7
Producing utterances

7.1 Introduction

7.1.1 *Preview*

In this chapter we shall see how far the naturally displayed evidence from language production (which we reviewed in ch. 3) can be used to establish components of a model of language production.

In our exposition, we shall start with the model proposed by Garrett (1982): this is a convenient starting point for us because it represents the culmination (thus far) of the error-based model discussed in Garrett's earlier work, but also attempts to pull together insights from the work of Fromkin, also based on error data, and that of Goldman-Eisler and her associates (Goldman-Eisler 1968; Beattie 1980; Butterworth 1980b) on the evidence of hesitation phenomena. In section 7.2 we start looking more closely at the topmost, or 'message-structure' level. In this, we make contact with the meaning-representation issues that we arrived at in the last chapter, in considering comprehension processing. Much of what we have to say about possible mental representation here derives from the work of Johnson-Laird (1983). We return to the matter of lexical access in production, briefly touched on in chapter 5, in section 7.3, and consider the form in which word meanings may be represented in the mental lexicon. Garrett's model recognises two levels of sentence structure, a deeper as well as a more superficial one, and their organising characteristics are investigated in the next section (7.4). Finally, we consider the serial nature of Garrett's model and examine some arguments (Dell and Reich 1981; Stemberger 1985) relating to the possible interactions between levels in such a model (section 7.5) which would lead to a parallel interpretation.

7.1.2 *Assumptions and preliminaries*

The Garrett model is based on a number of considerations and assumptions, as follows.

1. *Computational decomposition reflects the grammatical decomposition of the language faculty.* This is an hypothesis according to which observations such as '*a* is adjusted to *an* after the anticipation error in

(1) I deserve *an around* of *plause* for that'

are held to tell us something fairly directly about the nature of the psycholinguistic processes involved in speech production. The statement is a 'computational' one, inasmuch as it assumes that one sort of process is dependent on the output of another – a type of 'information-flow' account.

These processes can be thought of as operating at familiar levels of linguistic description, so the hypothesis really involves a fairly direct relationship between these levels (as established by formal, reflective linguistic analysis) and stages of speech production.

As Garrett points out (1982: 21), the hypothesis might prove untenable if the exacting effects of producing (planning and executing) speech under normal time pressures force the language faculty into patterns of behaviour that are radically different from those involved in yielding contemplative linguistic analyses. As we shall see, the linguistic levels that are implicated in current models of production are fairly general ones – message formulation, lexical specification, syntactic configuration and so on – so we can say that the hypothesis is not yet tested in very great detail.

2. *The processes of speech production are independent of, and may interact with, general cognitive and motor-control factors.* This is really the same principle, on a larger scale. We assume that, just as there are, for example, (foreign) language articulatory targets that we cannot normally achieve without training, even though our motor control system is in principle capable of them, so there are certain language-specific grammatical properties that are not simply determined by the ultimate form of the cognitive plan that forms our intention to say something. That is, a foreign-language learner must learn not only speech sounds in the target language, but also specific grammatical properties, which are nevertheless within the general capability of the learner by virtue of their existing faculty (their ability to speak their native language).

This seems a reasonable assumption, notwithstanding the possibility that certain modes of thought, 'ways of viewing the world', may be characteristic of one language community rather than another. This assumption also allows for the possibility that cognitive factors may influence speech processes in ways that may be observed to hold across languages (Gazdar 1980).

3. *Speech errors may be revealing of levels of processing.* This assumption is basic to the research traditions involved in analysing naturally occurring speech errors. The sorts of levels envisaged may be summarised as: (a) the form of the utterance may fail to represent the intended message-level structure; (b) less commonly, the abstract form of the utterance may not be appropriately represented in its expression; and (c) at the level of abstract forms, we

may occasionally have the experience of 'talking ourselves into a syntactic corner' a form of what may be called a 'maze'.

Thus we have: message vs form; and abstract form vs concrete expression.

4. *Speech production is subject to real-time constraints.* At the output end of the process of production, the rate must be sufficiently rapid that our memory for what has been planned, and what has been excluded, is still available.

Memory for topic is relatively long, and even a slow output rate rarely exceeds its limitations – although most speakers have experienced the 'Where was I?' phenomenon, and not only after some extraneous interpolation.

Memory for specific grammatical form is much shorter, and we more frequently encounter a point in output where one's choice of how to continue an utterance may be in doubt because memory for the initial form has been lost. A common speaker-strategy in such cases is to restart the utterance; but if this does not happen, the latter part may be output and may fail to construe with the earlier, in violation of grammatical (but not message-level) constraints. The speaker may be unaware of this failure, or may carry on regardless.

