$ in any future developments of the tree.

A corollary of structural underspecification, and another essential feature of the model, is the existence of requirements for update. This is central to the ability to reflect the time-linearity involved in building up trees in stages (i.e. through partial trees). For every node, in every tree, all aspects of underspecification are twinned with a concept of *requirement*, represented as $?X$ for any annotation X on a node. These are constraints on how the subsequent parsing steps must progress. Such requirements apply to all types of decoration, so that there may be type requirements, e.g. $?Ty(t)$ (or $?t$), $?Ty(e)$ (or $?e$), $?Ty(e \rightarrow t)$ (or $?e \rightarrow t$); treenode requirements, e.g. $? \exists \mathbf{x} Tn(\mathbf{x})$ (associated with underspecified tree-relations); and formula requirements, e.g. $? \exists \mathbf{x} Fo(\mathbf{x})$ (associated with pronouns and other anaphoric expressions). These requirements drive the subsequent tree-construction process, because unless they are eventually satisfied the parse will be unsuccessful.

Such structural underspecification and update can then be used to define core syntactic notions in a way that follows insights from parsing, and the time-linear dimension of processing in real time. For example, they notably lend themselves to analysis of the long-distance dependency effects which, since the late 1960's, have been taken by most to be diagnostic of a syntactic component independent of semantics. When first processing the word *Mary* in (39) below, it is construed as providing a term whose role isn't yet identified. The parse is then taken to involve the application of a computational action which introduces a structural relation to the topnode (the initial root node decorated with $?t$) which is underspecified at this juncture: it is identifiable solely as being dominated by the topnode, and requiring type $\langle e \rangle$, i.e. bearing the requirement $?e$:

(39) Mary, John upset.

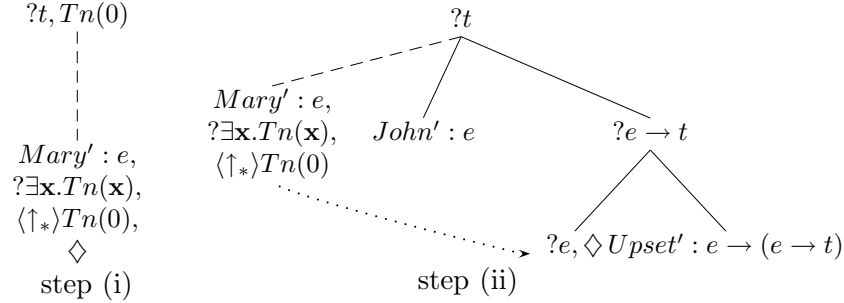
The expression *Mary* is thus taken to decorate an as-yet unfixed node: this is step (i) of (40). Accompanying the under-specified tree relation is a requirement for a fixed treenode position: $? \exists \mathbf{x}. Tn(\mathbf{x})$. All of this technical apparatus provides a formal reflection of the intuitive sense in which a string-initial expression of this kind finds its proper role in relation to form and meaning only after other aspects of the overall structure have been put in place.

The update to the relatively weak tree-relation in (40, (i)) becomes possible only after processing the subject and verb, which jointly yield the

¹²This is a standard tree-theoretic characterisation of *dominate*, used in LFG to express *functional uncertainty*; see Kaplan and Maxwell (1988)

two-place predicate structure as in step (ii) of (40). The simultaneous provision of a formula decoration for this node and update of the unfixed node is provided in the *unification* step indicated there, an action which satisfies the update requirements of both nodes to be unified, leaving no trace in the output of the word *Mary* having been string-initial.

(40) Parsing *Mary, John upset*:



This process feeds into the ongoing development in which, once all terminal nodes are decorated, bottom-up modalised application of labelled type deduction leads to the creation of the completed tree indicated in (38). Note that this is the same intuition that lies behind the syntactician’s notion of ‘displacement’, but captured without resort to abstract notions of movement or other syntax-specific mechanisms.

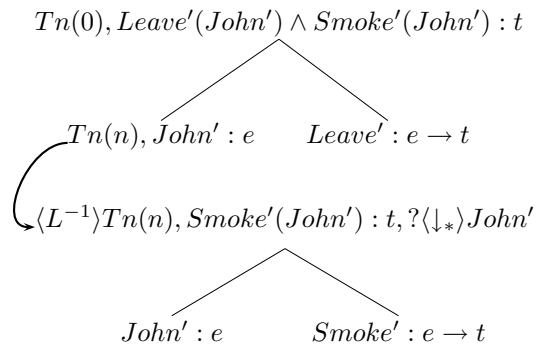
Such an account of structural underspecification and update is indeed not contentious as a parsing strategy; what is innovatory is its application within the grammar-mechanism to provide the central core of the syntactic generalisations and the characterisation of wellformedness. Discontinuity effects can now be seen on a par with anaphora construal, the latter an underspecification of interpretation to be resolved as part of the interpretation process, the former an underspecification of structure, equally to be resolved as part of that process. In the case of discontinuity, this construction of partial representation is a generally available option for tree growth. In the case of anaphora, it is the lexical projection of a place-holding formula value along with a requirement for its update which induces an element of underspecification, with both types of underspecification requiring update within the construction process.

