

## Cambridge Books Online

<http://ebooks.cambridge.org/>



The Cambridge Handbook of Pragmatics

Edited by Keith Allan, Kasia M. Jaszczolt

Book DOI: <http://dx.doi.org/10.1017/CBO9781139022453>

Online ISBN: 9781139022453

Hardback ISBN: 9780521192071

Paperback ISBN: 9781107558670

### Chapter

26 - The syntax/pragmatics interface pp. 529-548

Chapter DOI: <http://dx.doi.org/10.1017/CBO9781139022453.029>

Cambridge University Press

# The syntax/pragmatics interface

Ruth Kempson

## 26.1 Linguistic competence: Syntax/semantics/pragmatics

Throughout the last half century since Chomsky 1965, there has been whole-hearted acceptance of the distinction between the linguistic knowledge humans display, their linguistic competence, and the demonstration of their ability in speaking and hearing, viz. performance. Accordingly, concepts of competence have been articulated solely with respect to language *sui generis*, the sentences of any one language being analysable in terms of syntactic and semantic properties that are articulated in formally specified syntactic and semantic rules of grammar. Following a wholly uncontentious formalist methodology that linguistic theories must be expressible as formal models, these rules, together with rules of phonology, constitute that linguistic competence which any speaker is said to possess.

Once this perspective had been set up by Chomsky, it was taken up by semanticists and philosophers, who despite fundamental differences between different positions coincided on the assumption that the semantics of a natural language had to be given in terms of truth conditions for sentences of the language. Many turned to Richard Montague's formal programme for semantics (Montague 1974a) to provide a detailed formal articulation of what it means to attribute truth-conditional content to sentences of natural language. While there was disagreement as to the nature of syntax to be articulated, even amongst those making this move, the rampant display of ambiguity in natural languages led to very general acceptance of the view that syntactic properties of individual sentences had to be distinguished from a characterisation of their meaning. Moreover, early on, there emerged evidence from the so-called island constraints (Ross 1967) confirming the distinction between syntax and semantics, in virtue of the non-reducibility of island-restriction constraints to semantic considerations (Pardee 1976: see also 26.2.2 below). And so, despite major exceptions,<sup>1</sup> it became part of the accepted wisdom that there had to be independent syntactic and

semantic components of the grammar, with semantics providing an articulation of truth-theoretic content expressible, for example, by a model-theoretic formulation (Montague 1974a). Such a view was adopted by all those working within generalised phrase structure grammars (HPSG as it developed: Pollard and Sag 1994) and within LFG (lexical functional grammar) (Dalrymple 2001). Within these frameworks, all aspects of language that couldn't meet the methodological requirement of formalisability were taken to fall outside its remit. In particular, pragmatics, being dismissed as the wastebasket of non-formalisable aspects of language use (Bar-Hillel 1971, Kamp 1978), was taken at the time to be merely part of some poorly understood concept of linguistic performance.

Right from the early 1970s, with the publication of Lewis 1970 as part of this semantic programme, an additional objective has been to reflect the way in which understanding of words in combination systematically depends on aspects of the context in which they are produced. Nonetheless, it has been assumed that this can be made commensurate with the competence/performance distinction and the retention of the view of grammar as inducing a pairing of well-formed strings of the language with interpretation to be defined in some sense independent of considerations of use. Semanticists have indeed taken it as their task to articulate a formal articulation of context and how meanings of expressions combine to determine context-dependent interpretability (Lewis 1970, Kamp 1981, Kaplan 1989a, Kamp and Reyle 1993 and many others); and, increasingly, they have grappled with the need to invoke pragmatic aspects of content in semantic characterisations of compositionality where these involve articulating truth conditions (Partee 1999, Chierchia 2004). Pragmatists, on the other hand, along with psycholinguists, have had to take on board the restriction upon their own theorising that the input to performance models has to be whatever some defined grammar defines as output; and, for pragmatists, in particular, this is presumed to be some defined characterisation of sentence-meaning.

However, as our understanding of the systematicity of context dependence displayed by natural languages has deepened, this clean division between data which are properly within the remit of grammar to explain and those which fall outside it has become increasingly hard to sustain. The narrative history of pragmatics since the 1980s is indeed a story of the struggle pragmatists have engaged with, in virtue of presuming this to be an assumption which they must retain. Leaving aside the concept of particularised conversational implicature, which is relatively uncontentious in being outside what an utterance explicitly conveys (Grice 1975), the burden of determining the syntax/pragmatics interface concerns how much of what an uttered sentence conveys is determined internally to the grammar. There are, on the one hand, the instances of generalised conversational implicature which convey some putative implicature so standardly that they are said by some to constitute a