Equally, however, output must not exceed the rate at which planning decisions are made, at any level. To the extent that this does happen, we would expect the output to be 'contentless' to some degree, at those points where the outrunning occurs; this may take the form of pausing, or of using stereotyped phrases, or of lexical forms such as proforms, lacking in specificity. Combinations of all these types would be expected to occur, for example, in fluent aphasic speech (see ch. 8); in normal speakers, the most usual response, e.g. to games where one is forced to keep speaking without undue pausing or hesitations, is to slow the rate of output so as to match that of planning, marked by longer durations given to articulation of words and syllables.

The nature of real-time constraints may be quite exacting, and their effects are found in the 'naturally displayed' evidence that we started considering in chapter 3. Normal speech errors and non-fluencies are therefore not merely evidence of how the system can go wrong (with the assumption that it will go wrong only in exceptional circumstances) – the assumption is rather that they are more directly evidence of how, and under what conditions, the system normally operates.

7.1.3 *The hesitation model versus the speech-error model*

In the following general terms, Garrett contrasts the hesitation model of Goldman-Eisler (1968) with the speech model of Fromkin (1971).

The nature of the evidence

For Goldman-Eisler the evidence comes mainly from the phenomena of hesitation, and is interpreted in terms of cognitive psychology (the planning phases of speech). For Fromkin, the primary source is speech errors, interpreted in terms of formal linguistics (the levels of structure in speech). These evidential distinctions are significant inasmuch as they direct attention to distinct aspects of the speech-production process.

Aspects of the speech-production process

Hesitation data are quite straightforwardly interpreted in terms of message-level planning. Goldman-Eisler envisaged creative planning at this level as a dynamic process involving conceptual relations (e.g. Actor–Action), lexical selection and generalised syntactic form. Syntactic-to-articulatory mapping was seen as *routine* planning, this term referring to processing that is not under direct planning-control. Thus, a traditional sort of dualistic psychological theory is involved, with inner 'cognitive' processes being distinguished from outer 'automatic' ones (see Campbell 1986 for a discussion of such a distinction in the context of language acquisition). The demands of the creative planning level reveal themselves in the way they force the operations of the routine level to wait for further instructions at certain points.

In Fromkin's model, data from the non-articulatory-based speech errors are interpreted mainly at sentence level, with certain assumptions (e.g. that semantic primitives are features rather than word-sized) being made about the nature of the meaning-level representation.

The difference in focus between the two approaches can be illustrated as in table 7.1 (from the discussion in Garrett 1982: 28–36). It is not very easy to make comparisons between two such very different sets of stages. What we are calling stage I for Goldman-Eisler could be argued to correspond to stages I–IV in Fromkin, but the fit is not complete. The important point to notice, at this step in the exposition, is that it is stages I for Goldman-Eisler, II for Fromkin that represent the major aspect of the production processes that we are concerned with here.

The place of lexical selection

In Goldman-Eisler's model, lexical selection is found both before and after syntax; presyntactically, it is involved in content-specification; and post-syntactically, under the control of abstract syntactic form, certain other lexical choices are made.

For Fromkin, lexical specification, but not selection of lexical items, is embodied in syntactic form, as bundles of syntactico-semantic features; it is not until stage IV that lexical insertion takes place.

Table 7.1 *Comparison of the Goldman-Eisler and Fromkin models of language production*

Goldman-Eisler	Fromkin
I Creative aspect (under voluntary control) (a) Content (b) Abstract syntactic form (c) Lexical specification	I Meaning representation II Syntactic structure, with semantic features at lexical sites III Intonation assigned IV Lexical insertion, by (a) meaning, then (b) form
II Routine aspect (a) Syntactic organisation details (b) Articulatory output	V Morphophonemic processes, and Articulatory output

Based on discussion in Garrett 1982.

The sources of non-fluency

Goldman-Eisler's model, as we have seen, traces hesitation phenomena in the routine aspects of speech production back to sources in the creative aspects. There are three such sources, as can be seen from our discussion above:

1. conceptual sources (content);
2. abstract syntactic form;
3. lexical choice.

In Fromkin's model, if we take the assignment of the intonation contour (stage III) as the initiation of the temporal configuration of the utterance, then only processes subsequent to this can be disruptive of it – principally, lexical insertion (stage IV).