This account might seem in principle to be skewed by focussing on parsing, but this is only superficial. Production follows the very same processes, with but one further assumption: that at every step in production, there must be some richer tree, a so-called *goal tree*, which the tree under construction must subsume, in the sense of being able to be developed into that goal tree, according to the system defined. So parsers and producers alike use strategies for building up representations of content, either to establish an interpretation for a sequence of words, or to find words which match the content to be conveyed. Both of these activities are alike also

in being context-dependent, so that structural and content choices may be determined on the basis of what has just been processed.

To achieve the basis for characterising the full array of compound structures displayed in natural language, DS defines in addition the license to build paired trees, so-called *linked* trees, linked together solely by the sharing of terms, such as may be established by encoded anaphoric devices like relative pronouns. Consider the structure derived by processing the string *John, who smokes, left*:

(41) Result of parsing *John who smokes left*:



The arrow linking the two trees depicts the *link* relation. The tree whose node is pointed to by the arrow is the *linked* tree (read $\langle L^{-1} \rangle$ as ‘linked to’). Such linked trees may be conceived of as subproofs over a term shared with the ‘host’ tree, whose content must be established in order to establish some property of that term, which is expressed by the overall proposition. In the above, non-restrictive case, the content of the linked tree merely adds the information that John smokes to the information that he left. But such structures may also provide restrictions on the shared term, as in non-restrictive relative clauses, or constrain the context within which some term is to be construed, as in the so-called Hanging Topic Left Dislocation structures (Cann et al. 2005):

(42) As for John, I dislike his style of painting.

Within any one such linked tree, the full range of computational, lexical and pragmatic actions in principle remain available, depending solely on the type requirements relative to which the pairs of linked structures are developed.

With this flexibility to allow the incremental projection of arbitrarily rich compound structures, the result is a formal system combining lexical, structural and semantic specifications, all as constraints on the growth of trees. As argued in Kempson et al (2001), Cann et al (2005), and Kempson and Kiaer (2010), this leads to the comprehensive DS claim that the syntax of natural languages does not involve a separate level of representation besides what is needed for semantics (Cann et al 2005, and elsewhere), not

because there is no level of semantic representation, but because there is no independent level of syntactic representation.¹³

So, despite the assumption that this progressive build-up of a semantic representation is a basis for doing syntax, syntax in this model is not taken to include a level of representation where there is structure defined over a string of words. DS trees are not inhabited by words and there is no notion of linear ordering expressed on the tree; the annotations on the tree are solely representations of conceptual content. Lexical specifications are defined in exactly the same terms of tree growth, as constraints on tree growth. Such tree growth actions can take place only if the condition triggering these actions matches the decorations on the node which the pointer has reached in the parse – this is a major determinant of word order effects.

A consequence of this methodology is the way concepts of structural underspecification and subsequent update replace the need to postulate multiple types of representation. Through the building and updating of unfixed nodes, a multi-level account of syntax is replaced with progressive growth of a single representational level; and this level turns out in the result to be nothing more than the representation of content, as established from processing the linguistic string in context. The characterisation of lexical specifications in the same terms enables seamless integration of lexical and syntactic forms of generalisation, so that discrete vocabularies for lexical, syntactic or semantic generalisation are unnecessary (and, indeed, precluded).

Constraints taken to be specific to natural-language syntax and not reducible to semantic generalisations are analysed as constraints on the same growth process. For example, the complex NP constraint, which precludes the dependency of an expression outside a relative clause sequence with some site within that relative, is analysed in DS via the licence to build linked-tree pairings. This imposes its own locality restrictions, in terms of limits on the direction of further tree growth. Any expression characterised as decorating an unfixed node, e.g. a relative pronoun,¹⁴ has to be resolved within the tree which that unfixed node construction step initiates. Hence it cannot be resolved in a tree that is merely linked to that tree. Thus, the island constraint is captured, not in terms of notions of subjacency (however realised) that are defined over putative hierarchical structures over strings of words, but in terms of the discreteness of subproofs within an overall proof of the content expressed by such a string.

¹³Analogous arguments apply to morphological structure (Chatzikiyiakidis and Kempson 2011) but we do not pursue these here.

¹⁴A relative pronoun in English is lexically defined to induce a copy of its antecedent ‘head’ at an unfixed node.

5.2 Ellipsis and context

With this sketch of the DS framework, we can now return to ellipsis, and see how a multiplicity of unrelated ellipsis types can be avoided within a system if it articulates a dynamic account of content accumulation. Recall the central problem regarding ellipsis: model-theoretic accounts are too weak to handle syntactic constraints, while syntactic accounts, required to feed interpretation rather than interact with it, freely posit ambiguity. In DS, though, syntax is expressed as growth of representations of propositional content relative to context. Within such a system, VP ellipsis construal and pronoun construal work in essentially the same way. Both project a placeholder for some value with its requisite type specification (provided by the lexical specification of the pronoun or auxiliary), for which the formula value is taken to be provided from the context (Purver et al 2006).