default meaning of the expression itself (Asher and Lascarides 2003, Jaszczolt 2005). There is also the problem of conventional implicatures which are said to be an encoded aspect of the use of the word even though not part of the truth conditions conveyed by sentences in which it occurs. And since Grice's work (Grice 1989), there have been numerous variants of a broadly Gricean programme (see Horn and Ward 2004). Nonetheless, despite controversies arising from distinct classifications which the various categories have made available, the concept of sentence meaning as some weak specification of truth-conditional content that has to be enhanced in some way in context has remained a core underpinning to nearly all pragmatic theorising (though see Atlas 1988, Recanati 2004a). Even relevance theorists, with their championing of the gulf between what is encoded and what is only inferentially established, sustain as a concept of sentence meaning a concept of logical form for a sentence that constitutes information established by the grammar, such logical forms being taken by hearers as evidence for a process of enrichment (or other modulation) that yields some propositional form (Sperber and Wilson 1986, Carston 2002, Wilson and Carston 2007, Carston forthcoming). This has led to the debate known as contextualism, with Cappelen and Lepore (2005a) arguing against a group they dubbed 'contextualists' and in favour of a very weak concept of meaning corresponding to truth-conditional content for utterance meaning as well as sentence meaning, with advocacy of a much richer speech-act content as the requisite additional broader notion corresponding to everything that utterances can be taken to convey.<sup>2</sup> The ensuing debate between such contextualists and Cappelen and Lepore has led to fierce disagreements; and there isn't agreement between the various protagonists as to what context amounts to. Nevertheless, all contextualists agree (against Cappelen and Lepore) that there is more to utterance understanding than what some concept of sentence meaning can be seen to provide and, furthermore, that the content expressed by an utterance is due to some interaction between information provided by the grammar and pragmatic processes.

What is less commonly noted is that independent evidence is accumulating in syntax and semantics which equally undermines the clean division of labour between what the grammar formalism provides as output and what pragmatic processes determine. The phenomenon of ellipsis illustrates this particularly dramatically. Informally, ellipsis occurs when the required interpretation is recoverable from context without need of any explicit overt expression. Indeed, precisely because of the total lack of explicit verbal expression other than something to trigger the elliptical construal, ellipsis arguably provides direct evidence of what context amounts to. This makes ellipsis directly pertinent to the context-dependency debate, hence also a window on how the grammar/pragmatics interface should be articulated. One might then think that one could use comparative success in explaining ellipsis in an integrated way as a criterion for evaluating accounts of

the grammar/pragmatics interface. However, this has not been realisable in practice. Not only do all accounts of ellipsis articulated from orthodox assumptions fail this criterion of evaluation but none even project ellipsis as providing any such yardstick. This is due to restriction on the remit of the grammar to sentence-internal properties only. Any phenomenon which is displayed both sentence-internally and across sentential boundaries will be bifurcated within this methodology; and ellipsis is certainly one such. Accordingly, some aspects of ellipsis are characterised grammar-internally, but others as an unrelated discourse phenomenon, with no unitary characterisation of ellipsis. Moreover, this problem is not an ellipsis-specific phenomenon. It applies equally to all context-dependent phenomena – ellipsis, anaphora, tense construal, domain selection etc. So it is a problem that lies at the core of explaining what it is that natural-language expressions encode that enables them to allow such flexibility in interpretation in context. As we shall see, there is a way to reinstate some reflection of the folk concept of ellipsis as being a window on the concept of context, but it involves redrawing the grammar/pragmatics boundary. And, in this chapter, we sketch the arguments as to why considerations of ellipsis lead to such a conclusion.

## 26.2 Ellipsis: Syntax vs semantics vs pragmatics

Taking a history-of-linguistics perspective, ellipsis is a remarkably accurate barometer of the way linguistic argumentation has developed over the past fifty years. First, ellipsis was presumed to be a syntactic phenomenon (Ross 1967); then a more inclusive semantic basis was identified (Dalrymple *et al.* 1991); and, most recently, pragmatic forms of ellipsis have been identified which resist either syntactic or semantic characterisations (Stainton 2006, Cann *et al.* 2007).

### 26.2.1 The syntactic basis for ellipsis

Ellipsis as an observable natural-language phenomenon was taken from early on in this period to fall within the remit of a competence theory of language. Conjoined sentences were used to display different elliptical forms in the second conjunct, whose interpretation in some sense matched that of the first conjunct; and with evidence that different forms of ellipsis were subject to somewhat different structural constraints, the distinct forms were taken to motivate distinct analyses involving distinct structures. For example, in transformational grammar (Ross 1967), deletion operations were proposed as part of syntax in order to yield the truncated form of the second elliptical conjunct by a process deleting that structure under some condition of identity with the structure of the antecedent conjunct (what has come to be known as PF deletion) (PF = phonological form):

- (1) John saw Mary, and so did Bill. (VP ellipsis)
- (2) John ignored Mary, and Tom Sue. (Gapping)
- (3) John ignored Mary. Tom too. (Stripping)
- (4) John ignored someone, but I don't know who. (Sluicing)

As part of this, there was recognition of a need to have mechanisms keep track of the bases of interpretation that were invoked through coindexing of expressions:

- (5) John<sub>i</sub> washed his<sub>i</sub> socks. And Bill did too./Bill too.