The scope of non-fluency in the hesitation model

Because conceptual structure is recognised as a potential source of non-fluency in the Goldman-Eisler model, it is possible to accommodate the distal effects we described in chapter 3, under the heading of 'the encoding cycle'. We have noted that the work of Butterworth (1980b) has provided evidence of a correlation between fluent and non-fluent phases in spontaneous speech on the one hand, and 'idea groups', as judged by naive subjects, on the other. It is argued that non-fluent phases precede the onset of new idea-groups, and may thus be interpreted as points of long-range forward planning. Further work by Beattie (1983) has argued that gestural and gaze behaviour occurring during pauses in spontaneous speech is also linked to the encoding cycle: gestures in the fluent phase tend to be iconic, i.e. expressive of the message in the idea group currently being communicated, and emphasise or extend the linguistic means of expression; but gestures in the non-fluent phase tend to be non-specific in their content. Similarly, speaker gaze tends to be differentiated in the two types of context: it tends to be listener-orientated during the fluent phase, but averted during the non-fluent. Such evidence points to a distinction between fluent phase (FP) pausing and hesitant phase (HP) pausing. The natural interpretation of HP pauses is that they are markers of the loci of planning over very large domains – represented by whole fluent phases (perhaps one or two dozen clauses) in the encoding cycle.

It is a more difficult matter to establish exactly what planning takes place at such loci, however. It seems least likely that lexical selection for the upcoming fluent phase is involved, since this may constitute several clauses, and lexical selection is usually regarded as a more local type of process. More likely, then, the planning at HP pause-loci is for content (which does not exclude some lexical selection), but this is still a vague notion. How much abstract syntactic form is processed at such loci? Indirect evidence comes from the observation that syntactic complexity in the later part of the non-fluent phase does not appear to correlate with length of the preceding HP pausing, suggesting that syntax is planned more locally.

Beattie (1980) has found evidence in the HP pauses for a relation between pause length and length of following clause in the HP pauses, suggesting that syntactic planning is taking place fairly locally there. That no such relation emerges in the fluent phase is consistent with the view that some syntactic planning may take place before these phases are initiated, i.e. in the HP pauses.

So the picture is quite complex. In the FP pauses, there is probably a mixture of lexical selection and syntactic planning. Syntactic effects appear to emerge in the way pauses gather at clause and constituent boundaries (see ch. 3), when pause incidence is being considered. Pause duration and pause incidence may be telling us different aspects of the same complex story.

If lexical selection is not the major activity that is carried on in the sorts of pauses we have been considering, it is nevertheless an obvious determinant of much non-fluency. It may account for certain aspects of pause behaviour during both the hesitant and fluent phases, and it is most likely to be solely responsible for pauses immediately preceding major lexical items within clause and phrase constituents. Consistent with this view, Goldman-Eisler (1968), Butterworth (1980b) and Beattie (1983) have reported prelexical pausing to be related to the predictability of following words in spontaneous speech. Butterworth has suggested that, where prelexical pauses are accompanied by iconic gestures, form-based lexical retrieval processes may be distinguished from meaning-based selection (see our discussion of the tip-of-the-tongue phenomenon in ch. 3).

The mechanism of nonfluency in the hesitation model

To conclude, Garrett (1982) suggests that we may relate the three sorts of pause determinants just outlined to their respective mechanisms.

1. *Overload*: distal pauses, linked to the encoding cycle, arise from the demands of conceptual planning which momentarily overload the system, yielding pauses that are associated with non-specific gestures, averted gaze, a break between idea groups and the onset of a fluent phase.

2. *Establishing frames*: a frame, as distinct from a plan, is specified for certain syntactic properties, and generally (we assume) corresponds to a clause. Establishing a frame yields pauses that, in terms of their incidence, tend to gather at major constituent boundaries (frame-joints), particularly those that are clause-initial (between frames).

3. *Filling the frames*: if a plan is a sequence of frames, specified locally on a frame-by-frame basis, then a frame is a sequence of elements of which some, particularly the major lexical items, are also specified locally. Filling a frame, word by word, yields pauses that tend to gather at prelexical points.

Finally, let us recall, first, that 'pause' in this discussion is a quite general term covering a range of distinguishable phenomena, filled and unfilled hesitations, and also phrase- and word-final segment lengthening; secondly, that all three mechanisms may combine, in complex determination of individual pauses; and, thirdly, that not all such planning mechanisms may occur actually during such pauses, since stereotypic stretches of speech, including such 'filler' words and phrases as *actually*, *well*, *I suppose*, *you know*, etc. may afford the speaker planning opportunities that would escape a traditional type of pause analysis.

7.1.4 *The speech-error model*

Levels of representation

We need, then, to recognise a number of potential levels of processing, which Garrett (1982) refers to as *the message level*, *the sentence level* and *the articulatory level*.

The message level

This is described as a real-time construct, compositionally built up from simple concepts that are linked to the speaker's perceptual and affective states and encyclopaedic knowledge, according to some 'conceptual syntax'. It is related to the pure-linguistic level of an utterance; but, unlike semantics in formal grammars, it embodies non-linguistic encyclopaedic knowledge as well. 'Pragmatic–semantic' might be a closer description.