Here, one's notion of context is clearly crucial. In DS, context is an evolving record of representations of meaning plus the process of their building (strictly, a sequence of triples: a decorated structure, a sequence of words and the update actions used in establishing the given structure over that word sequence (Cann et al 2007)). Given this notion of context, any aspect of it is expected to be re-usable as a basis for the construal of ellipsis, and this encompasses all that is required to account for the various kinds of ellipsis.

First there is the availability of meaning annotations from some context tree, re-using a formula just established by a simple substitution process. This direct re-use of a formula from context is illustrated by the strict readings of VP-ellipsis, where the meaning assigned to the ellipsis site matches that assigned to the antecedent predicate (see section 1). In the sloppy readings, where there is parallelism of MODE of construal but not matching of resultant interpretation, it is the structure-building actions that are replicated and applied to the newly introduced subject. (43) provides such a case:

- (43) A: Who hurt himself?
B: John did.

Processing the question in (43) involves the construction of a two-place predicate, as indicated by the verb, plus the construction of an object argument; and then, because this object contains a reflexive pronoun, it is obligatorily identified with the argument provided as subject. Re-applying these very same actions in the new tree for B's reply, gives rise to a re-binding of the object argument to *john*, which already decorates the subject node of the new tree, thanks to the elliptical fragment. The effect achieved is the same as the higher-order unification account but without invoking any mechanism beyond what has already been used for the processing of the previous linguistic input. All that has to be assumed is that the meta-variable con-

tributed by the anaphoric *did* can be updated by some suitable selection of a sequence of actions taken from the context. This license to re-use actions stored in context is equally made use of in anaphora construal, giving rise to the so-called “lazy” use of pronouns (see section 1).

Finally – and now falling within just the same general mode of explanation – there are those cases from dialogue where what the context provides is structure, to which the words of the follow-on speaker provide an extension. Canonical cases of this are question-answer pairs, the answer providing the update to the very structure provided by the question.¹⁵

- (44) A: Who did John upset?
B: Himself.

But this pattern is characteristic of dialogue: quite generally, as we saw in section 1, one speaker can provide words which induce some structured representation of meaning, often one that is in some sense incomplete, to which their interlocutor can provide an extension. Here the modelling of production in DS, using the very same procedures as in parsing, comes into its own. This predicts the naturalness of split utterance phenomena, of which question and answer are a subtype, since both the speaker and the hearer in an exchange are presumed to be building comparable structures, continuing to use the very same procedures whether in production or in parsing. So when an interlocutor switches roles, they continue with the very same structure which they have just been building in the previous role. The immediate consequence of this is that tight coordination between the parties is expected, as is the availability of a procedure of setting up the first part of some structural dependency within one speaker/hearer role which is subsequent resolved in the other role. Thus, ellipsis construal can take aspects from immediate context, whether these be representations of meaning, or the actions used to build such representations, or indeed the partial structures that formed that context. The breadth of effects achieved in ellipsis construal need not be stipulated; it is grounded in the richness of the dynamic, structural concept of context.

Perhaps the most significant part of this rich attribution of structure to context, in relation to issues of representationalism, concerns the interaction of structural constraints with the general process of building interpretations in an evolving, structured context. For example, the supposed island constraints displayed in antecedent-contained ellipsis are naturally captured without any syntactic level of representation. Recall that the complex NP constraint is said to underlie the contrast between (31) and (32), repeated here, in which the construal of the ellipsis site precludes interpretation across any additional relative clause boundary:

¹⁵This is on the assumption that *wh* expressions project a particular form of metavariable (Kempson et al 2001).

- (31) John interviewed every student who Bill already had.
 (32) *John interviewed every student who Bill ignored the teacher who already had.

The crucial point here is the nature of the relative pronoun, which initiates the construal of the expression containing the ellipsis site. From a processing point of view, an English relative pronoun intuitively does two things: it initiates some type of sub-structure and promises a term within that sub-structure that is co-referent with the term that the relative clause sequence is taken to modify. The implementation of this in DS is that the relativiser triggers both the construction of a linked tree and the presence of an initially unfixed node within that linked tree, to which it adds a copy of the term that the relative clause modifies (informally, the ‘head’ of the NP+relative sequence). It then follows from general properties of the structural dominance relations defined in DS that an unfixed node must ultimately be fixed locally within that structure; in particular, it cannot ‘look beyond’ any additional link relation for a place to fix itself (Cann et al 2005). The effect that (in more conventional vocabulary) a relativiser cannot be coindexed with a trace across another relative clause boundary thus follows from the parsing mechanism. The fact that such constraints are apparently non-semantic therefore need not be taken as evidence for a distinct syntactic level of representation: in this new perspective they are locality constraints on tree growth. Accordingly, examples like (32) do not, on a dynamic view, preclude a unified account of ellipsis, even though on other accounts they are taken to be diagnostic of forms of ellipsis requiring syntactic analysis, while others require quite distinct semantic analysis .