But problems emerged with the presumption that sentence strings should have indexing as part of the datum indicating mode of interpretation. Under this assumption, elliptical forms need two distinct indexings, even for a single interpretation. In particular, as (5) shows, the first sentence of (5) gives evidence of needing more than one type of indexing even under the indicated interpretation. This is because the elliptical fragment supposedly displaying the structure of the antecedent clause from which its interpretation is based can be construed either as Bill washed his own socks (the so-called 'sloppy' interpretation of ellipsis), or as Bill washed John's socks (the 'strict' form of ellipsis), thereby indicating two discrete indexings of its antecedent string, and hence ambiguity of structure in that antecedent relative to the one interpretation. Thus, on any such grammar-internal indexing, there have to be two distinct structural sources for the single form of interpretation of that antecedent string and hence deletion in the derivation of the second conjunct of distinct structures (Fiengo and May 1994 and many others). And this isn't the end of the problem, as sentences such as these in their turn led to the recognition that ellipsis is not simply a matter of deleting words in a string, even as indexed, as the appropriate construal may need to be grounded in some replacement of words, as in:

- (6) A: You're sitting on my chair. B: No I'm not.

What A and B are disagreeing about is whether B is sitting on A's chair, irrespective of the pronouns each uses to express that. Moreover, since syntax is defined in exclusively hierarchical terms, there is no basis for imposing a restriction of strict parallelism between the structure under deletion and the antecedent: this is granted to require independent stipulation (Fox 2002).

### 26.2.2 The semantic basis for ellipsis

This postulation of multiple ambiguities, and unclarity as to the level of structure over which the required concept of syntactic identity is defined, led semanticists in response to explore the use of formal-semantic tools for the projection of content directly from the surface sequence of the elliptical expressions. Dalrymple *et al.* (1991), in a very influential paper, defined a mechanism of construal which applied directly to the surface syntactic

structure displayed by the fragment expression itself to yield its interpretation, on the basis of the denotational content of the previous conjunct in the paired coordinate structure. The VP ellipsis site, that is, is seen as projecting a predicate variable whose value has to be identified with some predicate constructible from the antecedent conjunct:

(7) John sneezed and Bill did too.

The core idea is that ellipsis involves a semantic equation involving unification (called ‘the higher-order unification account’ because it involves unifying predicates). The mechanism for achieving it is to construct some appropriate lambda-defined predicate on the basis of the derived content of the antecedent conjunct, by binding some position or positions within it. This involves applying an abstraction operation to the content of the antecedent conjunct, *John sneezed*, to yield a predicate abstract that could be applied to the parallel subject, Bill, in the second, ellipsis-containing conjunct. A possible solution of the requisite equation for (7) that provides a predicate to apply to the content assigned to Bill would be as follows:

(8)  $P = \lambda x.Sneeze'(x)$

The abstract  $\lambda x.Sneeze'(x)$  is then predicated of Bill, yielding the overall parallel construal of both conjuncts. This process is not unrestricted: all selected abstracts must involve a presumption of parallelism between first and second conjuncts, and this must involve one ‘primary’ argument – the subject; and either only that or the subject plus all occurrences of any pronouns construed in that antecedent clause as picking out the subject are replaced in the construal of the second conjunct. Accordingly, (9) is ambiguous according to whether Bill is thinking of taking John’s mother to John’s sister (the *strict reading*) or whether he is thinking of taking his own mother to see his own sister (the *sloppy reading*), but there are not more readings than this:

(9) John is thinking of taking his mother to see his sister, and so is Bill.

(10) = ‘John is thinking of taking John’s mother to see John’s sister and  
Bill is thinking of taking John’s mother to see John’s sister’

OR

(11) = John is thinking of taking John’s mother to see John’s sister and  
Bill is thinking of taking Bill’s mother to see Bill’s sister.

This particular parallelism between conjuncts and the specification that the subject must be involved are independent stipulations in this account. Nevertheless, if we grant these stipulations, the result is, as can be seen in the strict/sloppy interpretations, that we can derive non-identical resolutions for ellipsis from a single semantic content, a clear advantage over any syntactic account.

However, there is reason to doubt whether a semantic explanation could ever be complete. Not only are there problems about more complex VP ellipsis cases, but there is evidence that at least some instances of ellipsis must be reconstructed syntactically. This is because some cases display sensitivity to the island constraints that, in holding over structural configurations, are taken to be diagnostic of there being such structure. The most robust of these is the so-called Complex NP constraint (Ross 1967), which precludes co-dependency of some argument position within a relative clause and an expression external to that relative clause; and it is this restriction to which the case of ellipsis called antecedent-contained deletion (antecedent-contained ellipsis on the semantic form of explanation) appears to be subject to. That is, in this type of ellipsis, the ellipsis site itself is part of a relative, appearing to be recoverable from the very predicate within which it is contained, as in (12). However, what it cannot allow is any dependency across an additional intervening relative clause boundary; and this is taken to indicate sensitivity to the Complex NP constraint. This constraint cannot be expressed within a higher-order unification account because the unification operation is defined within lambda calculus terms, hence over denotational contents, and cannot make reference to structure-specific details.