Garrett also suggests that the word meanings, as found in the lexicon of a language, might form the basic vocabulary also of message-level structures and processes. Thus, when the message level is described as (de)compositional, this means that individual messages can be thought of as being built up out of, or decomposable into, meaning components that essentially map onto the word meanings in the lexicon. These word-sized meaning components are themselves, in certain semantic descriptions, seen as decomposable into *semantic features*; but Garrett is at pains to suggest that these features seem to play no role in speech production. One reason for this is that there is little or no evidence in speech production data for the view that 'bachelor', 'die', 'empty', etc. involve a negative feature, as in 'not + married', 'cause + not + live', 'not + full', etc. Clearly, the assumption that the basic building blocks of the message are 'word-sized' rather than atomistic elements such as semantic features prepares for the stage in the model where message-level formulations are mapped onto word meanings in the lexicon.

The sentence level

This is described as the real-time construction of representations wherein lexical items (selected through their meanings) and abstract utterance-markers such as tense, number, mood, etc. are grouped into phrasal constructions. These elements all have phonological forms, so that this level establishes a mapping between message structure and articulation.

The articulatory level

This is the level which is addressed in the model we have already discussed in chapter 4, working from sensory goals for articulatory targets down to monitored innervation and control of specific muscle groups in the articulatory system. We shall not consider it further in this chapter.

7.2 The internal structure of the message level

Our discussion of the Garrett model has introduced certain levels of representation. Of these, Garrett's own discussion focusses particularly on the sentence level, but we shall start with a closer consideration of the message level, and its relationships with the internal structure of the lexicon. For this purpose, we shall not pursue further the line of evidence that derives from hesitation phenomena, but turn to consider a different body of evidence, from certain sorts of conceptual operations. In this connection, we shall refer to the work of Johnson-Laird (1983) which, while it owes no allegiance to the Garrett model, serves as a useful framework within which some of the important issues may be raised.

7.2.1 *Models of meaning*

We have earlier (ch. 3) made reference to the relatively inaccessible nature of meaning; but having pushed the problem aside in this way for the intervening discussion, we now have to face it more squarely. We have, according to Garrett (1982), to deal with a message level having the following properties:

1. it is a real-time conceptual construct;
2. it is compositionally built up from simple concepts by some sort of 'conceptual syntax';
3. it uses pragmatic (i.e. real-world knowledge) as well as semantic input;
4. the primitive elements that constitute its basic vocabulary are word-sized units rather than semantic features.

While these points help to set the scene, they do not take us very far towards an understanding of the structures and processes that might underlie the linguistic forms of representation in the Garrett model.

The issue that confronts us concerns the nature of thought processes, and how they are related to the sorts of representations that constitute the linguistic bases of language production. Johnson-Laird (1983) provides a wide-ranging review of the problems and approaches to this area in a form which allows contact to be made, in general terms, with the sort of model we have been discussing. In particular, Johnson-Laird has been interested in *reasoning*, the sorts of thought processes that are said to be 'logical', moving from evidence to conclusion. Of all the various types of thought processes, these might appear, by their apparent systematicity, to be among the most amenable to description and analysis. The steps in reasoning involve *inferences*, and these may be studied most easily in fairly formal conditions, as in *syllogisms*, such as:

(2) (a) Some of the artists are beekeepers
(b) All of the beekeepers are chemists

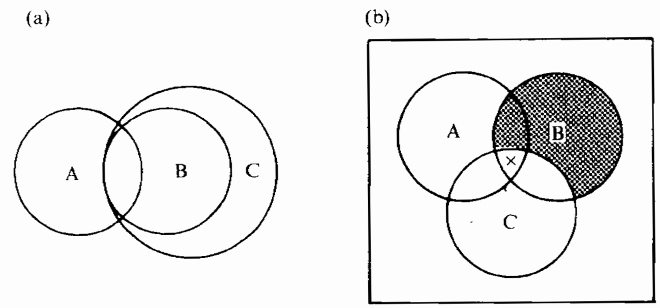

Figure 7.1 (a) Euler circle and (b) Venn diagram representation of 'Some A are B, all B are C'. The Euler circle representation is not the only one possible. The shaded area of the Venn diagram indicates subsets that are ruled out by the premises, and x marks the area that conforms to the conclusion '∴ some A are C'. (Based on discussion in Johnson-Laird 1983: pp. 77–93.)