Over all, then, the DS perspective indicates the potential to meet the threefold challenge posed by ellipsis: capturing the very diverse forms of interpretation, providing a unitary base from which this diversity can be obtained, and enabling the articulation of a concept of context that is rich enough to make an integrated account of ellipsis possible. We have developed the case of ellipsis here in preference to that of anaphora, not only because it brings out the representationalism issues more strikingly, but also because of the extensive parallelisms between anaphora and ellipsis: both are taken to involve inputs of meaning which are lexically defined as being underspecified with respect to an output meaning, and rely on context to provide the necessary update.¹⁶ This is made possible precisely because the core concept of the natural-language grammar formalism is that of growth of representations of content; and all explanations are accordingly expressed in these terms. The diversity of structure-sensitive interpretation mechanisms

¹⁶We set aside here cataphoric effects, but these can be handled without stipulation in the framework, using appropriate delay mechanisms. Initial lexically determined specifications of type but not content allow for later specification of the content, as with expletive pronouns.

which have on the conventional syntactic account of ellipsis to be expressed as independent structures, hence as stipulations of ambiguity of the strings themselves, can be seen as the application of processes to some underspecified input, with different procedures for growth from that input giving rise to the divergent construal. So the shift into the dynamic perspective is essential to the resulting integrated account.

While here we have pursued the significance of anaphora and ellipsis specifically, we should not lose sight of the endemic nature of context dependence in natural language, which extends far beyond these two particular phenomena. Indeed, context dependence is now recognised as permeating the whole of the lexicon (and not just some quasi-grammatical subset of lexical items such as pronouns and auxiliary verbs). In recent years, important work has begun on the general phenomenon of lexical meaning and the extreme flexibility of construal which content words in context seem to allow. In particular work being developed within Type Theory with Records (Cooper 2005, this volume, Larsson 2008), building on Martin-Löf's type theory, provides a general framework for natural language interpretation which, like DS, takes a basic semantic vocabulary and gives it a proof-theoretic twist, so it is model-theoretically grounded but makes essential use of proof-theoretic dynamics (see Cooper this volume for detailed exegesis). Notwithstanding the proof-theoretic underpinnings of the framework, one important application has been to provide a formal articulation of what it is about a word that enables it both to be the bearer of an identifiable concept of meaning (in some sense) but nevertheless to model the full range of variation in truth conditional contents which the word makes expressible.

We now find ourselves in a situation that turns on its head the general view of linguistic representation from the later twentieth century. The position developed in the early nineteen-seventies onwards was one of indispensable representations of syntactic structure defined over strings of words, over which semantic interpretations can be directly stated. From this position, we have reached one in which syntactic representations are replaced by processes that articulate semantic representations – and these are, in their turn, conceived of as being the only necessary level of representation.¹⁷ So it is syntax, traditionally conceived, that has become superfluous, in the sense of not requiring any special vocabulary other than that of inducing the growth of representations of meaning – a view which is notably close to the ontology of categorial grammar (Morrill 1994), but avoids the Montago-

¹⁷It might be argued that this position can be attributed to the minimalist conception of transformational grammar where LF (Logical Form) is the only level of representation (Chomsky 1995). The difference between this approach and that of Dynamic Syntax is that minimalism retains its structuralist foundations by defining LF to be inhabited by words, categories and their hierarchical combinations, rather than by concepts and their proof theoretic ones. As we have argued, an inability to accommodate the context dependence of natural language phenomena ultimately follows from these foundations.

vian limitations of the latter with respect to issues of context dependence in natural language. Thus, we are reaching the ability to express directly the folk intuition that in understanding utterances of language, it is representations of content that are being built up in context, and that language is a vehicle for recording or expressing our thoughts.

6 Implications for dynamic perspectives

6.1 Compositionality

With this move away from direct mappings of natural language strings onto denotational contents, in much the same spirit as DRT, it is important to address potential objections in the same style as those levelled at DRT by Groenendijk and Stokhof (1991): that the account presented violates the compositionality of meaning for natural languages (see also Dekker 2000, and the response by Kamp 1996).

Compositionality of content has indeed been very generally presumed to be sacrosanct as a working methodology for formal specifications of natural-language semantics. However, consider again the standard construal of compositionality: ‘The meaning of an expression is a function of its component parts and the way they are put together’. Far from being a tight restriction on natural languages, this general form of compositionality is (as has intermittently been noted) extremely weak, requiring at least a specific definition of ‘be a function of’, since functions may delete or arbitrarily re-order symbols or other object language objects.

In essence, the principle of compositionality is an attempt to capture the intuitive idea that meanings are not computed arbitrarily on different occasions, but are constrained to the meanings of the basic expressions of a language (words or morphemes, in natural languages) and some known and determinate means of combining them to construct propositional meanings. A minimal assumption (often unexpressed, but see Cann 1993) is that no mapping from syntax to semantics may delete already established information, so that any compositional function should be monotonic. Additionally, there is typically an assumption that each word contributes something to the meaning of the expression in which it appears. But this notion of ‘contribute’ is also in need some interpretation, given the existence of fully grammaticalised expressions like the complementizer *that* in English (which at most operates as an identity function over the proposition expressed by its associated clause) and pleonastic expressions like *it* in weather verb constructions or expletive uses (like that in (17), above). Arguably, these contribute nothing at all to the interpretation of a string; at most they can be claimed to contribute to some form of ‘constructional meaning’, and this is likely to be of a sort that does not fall under the purview of conventional denotational semantics. In general, therefore, the precise interpretation of

compositionality depends on the theory in which it is taken to apply and to prior assumptions made about the nature of the relations between words, syntax and expressible meanings.