- (12) John interviewed every student that Bill had.  
 (13) \*John interviewed every student that Bill ignored the teacher who had.

Hence the granting by semanticists that not all syntactic generalisations are reducible to semantic ones.

The sensitivity of some ellipses to such ‘strong island’ constraints is far from being the only type of syntactic specification to which elliptical construal has to be sensitive. In many languages, the fragment provided may bear a syntactic specification which the antecedent has to match in order to provide a well-formed pairing of fragment and content. For example, German elliptical fragments must bear the case appropriate to the verb and syntactic position that has to be reconstructed in order to resolve the ellipsis:

- (14) *Hat er nicht den Brief geschrieben?*  
 has he<sub>NOM</sub> not the<sub>ACC</sub> letter written  
 ‘Didn’t he write the letter?’  
*Nein. Ich/\*Mich.*  
 no I<sub>NOM/IACC</sub>  
 ‘No. I did.’

Thus, whatever the basis for ellipsis, it is seen as having to be sensitive at least to syntactic structure, and in languages with rich case morphology, also to morphological form.



### 26.2.3 Pragmatic forms of ellipsis

Finally, there are cases which appear to resist either syntactic or semantic accounts of ellipsis, requiring an independent pragmatic form of explanation. These are fragment expressions which do not fall into the pattern of using some clausal antecedent from which to build up interpretation (as both syntactic and semantic accounts require). Rather, these are freely interpreted from the utterance scenario directly:

- (15) A (coming out of the lift): McWhirters?  
 B: Second on the left.
- (16) A (seeing a woman enter): Sue's mother.

On the basis of evidence such as this, it has been argued in detail by Stainton (2006) that fragments of this type have to be seen as subsentential assertions. In consequence these cannot be taken as either syntactically or semantically of the type that corresponds to propositions, and the grammar itself is not sufficient to license them: some form of pragmatic reasoning has to be involved. There have been counter-arguments to this view proposing complex covert syntactic structure (Merchant 2010), but even if the force of Stainton's specific argumentation for such cases is accepted, it should be noticed that Stainton only takes his arguments to apply to these arguably peripheral cases, leaving all centrally argued cases of ellipsis intact as different forms of grammar-internal specifications, hence ambiguity. In sum, the consensus is that ellipsis is not a unitary phenomenon. Indeed the disparate nature of ellipsis construal has been graphically labelled its 'fractal heterogeneity' (Ginzburg and Cooper 2004). Ellipsis apparently makes use of whatever information the grammar may provide, with morphological, syntactic, semantic, even phonological information yielding different bases for ellipsis, a phenomenon suggestive of the need to employ a rich multi-level form of analysis.

The problem for all these accounts is that there is no commitment to an integrated explanation of the phenomenon of ellipsis itself. To the contrary, there is a plethora of ambiguities, apparently discrete structures, and failure to see just how context might be seen as providing input in each case of ellipsis construal. As with other context-dependency phenomena, it would seem that some forms of context-dependency can be characterised within the remit of the sentence-based grammar, while others cannot, so the phenomenon itself is necessarily bifurcated. The net effect is that context-dependency – arguably the core datum to be captured within any account of natural language understanding – continues to lack a principled explanation.

An alternative, more radical view that we now turn to is to take ellipsis and the range of types of construal it gives rise to as the right set of data for exploring context dependency as manipulated by natural languages. With this in mind, we turn back to a descriptive classification of the types of ellipsis construal to see what the range of data are.

### 26.2.4 Ellipsis: towards a unitary account

It is uncontroversial that ellipsis construal can become available through reiteration of content directly, for example, in strict interpretations of VP ellipsis where the construal of the antecedent VP and that of the ellipsis site are identical:

(17) John saw Mary, and Tom did too.

In such cases, the immediate context for the construal of the fragment provides a predicate content for re-use without modification.

There are, in addition, cases where the fragment is an add-on to what is in the context, building on what has been started in the context. Question-answer pairs might be seen in this light:

(18) A: Where are you going?  
B: To London.

The interpretation of (18) contains the construal of the entire containing structure of the question modulo the replacement of the *wh*-term by what is proffered as its answer. Question-and-answer exchanges are in fact illustrative of a very broad phenomenon displayed in dialogue, where speakers freely take over from their conversational partner, switching easily between the role of hearer and the role of speaker, with no restriction as to whether or not what is provided in context is a complete sentence in its own right as in (19) or not, as in (20).

(19) A: We're going to Casa Plana.  
B: To show my mother what we've done there.

(20) A: We're going to . . .  
B: (to) Casa Plana.

Just as in *wh*-question-and-answer pairs, the fragment in such subsentential cases has to be seen as a development of the structure made available in the immediate previous context. Each party is simply extending the structure they have just processed as a speaker/hearer, using that as their point of departure for the other processing mechanism to which they are switching.