(ci) Therefore some of the artists are chemists
(and (cii) Some of the chemists are artists)

but they may also be found in more everyday situations, such as:

(3) A: Do you know where the squash courts are, please?
B: I think that's the sports hall over there.
A: (goes to indicated building)

7.2.2 *Mental models versus mental logics*

A traditional view of the sort of reasoning just illustrated – at least in its more successful moments – is that it proceeds according to universal laws that hold just as well for that part of the universe that lies inside our heads as for everything outside. The laws governing thought are essentially those that govern the operation of physical machinery, whether (as in a typical nineteenth-century analogy) steam engines or (more recently, and more compellingly, perhaps) computers. Particular systems of such laws have been proposed, e.g. 'reasoning is nothing more than propositional calculus itself' (Inhelder and Piaget 1958: 305, quoted in Johnson-Laird 1983). For syllogisms, it has been suggested that mathematical notations such as Euler circles or Venn diagrams provide an insight into the nature of the reasoning processes involved, as illustrated in figure 7.1. Johnson-Laird notes that this 'mental-logic' approach has a number of problems:

1. it has difficulty with quantified expressions, such as '*every* X has *some* Y';

2. a very high number of formal representations typically have to be computed and evaluated to arrive at the set of possible outcomes – sufficiently large to raise suspicions about their psychological plausibility;
3. it cannot account naturally for certain aspects of human performance, e.g.
   (a) the extent and varieties of normal errors made on particular syllogisms;
   (b) certain 'figural effects', as in the clear preference for (2ci) above over (2cii) (though each is equally valid; see below for further discussion);
   (c) the effect of content, by which it is easier to arrive at conclusions when the elements in a syllogism are represented by familiar, imageable words rather than by the abstract symbols A, B, C.

Johnson-Laird refers to the work of Craik (1943) for an alternative concept of a 'mental model'. This has the following general properties:

1. it is a small-scale, internal, model of external reality;
2. it has a *direct* relationship with its external counterpart;
3. it is a *functional* rather than a physical entity;
4. it is not necessarily complete or accurate;
5. it has *recursive function*.

Some of these points require some comment.

Direct representation of reality: What is emphasised here is that, for example, an engineer's mental representation of a particular bridge, or a sculptor's idea of a statue, is couched in terms of the materials and proportions of the actual object. Model stresses will run through the model bridge (guided by the engineer's knowledge) like real stresses through the real bridge; and the spatial relations between different parts of the idea of the statue will appear in the final result.

Functional representations: Nevertheless, the model is essentially a relational structure rather than a physical entity. We must assume that the mental model, as a complex neurophysiological entity, or state of affairs, exists in the brain, but it is not this brain construct, in terms of the cells and pathways involved, that models the real world in a direct way. The mental model is a construct of the mind, rather than the brain. So, just as in part I, chapter 2 we were able to discuss the architecture of the brain from the point of view of language, without making contact with language itself, so here we assume that thought processes are not simply reducible to brain states. In this sense, the

mental model is abstract: Johnson-Laird likens this abstractness to that involved when we say that a particular computer program (in the abstract sense) is available in different versions (the physical sense), to run on different pieces of hardware.

Completeness and accuracy: Human beings are constantly manipulating objects that they barely understand, both natural (e.g. garden plants) and artificial (e.g. cars and television sets). Hence, we cannot require people's understandings to be utterly faithful representations of these objects. Also, there must be individual differences in degree of understanding: the television set is a box with knobs on to most people, but much more than this to the repairman.

Recursive function: We can illustrate this properly by reference to one of the earliest modern attempts to understand the conceptual structuring that underlies complex behaviour, presented in Miller, Galanter and Pribram (1960). They envisaged such behaviour as governed by *plans*, which are built up from basic units that can be likened to simple flow-charts of information control – Test–Operate–Test–Exit (or TOTE) units. (See fig. 7.2.) Aspects of complex performance, such as hammering a nail into a block of wood, can be viewed as sequences that are guided by *plans*: it is such guidance that distinguishes complex activities from random acts. Further, it is not just the succession of individual acts of hammer striking that is complex: each hammer strike has its own components, viz. of raising, then lowering, the hammer. The complex plan thus allows for recursion in two senses: first, operationally, it allows for recirculation of control around the TOTE structure any number of times prior to Exit; secondly, in its structural configuration, it allows for TOTE units to be embedded inside larger TOTE units, thus enumerating a hierarchy of control processes that is dominated by a single plan. The essence of such an arrangement is captured in the following:

> More complicated Plans . . . can similarly be described as TOTE units built up of subplans that are themselves TOTE units. A bird will take off, make a few wing strokes, glide, brake with its wings, thrust its feet forward and land on the limb. The whole action is initiated as a unit, is controlled by a single Plan, yet is composed of several phases, each involving its own Plan, which in turn may be composed of subplans, etc. (Miller, Galanter and Pribram 1960: 37).