Furthermore, the familiar rule-by-rule concept of compositionality threatens to preclude any characterisation of the systemic potential of all natural language expressions for context-dependent construal. Yet this is arguably the core property of the expressivity of natural languages and as such should be diagnostic of successful characterisations of natural language content. It follows that the common understanding of compositionality has to be modified, if it is to be sustainable as part of a properly explanatory account of natural language interpretation.

In the face of this challenge, we suggest that discussions of compositionality of content for natural language strings have conflated two concepts. There is, on the one hand, the essential contribution to be made by each word to the business of interpretation. This, we have argued, should be conceptualised as a contribution to a process: the progressive construction of a structural representation from a sequence of words. On the other hand, there is the compositionality of the content of each structure that results from that process. In teasing these two notions apart, we have two relatively simple concepts of compositionality, one defined in terms of monotonic incrementality, but lacking any notion of content; the other defined in terms of compositionality of content, but making no reference to the contributions of individual words.¹⁸ The first of these is defined in Kempson et al 2001 (chapters 2-3, 8) as monotonic growth over partial trees. The second requires a concept of semantic content defined for whichever logical system is used to represent the structure of the Language of Thought, taken in DS to be a typed version of the epsilon calculus.¹⁹

In any case, it is important to bear in mind that the criticism of DRT with respect to compositionality primarily concentrated on the introduction of a level of representation that is intermediate between the structure inhabited by strings and the level of assignable denotational content, thus apparently flouting Occam's Razor. On the DS perspective, this criticism is simply deflected, since DS posits but a single level of representation: that of meaning which is, in principle, amenable to model-theoretic interpretation.

What is abandoned in this move is the assumption of there being independent syntactic representations: these are replaced by systematic incremental processes which are defined for all wellformed strings of some language. The order of words as given must induce at least one monotonic

¹⁸Cf. Fodor's (2001) position, in which the representation of meaning (the Language of Thought) is similarly claimed to be the locus of compositionality, rather than natural languages themselves displaying compositionality. Fodor remains inexplicit on the matter of just how words do make systematic contributions to representations of meaning.

¹⁹See Meyer-Viol (1995) for discussion of formal properties of the epsilon calculus in relation to predicate logic.

process of tree growth to yield a complete, and compositionally defined, representation of content as output. Knowledge of language, then, resides in the systematic capacity to build up representations of thought from sequences of words, relative to the context within which such incremental processing takes place.

6.2 Knowing how and knowing that

The view of grammar that we have advocated in this article has implications for another fundamental assumption that underpins linguistic methodology. Conventional models of syntax and semantics assume that knowledge of language consists in ‘knowing that’ language is certain way, and not in any sense ‘knowing how’ to use language. A general distinction between knowing-that and knowing-how was introduced into the philosophy of mind by Ryle (1949). Though Ryle himself advocated a certain kind of knowing-how view, the characterisation of knowledge in a knowing-that sense has overwhelmingly held sway in linguistics and in cognitive science more widely (it has also been dominant in philosophy, though see Bengen and Moffett 2007²⁰).

In contrast to this, the DS conception of grammar, as a process of constructing representations of meaning, clearly constitutes a model of ‘knowing how’, in a certain sense. In some ways, this is a striking departure from long established assumptions in linguistic theory, and this difference is key to the potential of the framework to account both for the general properties of language as an inherently context-dependent system and for more specific phenomena such as ellipsis and anaphora. However, it is important to clarify the particular ways in which DS is a knowing-how approach, in order to pre-empt certain potential objections.

The view of grammar that we have advocated is a knowing-how view in the following sense: knowledge of language consists in having a set of mechanisms for language use, specifically a means of retrieving a meaning from a string of words (or their sounds) uttered in context (and, by the same token, a means of producing a string in context to convey some meaning), reflecting a ‘commonsense’ view of language (see also Phillips 1996, 2003).

²⁰From a philosophical perspective, Stanley and Williamson (2001) claim that the concept of knowledge-how is simply a species of knowledge-that – in both cases, a relation between an agent and a proposition. However, a significant part of the Stanley and Williamson case depends on parallel assumptions within linguistics of overwhelmingly sentence-based methodology, in which putatively nonsentential/nontruthconditional phenomena are analysed wherever possible as sentential structures and truth-theoretic contents respectively. This leads to a near-circularity of reasoning: the conclusion that ‘knowing how’ phenomena are reducible to sub-species of ‘knowing that’ is derived from syntactic and semantic methodologies which presuppose a strictly ‘knowing that’ approach, being themselves reflections of static, truth-based methodology. In addition, in order to explain their distinctiveness, they have to invoke sub-divisions of type of ‘knowing that’ based on pragmatic factors of very unclear provenance, for which no pragmatic theory as currently envisaged would provide substantiation.

Linguistic competence, then, is no more than the possession of mechanisms that make possible the doing of what one does in language performance.