Then there are cases where it appears to be only the *PROCESS* of building up interpretation that is replicated from context, leading to a different content, but established in the very same manner. These are the sloppy construals of VP ellipsis:

(21) John washed himself. Sue refused to, until I told her she must.

As we have already seen, it is these that appeared to yield the need of ambiguity of structure for the antecedent in structural explanation. To add to the complications, there is interaction between such sloppy and strict ellipsis construals. For example, a single basis for interpretation established as

sloppy can also be used to provide a strict construal of what follows as in an interpretation of (22) in which Harry thinks that Harry is a fool but Harry's wife doesn't think that Harry is a fool:

(22) John thinks he's a fool. Harry does too, although his wife doesn't.

For such an interpretation, the first ellipsis site has to be sloppily construed, and on the syntactic account represented as such; but the second ellipsis site then presumes on that first ellipsis site taking the form of a strict construal. If we are to make sense of this interaction between types of interpretation, and yet retain an integrated perspective on context as providing the wherewithal in each case to establish the content of the fragment, we have to have a concept of context which is rich enough to encompass all of these as an integral part. In particular, such an account has not merely to allow for the attribution of structure to such a context, but also the dynamics of how interpretation is built up. For this, the novel framework of Dynamic Syntax provides a candidate formalism, for the core of the notion of structure it articulates is this very dynamics.

### 26.3 Dynamic Syntax


Dynamic Syntax (DS) is an avowedly representationalist model of interpretation and interpretation growth. It is a model of how interpretations, represented as binary tree-structures of predicate–argument form, are built up relative to context, and 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. The core concept is that of underspecification and its update, with underspecification of structure as well as of content. Indeed, this process of building up structure is taken to be what constitutes the syntax of natural-language grammar. With the dynamics of structural growth built into the core grammar formalism, natural-language syntax is defined as a set of principles for articulating growth of such structures. Syntactic mechanisms are thus meta- to the representations themselves: they are procedures that define how parts of representations of content can be incrementally introduced and updated. Furthermore, all procedures for structural growth are defined relative to context; and context is defined to be just as structural and 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; see also section 26.4 below). The bonus of such explicit adoption of representationalist assumptions and the shift into a perspective in which the grammar reflects key properties of the dynamics of how language processing takes place is, as we shall see, that a natural basis for a novel grammar/pragmatics articulation emerges that is fully commensurate

with an integrated account of ellipsis and context-dependency more generally.

### 26.3.1 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, 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 (23) constitutes some final end result: it is a tree in which the propositional formula itself annotates the top node, 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 (see below). The input to that process is an initial one-node tree (as in the tree representation to the left of the  $\mapsto$  in (23)) which simply states the goal as the requirement to a formula of appropriate propositional type (shown by  $?Ty(t)$ , the ‘?’ indicating that this is a goal not yet achieved):

(23) John upset Mary.



Parsing *John upset Mary*

These DS trees are invariably binary, and, by convention, the argument always appears on the left branch, and the functor on the right branch (a pointer *pointer*,  $\diamond$ , identifies the node under development). Each node in a complete tree is decorated not with words, but with terms of a logical language, these being subterms of the resulting propositional representation. The parsing task is to use both lexical input and information from context to progressively enrich the input tree to yield such a complete output following general tree-growth actions.

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).<sup>3</sup> The language of LOFT makes available not only a vocabulary for describing fixed tree relations, but also a basis for defining concepts of structural underspecification. Concepts of *dominate* and *be dominated by* are defined (using Kleene star operators), indicating some possible sequence of mother relations, or conversely a possible sequence


of daughter relations; and these can be licensed even before there is some fixed number of such mother or daughter relations. For example,  $\langle \uparrow_* \rangle T n(a)$  is defined as a decoration on a node indicating that somewhere dominating it is the node  $Tn(a)$ .<sup>4</sup> 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. This structural underspecification is analogous to the more familiar underspecification displayed by anaphoric expressions, which are taken to project place-holding, metavariable formula decorations, to be substituted by pragmatic substitution actions from context. A second core concept in the explanation is that of requirements for update. This is essential to get appropriate reflection of the time-linearity involved in building up trees in stages (partial trees). For every node, in every tree, all aspects of underspecification are twinned with a concept of *requirement*,  $?X$ , for any annotation  $X$  on a node; and 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,  $?Ty(t)$ ,  $?Ty(e)$ ,  $?Ty(e \rightarrow t)$  etc.; tree-node requirements,  $?\exists x Tn(x)$  (associated with underspecified tree-relations), and formula requirements  $? \exists x Fo(x)$  (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. In particular, the long-distance dependency effects which, since the late 1960s, have been taken by most to be diagnostic of a syntactic component independent of semantics are recast in terms of structural underspecification plus update. For example, when first processing the word *Mary* in (24) below, which is initially construed as providing a term whose role isn't yet identified, the parse is taken to involve the application of a computational action which introduces from the initial root node decorated with  $?Ty(t)$ , a relation to that top node which is *UNDERSPECIFIED* at this juncture, identifiable solely as dominated by the topnode, and requiring type  $e$ , i.e. with requirement  $?Ty(e)$ :

(24) Mary, John upset.