7.2.3 *Propositions, images and models*

There are basically three ways in which abstract problems or properties can be represented in terms that are compatible with what is

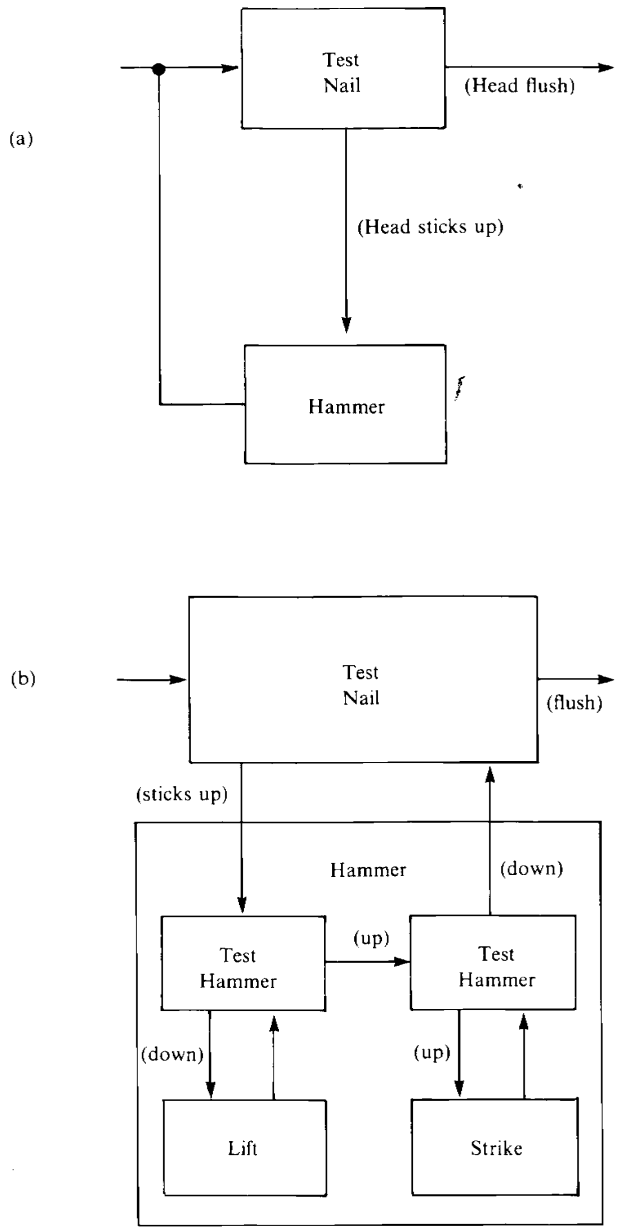

Figure 7.2 TOTE-units (a) simple, (b) hierarchical. (From Miller, Galanter and Pribram 1960: figs 3, 4, 5, pp. 34–6.)

observed of human performance (as opposed to abstract systems of logic): mental *propositions*, *images* or *models*. What we have said thus far does not address the issue of how far they are to be distinguished from each other, so we shall briefly compare them here.

Reasoning with propositions

Propositions have the virtue, from our point of view, of being compatible with and close to the utterance of language. They are, in Garrett's terms, the ultimate product of the message level in language production, available for fairly straightforward mapping into the first linguistic level, the functional level of representation. Propositional vocabulary is close to the surface vocabulary of language; propositional syntax is close to the functional-role frames of the functional level. If a speaker wants to describe a sequence of related events, then the sequencing, and the relation between the events, must be part of the message-level structure (Garrett 1982) or plan (Miller, Galanter and Pribram 1960), or propositional representation (Johnson-Laird 1983) that underlies the utterance.

Reasoning is involved in the exercise of the speaker's judgement in how much detail to put in, or leave out. For instance, in the utterance sequence (adapted from Johnson-Laird 1983: 52):

(4) (a) The victim was stabbed to death through the throat
(b) The suspect was on a train to Edinburgh at the time

the speaker has judged appropriately if we are intended to infer that the suspect has an excellent alibi. But we shall feel deprived of information if it is revealed subsequently that the murder in question took place on that train. Johnson-Laird reports that listeners who are asked whether the suspect's alibi is really a good one will try to reinterpret the utterance sequence in an effort to examine certain (possibly erroneous) assumptions that may have been involved in their understanding (see our discussion at the beginning of ch. 6), even to the extent of speculating about extremely long knives. In this process, Johnson-Laird suggests that what they have been doing is going back and forth between the original utterance and their *propositional representation* of it; the latter is sufficiently close to the former to permit accurate verbatim recall of the utterance sequence, at least over short periods of time. Also, they have been attempting to eliminate those possibilities that are ruled out by the premises, until only one (apparently) is left.