This view ostensibly flies in the face of Chomsky's articulation of the competence/performance distinction, which espouses a knowing-that concept of the capacity for language, and which is wedded to a purely static kind of representationalism, with no reflection of the dynamics of language processing. Yet the suggested shift does not amount to a collapse of a competence model into a performance model, with the loss of necessary scientific abstraction that this would imply. The strict concentration on knowledge of language that Chomsky advocated is not lost simply by reconceptualising the nature of this knowledge. The system as defined in DS articulates a set of constraints on tree growth, a set of principles proposed as underpinning processes of interpretation in real time. It does not bring the APPLICATION of these principles into the domain of competence. There remains a sharp distinction between grammatical knowledge and independent, potentially intractable extra-grammatical factors, though the dynamic approach opens up new and explanatory possibilities for interaction between the two. In so doing, this approach makes available ways of articulating underspecification and update in natural language interpretation. Crucially, the DS formalism makes no attempt to model what determines the particular choices made by language users in resolving such indeterminacies on particular occasions; these matters do indeed belong to performance.

Therefore, the way in which DS constitutes a knowing-how approach does not entail the abandonment of a competence theory, in the sense of a theory that aims strictly at characterising knowledge of language. The notion of competence that we have criticised elsewhere in this article is an altogether more specific notion, encompassing a sentence-based methodology, a commitment to static representations of linguistic structure, and a self-imposed blindness to the pervasive context-dependence of natural languages. One of the lessons of re-assessing the abstract notion of representation in grammatical theory has been to show that a grammar, as a characterisation of knowledge of language, need not have these features.

Just as 'knowing how' in our sense must be distinguished from performance theory, so it must also be distanced from certain other possible associations, in particular any connection to behaviourism. Ryle's original concept of 'knowing how' was grounded in terms of dispositions and this has led to the perception among some analysts that Ryle advocated some form of (quasi-)behaviourism. Whether or not this is a fair reading of Ryle (an issue that we leave to philosophy), it would be wrong to tar all kinds of knowing-how approach with the behaviourist brush. This is arguably what has happened in the development of modern linguistics, in which the adoption of a mentalist perspective on language is often equated, without argument, with a knowing-that conception of knowledge of language. For example, Chomsky (2000, 50–52) contrasts "the conception of

language as a generative procedure that assigns structural descriptions to linguistic expressions”—an intrinsically knowing-that characterisation of linguistic knowledge—with the idea that knowledge of language can be reduced to an ‘ability’. No other possibility is entertained in this discussion, despite the fact that Chomsky actually notes that “*knowing-how* [...] cannot be characterised in terms of abilities, dispositions, etc.” (2000, 50–52; [italics in original]). Logically, this leaves open the possibility of a knowing-how model which characterises not dispositions, abilities and behaviours but the principles underlying them, which is suitably abstracted from the data and which in all other ways satisfies the need for a coherent and tractable account of linguistic knowledge. More generally, one may be committed to the study of language as an abstract system of knowledge but still question the best ways to conceive of and represent this knowledge. Formal modelling of the structural dynamics of language is as far from behaviourism as any more familiarly static representationalist model is.

Nevertheless, one thing that we have tried to convey is the fact that the choice between kinds of knowledge representation is a significant one. Different conceptions of linguistic knowledge are not mere ‘notational variants’ of one another, nor is the choice between them merely a matter of perspective. To the contrary, different approaches have very different empirical coverage, display different degrees of overall theoretical economy and contrast in how they relate to the reality of language as a system that is interpreted relative to context. As we have argued, it is not only possible to define a coherent representationalist system that gives the dynamics of structure a central role; it is also highly desirable from both empirical and conceptual points of view. Therefore, while we would distance ourselves from some aspects of Ryle’s characterisation of knowing how – and certainly from some associations it has gained over the years – we maintain that a knowing-how model, at least in the limited sense that we require, is still very much available for linguistic theory to pursue. There remain two different types of perspective for natural language analysis. One broad type is a characterisation of language as a declarative multi-level body of knowledge for which numbers of stipulations and imposed ambiguity postulations are required even to express the data. The other perspective is one in which a natural language is characterised as a set of mechanisms for proposition construction. Recalcitrant problems of the former perspective promise to dissolve away in the latter.

We close with the observation that our arguments, for all their questioning of conventional methodologies, have led us to a position that is highly intuitive and has a long and respectable (even in some ways Fregean) heritage: that thoughts have denotational content and constitute the sole focus of semantic enquiry. Languages are used to express our thoughts, and these have the content they do through a semantics of resulting structures defined to yield denotational contents grounded in the primitive concepts of individ-

ual and truth.²¹ What is new is the explanation of how language performs this function. Language is neither identical to thought nor ‘mapped onto’ thought via a series of representations of uncertain status. It is instead intrinsically dynamic, a vehicle for the construction of objects over which inference is definable. As a mechanism, it is built to interact with the environment in which it is employed, but is not defined by it. A representationalist commitment, then, is ineliminable. Yet it is minimal, in invoking only the forms themselves and representations over which inferences are derivable from their use. In sum, language is a vehicle for constructing thoughts from the building blocks which our words provide. But, at one and the same time, given its reflection of real-time language use, it is a vehicle that allows interaction with others in the construction of such thoughts, hence a vehicle for the interactive and coordinated construction process that constitutes communication.