This enables the expression *Mary* to be taken to decorate this node: this is step (i) of (25).<sup>5</sup> Accompanying the underspecified tree relation is a requirement for a fixed tree-node position:  $? \exists x. Tn(x)$ . The update to this relatively weak tree-relation becomes possible only after processing the subject-plus-verb sequence, which jointly yields the two-place predicate structure as in step (ii) of (25). 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:

(25) Parsing *Mary, John upset*:




This process feeds into the ongoing development in which, once all terminal nodes are decorated, bottom-up application of labelled type deduction leads to the completed tree indicated in (23). Such an account of structural underspecification and update is not contentious as a parsing strategy; what is innovative is its application within the grammar mechanism as the basic underpinning to syntactic generalisations.

This account might seem in principle skewed by focusing on parsing, but this is only superficial. Production also 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 by rules of the system. For the production of both (23) and (24), for example, each selected strategy for update has to be checked for subsumption with respect to the goal tree representing the content to be conveyed. These indeed share such a goal tree, illustrating how more than one sequence of strategies is licensed for any string–content pairing, both in parsing and production (to the advantage of real-time processing: Ferreira, V. 1996). So parsers and producers alike use strategies for building up representations of content, either to establish interpretation for a sequence of words, or to find words which match the content to be conveyed.

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

(26) Result of parsing *John, who smokes, left*:



The arrow linking the two trees depicts the so-called link relation. The tree whose node is pointed by the arrow is the *linked* tree (read  $\langle L^{-1} \rangle$  as 'linked to'). Within any one such linked tree, the full range of computational, lexical and pragmatic actions remain available;<sup>6</sup> and 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 others, 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, not because there is no level of semantic representation, but because there is no level of syntactic representation other than that of growth of semantic representation.<sup>7</sup>

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 over a string of words. These trees are not inhabited by words and there is no notion of linear ordering expressed on the tree. Furthermore, lexical specifications are defined in exactly the same terms of actions inducing tree growth, and these actions can take place only if the condition triggering these actions matches the decorations on the node at which the pointer has got to in the parse. So all structural restrictions are stated in terms of the interaction of constraints on tree growth.

A consequence of this methodology of incorporating the dynamics of incremental growth into the syntactic formalism itself is the way concepts of structural underspecification and subsequent update replace the need to postulate multiple levels of representation. The building of unfixed nodes and updating them replaces a multi-level account of syntax with progressive growth along a time line towards just one type of representation, hence a single representational level. 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 and syntactic generalisation are precluded. And constraints that, in other frameworks, are 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 associated with a precluding of dependency of some expression outside a relative clause sequence with some site within that relative is analysed in DS via the locality imposed by the licence to build linked-tree pairings. Any expression characterised as decorating an unfixed node, e.g. a relative pronoun,<sup>8</sup> has to be resolved within the tree which that unfixed node construction step initiates. Hence it cannot be resolved in some tree only linked to that tree, and the island constraint is captured, albeit in less familiar terms than is standard.

Such a system might appear to face the challenge of characterising quantification, often thought to constitute a second core case where the syntax of natural languages is disjoint from what is required for the semantics of

quantification. Notoriously, no natural language overtly displays quantification following the pattern of predicate logic's propositional quantifying operators. But in this framework, this problem is addressed by grounding the account in the so-called epsilon calculus. This is a logic that provides the formal account of the so-called arbitrary names of natural deduction systems for predicate logic. The heart of such names is that their syntax is simple: they are a naming device like all other individual-denoting expressions of the logic.<sup>9</sup> It is the semantics for such names that is complex, for they are terms denoting witness sets for the entire proposition in which they occur; and this means that a rule of semantic evaluation is defined to determine, as output, that their internal structure reflects the environment in which they occur.<sup>10</sup> There is thus a concept of growth in this aspect too, in growth of the restrictor from what the incremental structural process provides (e.g. that projected by the nominal) and that of the predicate structure within which it is contained (see Kempson *et al.* 2001, Cann *et al.* 2005 for all details). Details aside, the bonus of this account in relation to multiplicity of levels of representation for natural-language grammar-writing is that another supposed dis-symmetry between natural-language syntax and its required semantics dissolves upon analysis, for the account presumes that these are terms of the same type *e* as all other argument expression.<sup>11</sup>

Overall then, the system involves but a single level of representation, the need for multiple levels replaced by the concept of growth of partial representations, these representations themselves being part of a denotationally interpretable system. In particular, the apparent multiplicity confronted by all grammar formalisms which posit independent, statically defined, syntax and semantics is resolved through articulation of the dynamics of how the one type of representation is incrementally built up. An immediate consequence is that the system is unencapsulated. In any one application to yield some derivation, application of general computational actions in interaction with lexical actions as driven by the sequence of words may be interspersed with pragmatic actions of substitution as the carrying out of the lexical actions creates underspecified formula values requiring update; and well-formedness is defined as the availability of at least one possible sequence of actions through from initial goal to some completed propositional output with no requirements outstanding, having used all the words and their actions in order.