This suggests that propositional representation enables reasoning to proceed without recourse to mental logic and formal rules of inference; but it does

Table 7.2 *A 'tableau' representation of the premises 'Some A are B', 'All B are C', with the conclusion 'Therefore some A are C'*

	artist				
'some,	artist =		beekeeper =	chemist	'therefore some
but not	artist =		beekeeper =	chemist	of the artists
all of			beekeeper =	chemist	are chemists, and
the artists		'all the	beekeeper =	chemist	some of the
are bee-		beekeepers		chemist	chemists are
keepers'		are chemists'			artists'

Based on Johnson-Laird 1983: 95.

not establish the relation between propositions and mental models, to which we now return.

Mental models

To go back to the artist/beekeeper/chemist syllogism in example (2): a way to reason without Euler circles or Venn diagrams is to set up what Johnson-Laird (1983: 137) calls a 'tableau', as illustrated in table 7.2. If we set the tableau out like this, using pen and paper, we are going straight from propositional representation to a pen-and-paper model in the real world; if we eliminate the pen-and-paper version, the tableau is set up in the mind – it is a mental model.

Such a model could account for the figural effects that derive from the order of its construction: artists come before beekeepers in the tableau, so the response bias towards (2ci) *some of the artists are chemists* (rather than the equally valid (2cii) *some of the chemists are artists*) derives from properties of the mental model (the term 'figural' refers to the *configuration* of the tableau or model used); these in turn can be argued to derive from properties of the propositions constructed from the utterance sequence embodying the syllogism.

But what of naturally occurring discourse? This can also be seen as involving speaker-judgements/listener-inferences, but they are not so explicit as those found in syllogisms, of course. Consider (again adapted from Johnson-Laird 1983: 128):

(5) The victim was stabbed to death through the throat. The pilot put the plane into a stall just before landing on the strip. He just got it out of it in time, and ran off in the confusion as soon as it had come to a halt. Wasn't he lucky?

In this constructed example, the speaker, we shall say, is describing a complex sequence of events, forming a complex whole. This is represented as a single mental model of the whole incident. It is the speaker's judgement that the fact that the victim was a passenger on the plane, the fact that the plane had a pilot, etc, need not receive explicit propositional representation. Such judgements depend on the speaker's awareness of the listener's knowledge. Equally, the intended meanings of the ambiguous items *pilot*, *plane*, *stall*, *strip*, are judged not to need explicit marking. The specific referential links, from the first *he* (*the pilot*), the first *it* (*the plane*), the second *it* (*a stall*), the ellipted subject of *ran off* (*the pilot*), the third *it* (the plane) and the second *he* (*the pilot*) are all likewise not explicitly marked. Finally, the force of the question *wasn't he lucky?*, which is 'rhetorical', more like a statement, is not explicitly marked.

What this suggests is that *certain* aspects, over others, of the mental model are being extracted and embodied into propositional form by the speaker; the propositions carry certain explicit markers of gaps, where the listener must build inferences – e.g. definite noun phrases (*the pilot*), pronouns (*he*), ellipsis (*ran off*) – but the nature of the links across these gaps is left implicit rather then explicit.

There is thus, underlying a linguistic utterance, a mental tableau within which everything is explicitly related to everything else (up to the limits of the speaker's knowledge), but not all these relationships appear in the propositional representation. Some relationships do not appear at all; others do appear, but implicitly. The nature of a propositional representation of a given mental model will vary, depending on the speaker's judgement of the listener's knowledge, the context and the earlier discourse. A mental model, for Johnson-Laird, appears to stand in the same sort of relationship to propositional representations as does Miller, Galanter and Pribram's (1960) master-Plan to the subplans that it dominates. But whereas subplans might simply be regarded as more locally determined specifications of aspects that are only covertly embodied in the master-Plan, the mental model is the construct that contains all the available detail, and the speaker's task is to judge how and what to select from it for building a propositional representation.

Images

It might appear from what we have said thus far that Johnson-Laird is using the term *mental model* essentially to refer to what are more generally called *images*, about which there has been a great deal of psychological research. It is well established that images are conscious phenomena, that they aid memory, can encode spatial information and be rotated in 'the mind's eye' at a certain speed, etc. But what the nature of images might be is less clear. They cannot be 'pictures in the mind' (as the expression 'in the mind's eye' suggests), because this would require an homunculus (possessing the 'eye') in the mind, whose mind receives the pictures by means of having an homunculus . . . and so on, *ad infinitum*.