References

- Baggio G. , van Lambalgen, M. & Hagoort, P. 2011. Language, linguistics and cognition. This volume
- Bengen, J. and Moffett, M. 2007. Know-how and concept possession. *Philosophical Studies* 36: 31–57.
- Blackburn, P. & W. Meyer-Viol. 1994. Linguistics, logic and finite trees. *Bulletin of Interest Group of Pure and Applied Logics* 2, 2–39.
- Brandom, R. 1994. *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard Univ Press.
- Brandom, R. 2008. *Between Saying and Doing*. Oxford University Press.
- Bresnan, J. 2001. *Lexical-Functional Syntax*. Oxford: Blackwell.
- Cann, R. 1993. *Formal Semantics*. Cambridge, Cambridge University Press.
- Cann, R. forthcoming. Towards an account of the English auxiliary system. In R. Kempson, E. Gregoromichelaki and C. Howes and (eds.) *The Dynamics of lexical Interfaces*. CSLI publications.
- Cann, R., R. Kempson, and T. Kaplan 2005. Data at the Grammar-Pragmatics Interface: the Case of Resumptive Pronouns in English, *Lingua* 115, 1475–1665.

²¹Note that this does not in itself entail a particular concept of individual. Hence this might take the form of the individual constant beloved of Russell or some concept of type $\langle e \rangle$ terms that are constructed for the purpose of inference-drawing to general conclusions, with arbitrary witnesses as their denotation (in the manner of Sommers 1982, Fine 1986, and others).

- Cann, R., Kempson, R & Marten, L. 2005. *The Dynamics of Language*. Dordrecht: Elsevier.
- Cann, R. Purver, M. & Kempson, R. 2007. Context and wellformedness: the dynamics of ellipsis. *Research on Language and Computation* 5, 333-358.
- Carr, P. 2011. The philosophy of phonology. this volume
- Chatzikiyriakidis. S and Kempson R. 2011. Standard Modern and Pontic Greek person restrictions: a feature-free dynamic account. *Journal of Greek Linguistics*.
- Dalrymple, M. ed. 1999. *Semantics and syntax in lexical functional grammar: the resource logic approach*. Cambridge MIT: MIT Press.
- Chomsky 1957. *Syntactic Structures*. Mouton: Dordrecht.
- Chomsky 1965. *Aspects of the Theory of syntax* MIT Press.
- Chomsky, N. 1986. *Knowledge of Language*. Praeger Press, New York.
- Chomsky, N. 1995. *The Minimalist Program*. Cambridge, MA. MIT Press.
- Chomsky, N. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Cooper, R. 2005. Austinian truth, attitudes and type theory. *Research on Language and Computation*, 5, 333–362.
- Cooper, R. 2011. Type theory and semantics in flux. This volume
- Dalrymple, M., Shieber, S. & Perreira, F. 1991. Ellipsis and higher order unification. *Linguistics and Philosophy* 14, 399–452.
- Dekker, P. J. E. 2000. Coreference and Representationalism, In K. von Heusinger and U. Egli (eds.), *Reference and Anaphorical Relations*, pp. 23. Kluwer: Dordrecht.
- Evans, G. 1980. ‘Pronouns’. *Linguistic Inquiry* 11, 337–62.
- Fiengo, R. and May, R. *Indices and identity*. Cambridge, MA: MIT Press, 1994.
- Fine 1986. *Reasoning with Arbitrary Objects*. Blackwell: Oxford.
- Fodor, J.A. 1975. *The Language of Thought*. MIT Press.
- Fodor, J.A. 1983. *Modularity of Mind*. MIT Press.

- Gargett, A., E. Gregoromichelaki, R. Kempson, M. Purver & Y. Sato. 2009. Grammar Resources for Modelling Dialogue Dynamically. *Cognitive Neurodynamics* 3(4).
- Ginzburg, J. forthcoming. *The Interactive Stance: Meaning for Conversation* Oxford University Press.
- Ginzburg, J. & Cooper, R. 2004. Clarification ellipsis and the nature of updates in dialogue. *Linguistics and Philosophy* 27, 297–365.
- Groenendijk, J. and Stokhoff, M. 1991. ‘Dynamic predicate logic’. *Linguistics and Philosophy* 14, 39–100.
- Hamm, F. Kamp, H and Lambalgen van. M. 2006. There is no opposition between formal and cognitive semantics. *Theoretical Linguistics* 32.
- Heim, I. 1982. The semantics of Definite and Indefinite Noun Phrases. Ph.D University of Massachusetts Amherst.
- Heim, I. and Kratzer, A. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.
- Heusinger, K. von, and U. Egli (eds.) 2000. *Reference and Anaphoric Relations*. Dordrecht: Kluwer.
- Hornstein, N., Nune, J. and Grohmann K. 2005. *Understanding Minimalism* Cambridge: Cambridge University press.
- Jackendoff, R. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Blackwell.
- Hans Kamp 1980. Some remarks on the logic of change. In Rohrer, C. ed. 1980: *Time, Tense and Quantifiers*. Niemeyer, Tübingen
- Kamp, H. 1981. A theory of truth and semantic representation. In Janssen, T. & Stokhof, M. (eds.) *Truth, Interpretation, and Information*, pp.1–34. Dordrecht: Foris.
- Kamp, H. and Reyle, U. 1993. *From Discourse to Logic*. Dordrecht: Reidel.
- Kamp, 1996. Discourse Representation Theory and Dynamic Semantics: representational and nonrepresentational accounts of anaphora. Published in French in: F. Corblin, & C. Gardent (eds.) 2005. *Interpréter en contexte*. Paris: Hermis.
- Karttunen, L. 1976. Discourse referents. In McCawley, J. (ed.) *Syntax and Semantics 7: Notes from the Linguistic Underground*, 363–85, Academic Press, New York