## 26.4 Ellipsis as a window on context

With this bringing together of syntax and semantics all reflecting the dynamics of how interpretation is progressively built up, the folk intuition about ellipsis can be modelled directly, opening up a whole new perspective on the syntax/pragmatics interface. For ellipsis can now be seen as making use of the different facets of context which the evolving build-up of



interpretation gives rise to. The problem about ellipsis, recall, is that model-theoretic accounts were too weak to handle syntactic constraints, and that syntactic accounts freely posit ambiguity. In DS, though, syntax just is the growth of representations of propositional content as established relative to context, and this together with context is an evolving record of representations of content plus the process of their building. More formally, a DS parse state is a triple of a sequence of words so far parsed, a (partial) structure and actions used to construct that structure. Accordingly, *context*, as a record of how such parsed states have developed, is a sequence of parse states each made up of a sequence of words, a complete or partial tree and the sequence of actions used to develop that structure (see Cann *et al.* 2007 for a formal definition).

Given this notion of context, any aspect of it is expected to be re-usable as a basis for construal of ellipsis, whether representations of content, actions used to induce some structure, or the structure itself. First there is the availability of content annotations as made available in some context tree, re-using a formula just established by a simple substitution process in the manner of anaphora. This direct re-use of a formula from context is illustrated by the strict readings of VP ellipsis, where the content of the ellipsis site matches that assigned to the antecedent predicate (see section 26.2). In the sloppy readings, where there is parallelism of mode of construal but not matching of resultant content, it is the actions that are replicated, applied to the newly introduced subject. (27) provides such a case.

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

The processing for the question in (27) involves the construction of a two-place predicate as indicated by the verb; 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 whose subject node has been decorated by the expression John of the elliptical fragment gives rise to the construal of the answer as involving a re-binding of the object argument to this new subject. The effect achieved is the same as the higher-order unification account but without anything beyond what has already been used for the processing of the previous linguistic input. All that has to be assumed is that the metavariable contributed by the anaphoric *did* can be updated by some suitable selection of some action sequence taken from the context. Finally there are the cases where what the context provides is structure, to which the follow-on speaker provides an add-on. Canonical cases of this are question-answer pairs, the answer providing the update to the very structure provided by the question.<sup>12</sup>

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

Indeed this is the phenomenon so characteristic of dialogue: quite generally, as we saw in section 26.1, one speaker can provide a structure, often one that is in some sense incomplete, to which their interlocutor can provide an extension. So, as expected, the diversity of ellipsis effects matches the richness of dynamically evolving contexts.

With this definition of syntax as the dynamics whereby interpretation is built up, problems that apply to other accounts of ellipsis do not apply to a DS form of analysis (see Cann *et al.* 2005, Purver *et al.* 2006). The ability to shift from sloppy construal to strict (and even back again), as in (22), is predicted to be possible because the context, in evolving along with the content, keeps at each stage a record of content (formula decorations), structure (the emergent tree representation) and actions (the retained record of the growth process). Even the supposed island constraints displayed in antecedent-contained ellipsis turn out to be expressible in view of the fact that the relative pronoun is morphologically present in the initiation of the construal of the expression containing the ellipsis site:

(29) John interviewed everyone who Bill hadn't.

It is the characterisation of this relativiser as decorating an unfixed node (dominated within the newly emergent structure) which determines that the resolution of this structurally underspecified relation must be satisfied within that individual structure (Cann *et al.* 2005), and not any property of the ellipsis site itself.

Moreover, the account has, as a bonus, the prediction of seamless switching between speaker and hearer roles that is diagnostic of conversational dialogue.

- (19) A: We're going to Casa Plana.  
 B: To show my mother what we've done there.
- (20) A: We're going to . . .  
 B: Casa Plana.

Unlike other frameworks, for which such split utterances pose very considerable problems (see Gregoromichelaki *et al.* 2009), on the DS account this phenomenon is predicted to be wholly straightforward, indeed their existence is a consequence of the DS account of production (Purver *et al.* 2006). According to the DS account, the very same mechanisms are used in production as in parsing. Tight coordination between the parties is expected. Each party is building up structure relative to their own context, so at any point, making use of that individually constructed representation whether as parser or producer, they can switch roles and take over the other role, the only difference between the two activities being the greater specificity of the goal to be achieved in production. So, even with a role switch and the first and second person pronouns having to be reinterpreted, the mechanism for processing them remains identical.

- (30) A: Did you give me back  
       B: your penknife? It's on the table.
- (31) A: I heard a shout. Did you  
       B: Burn myself? No, luckily.

As these display, the context used by a participant as a producer/hearer is exactly that of the context they use in their shifted role as a hearer/producer. So with the incorporation of the dynamics of structure built in to the grammar itself, a very considerably larger data set becomes characterisable. The split utterances, so signally ignored in accounts of ellipsis that purport to be a sub-part of sentence-based grammar, become core data, relative to which competing grammars can be compared.

## 26.5 Redefining the syntax/pragmatics interface

To see why the ellipsis issue has led to the consequence that natural-language grammars demand a dynamic perspective, we need to look back on the root of the problems posed by ellipsis. The heterogeneity of ellipsis arose because, despite the attempt to see the concept of context as grounded in some operation of lambda-abstraction on some model-theoretically definable concept of content, there appeared to have to be invocation of some independent concept of syntax. This was because restrictions on availability of elliptical forms of construal, in at least some instances, had to be defined in terms of the so-called island constraints. But these constraints were set up in order to explain structural co-dependencies between expressions that are discontinuous within some string, and so were not definable either in terms of the linear sequence of words themselves or in terms of their attributable content. Moreover, these structural co-dependencies across discontinuous sequences and their formulation independent of any level of semantic characterisation have special significance, as, in virtue of not being reducible to any domain-general form of explanation, they were taken as a basis for innately specified, and encapsulated, forms of syntactic architecture as the core of the syntax module. With these being apparently applicable to at least some forms of ellipsis, ellipsis too was taken to fall within the remit of syntactic explanation, hence within the grammar. And with ellipsis being seen as, at least in part, grammar-internal, the phenomenon of ellipsis was presumed to be bifurcated.

One side effect of this conclusion was that ellipsis played no role in the contextualism debate as driven by Cappelen and Lepore 2005a. However, once we incorporate the concept of growth into the grammar mechanism itself, with its ancillary notion of underspecified structural relations and growth of all aspects of such structure as part of the process defined by the grammar, then the mapping of a linear sequence of words onto some corresponding semantic representation is seen as a property of incremental growth along

the timeline of processing such linear sequences. With this move, the supposedly innate uniquely determinative properties of human language are seen in terms of progressive growth of semantic representations, and not some separate domain-specific capacity. Moreover, if syntax is defined to be a system of such procedures, a record of how such representations are built up yields precisely the right degree of richness for the concept of context for such processes of construal. Context simply is a record of previous parse states, hence a record of how progressive transitions from one partial structure to the next yield the current parse state. With this concept of context, an integrated account of ellipsis becomes available, notwithstanding the very considerable diversity in contents for an individual string, and its display both within and across sentence and utterance boundaries. This is a big advance over all sentence-based grammar formalisms, since these cannot do more than list the various forms of ellipsis without further explanation. On the DS characterisation, the range of diversity displayed in forms of ellipsis matches exactly the richness of context to be invoked, viz. building up on the basis of established content, established process of interpretation or established structure. However, with this recognition of what context amounts to, the simple setting aside of phenomena such as ellipsis as of no consequence to the contextualism debate is no longer warranted. To the contrary, ellipsis data have a key role to play in the contextualism debate, as they provide such a clear window on the requisite concept of context. But the account of ellipsis that emerges on the DS perspective demands a dynamic basis to the grammar, for it is this which makes possible the integrated explanation of ellipsis as a grammar-internal mechanism, despite contextual provision of values to be assigned to the ellipsis site.

This account of grammar opens up a radically new perspective on the interface between syntax and pragmatics. With syntax now envisaged as a system of procedures for building up representations of content, the implementation of any such procedures can only take place in interaction with whatever pragmatic, grammar-external constraints there may be that determine how selections from context are made. And, of fundamental significance, the concept of encapsulation associated with any such explanation breaks down. Lack of encapsulation is essential, as the choice mechanism has to pick out some value from the particular context and, whatever the form of this mechanism, it is one that is subject to wholly general constraints that apply to all cognitive processing. Furthermore, on this view, there is no fixed interface at which the system of natural-language syntax stops and pragmatic mechanisms take over. Pragmatics simply is the articulation of constraints that determine how mechanisms definitive of natural-language syntax are implemented. Thus the concept of an interface between syntax and pragmatics as some fixed level of representation constituting the feeding relation into implementation of some pragmatic procedures no longer holds. In its place, we have the articulation of grammar as a set of mechanisms making language processing in context possible, a theory of pragmatics being the

articulation of constraints that determine what and how particular choices are made. This interface can take place at any point in the construction of propositional forms from the sequence of expressions provided.

In closing, it is worth noting that core pragmatic phenomena such as conversational implicature do not pose a problem of principle for this account. With grammar defined as providing the architecture in virtue of which humans can build up propositional forms,<sup>13</sup> nothing excludes the characterisation of inference in terms of building further propositional structures leading to additional derived information, both constructed propositional form and additional structures combining through steps of inference defined over the logical formulae which the processing mechanisms license. The mechanisms for such definition are available in the Dynamic Syntax framework in the form of linked tree structures. Moreover, from this perspective, lexical items in language might specifically guide the type or direction of inference to be drawn by enrichment of the minimal context. The result is that the framework retains a competence/performance gap, though radically narrowed; and the grammar specifications may interface with general pragmatic/cognitive constraints at every step of the understanding process.