- Gregoromichelaki, E., Sato, Y., Kempson, R., Gargett, A., Howes, C. 2009. Dialogue modelling and the remit of core grammar. *Proceedings of IWCS 8*, pp. 128–139.
- Kempson, R. Meyer-Viol, W. & Gabbay, D. 2001. *Dynamic Syntax: The Flow of Language Understanding*. Oxford: Blackwell.
- Kempson, R. and Kiaer, J. 2010. Multiple long-distance scrambling: Syntax as reflections of processing. *Journal of Linguistics*, 46 , pp. 127–192.
- Larsson, S. 2008. Formalizing the dynamics of semantic systems in Dialogue. In Cooper, R. and Kempson, R. (eds.) *Language in Flux – Dialogue Coordination, Language Variation, Change and Evolution*. College Publications, London.
- Lewis, D. 1972. General semantics. In D. Davidson and G. Harman (eds.) *Semantics for Natural Language*. Dordrecht, Reidel.
- Merchant, J. forthcoming. Ellipsis. In Kiss, T., Alexiadou, A. (eds.) *Syntax: an International Handbook, 2nd edition*, Berlin: Walter de Gruyter.
- Montague, R. 1974. *Formal Philosophy: Selected Papers of Richard Montague*, editor R. Thomason. Yale: Yale University Press.
- Moortgat, M. 1988. *Categorial Investigations: logical and linguistic aspects of the Lambek calculus*. Dordrecht, Foris Publications.
- Morgan, J. 1973. Sentence fragments and the notion *sentence*. In Kachru, Lees and Malkiel (eds.), *Issues in Linguistics: Papers in Honour of Henry and Renee Kahane*, pp. 719–751. University of Illinois Press.
- Morrill, G. 1994. *Type-Logical Grammar*. Dordrecht: Springer.
- Morrill, G. 2010. *Categorial Grammar: Logical Syntax, Semantics, and Processing*. Oxford: Oxford University Press.
- Muskens, R. 1996. Combining Montague Semantics and Discourse Representation Theory. *Linguistics and Philosophy*, 19, 143–86.
- Partee, B. H. 1973. Some structural analogies between tenses and pronouns in English. *Journal of Philosophy* 70: 601–9.
- Partee, B. H. (ed.) 1976. *Montague Grammar* New York: Academic Press.
- Phillips, C. 1996. *Order and Structure*. Ph.D thesis. MIT.
- Phillips, C. 2003. Linear Order and Constituency. *Linguistic Inquiry* 34:1.
- Pollard, C. and Sag, I. A. 1994. *Head-driven phrase structure grammar*. Chicago: University of Chicago Press.

- Pullum, G. K. and Scholz, B. 2001. On the Distinction between Model-Theoretic and Generative-Enumerative Syntactic Frameworks. In de Groote, P. Morrill, G., Retoré, C. (eds.), *Logical Aspects of Computational Linguistics: 4th International Conference, (LNAI 2099)*, pp. 17–43. Berlin: Springer.
- Purver, M. 2004. The Theory and Use of Clarification Requests in Dialogue. Ph.D King’s College London.
- Purver, M., Cann, R. and Kempson, R. 2006. Grammars as parsers: the dialogue challenge. *Research on Language and Computation* 4, 289–326.
- Ross, J. 1967. Constraints on Variables in Syntax. Ph.D MIT.
- Ryle, G. 1949. *The Concept of Mind*. Chicago, Chicago University Press.
- Sag, I., Wasow, T. and Bender, E. 2003. *Syntactic Theory: a Formal Introduction*. Chicago: University of Chicago Press.
- Sommers, F. 1982. *The Logic of Natural Language* Clarendon Press. Oxford.
- Stainton, R. 2006. *Nonsentential utterances*. Oxford: Oxford University Press.
- Stalnaker, R. 1974. Pragmatic presuppositions. In Munitz, M. and Unger, P. (eds.), *Semantics and Philosophy*, pp. 197–213. New York University Press.
- Stalnaker, R. 1999. *Context and Content*. Oxford: Clarendon Press.
- Stanley, J. and Williamson, T. 2001. Knowing how. *Journal of Philosophy* 98.8, 325–328.
- Steedman, M. 1996. *Surface Structure and Interpretation*. Cambridge MIT: MIT Press.
- Steedman, M. 2000. *The Syntactic Process*. Cambridge MIT, MIT Press
- van, Valin, R. D. and LaPolla, R. J. 1997. *Syntax: structure, meaning and function*. Cambridge, Cambridge University Press.
- Wittgenstein, L. 1953. *Philosophical Investigations*. Blackwell.