Johnson-Laird distinguishes two views on the nature of images: one, attributed to Paivio (1971) and Kosslyn (1980), is *realistic*, inasmuch as it assumes that mental images share certain perceptual processes that are involved in the visual perception of external objects, that mental-image rotations are gradual, and analogous to the visual perception of rotating objects, and that images are representational of (temporarily absent) objects.

The other view, associated with Pylyshyn (1973, 1981), also regards images and the results of visual perception of objects as similar, but sees them each as being mediated by propositional representations, such as we have discussed above. According to this, *propositional*, hypothesis, there is a many-to-one propositional-to-element relationship in visual arrays, rotations are incremental rather than gradual and images are *epiphenomenal*, i.e. they may be constructed from propositional representations of (absent) objects, but they do not have to be, and, unlike propositional representations, they are not our primary means of reflecting upon those objects.

For Johnson-Laird, mental models can be seen as underlying not only propositions (as discussed above), but also images. Propositions are distinct from mental models, as we have seen, in having an arbitrary conceptual syntax, i.e. they are linearised, in real time. By contrast, mental models are multi-dimensional, non-linearised constructs; and images may be thought of as particular, two- or three-dimensional 'views' of mental models. Whereas we may have a mental model of *planes* (aircraft) or *pilots* (of aircraft) in general, we may also have *images* of particular examples of these.

Finally, we may put mental models, images and propositions in relationship to each other by adapting an observation of Pylyshyn (1973): we may rotate a mental model of, say, a room, by adopting a particular *imaginary* standpoint from which it is to be viewed; this establishes the spatial representations in the model in a fixed way – the bed is now to the left of the cupboard (from the opposite viewpoint, it would be to the right). Such a specification is not a fixed property of the mental model, but it is consistent with the multi-dimensional nature of that model. A propositional representation of this image may leave this spatial relationship between the bed and the cupboard relatively unspecified – using a spatial term such as *next to* or *beside*. Thus, in principle, a propositional representation may encode aspects of the mental model as formed in a particular mental image of it, or directly from the mental model itself, or

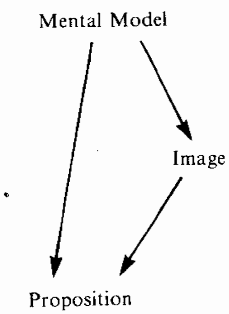

Figure 7.3 Possible encoding relationships between mental model, image, and propositional representation of meaning. (Based on discussion in Johnson-Laird 1983: pp. 146–66.)

both: see figure 7.3. We may say, in these terms, that where the contribution of the mental image is minimal, the resulting description may be radically indeterminate, as in:

> I have a very small bedroom with a window overlooking the heath. There is a single bed against the wall and opposite it a gas fire . . . The room is so small that I sit on the bed to cook. The only other furniture in the room is a bookcase on one side of the gas fire next to the window . . . and a wardrobe. It stands against the wall just near to the door, which opens almost directly onto the head of my bed. (Johnson-Laird 1983: 162)

A passage like this, low in image-based elements, is construable in many different mental representations; the listener will have no difficulty in constructing a propositional representation, but may find it difficult to go beyond this. Different listeners would vary a good deal in their attempts to draw the room from their (perfectly competent) understanding of the passage, and memory for such a description will be less good than for one which permits the construction of a single mental model. In terms of Grice's (1975) conventions, the speaker (or writer) has saddled the listener with an unfair burden of understanding.

Thus far, we have been considering the 'syntax' component of our general processing-model in figure 4.1. Propositional representations of meaning are essentially syntactic in nature, since they are linearised structures involving classes of elements (i.e. the syntax of semantics). We shall now turn to the nature of the elements concerned, which comprise those aspects of meaning for which the language in question provides *lexical* representation.

7.3 Lexical access: the nature of stored word-meanings

In a propositional representation of meaning in message structure, such as might be derivable from a mental model, there are 'elements' as well as 'relations between elements'. The existence of elements raises the issue of access to those meaning complexes that are subparts of the message structure and stored as stable representations of word meanings in the mental lexicon.

7.3.1 *Semantic features or meaning components*

On this matter, we have observed that Garrett (1982) views with scepticism the possible role of atomistic semantic features or meaning components – and Johnson-Laird (1983: 207) is equally dismissive of them. Before we move on, however, a word about the terms used for these atomistic primitives may be in order. 'Semantic features' and 'meaning components' are two terms that are often used together; we should bear in mind, however, that our primary concern is with meaning, and linguistic semantics is best thought of as a particular approach to the study of meaning. Thus 'semantic features' strictly means features that have been proposed within semantics, whereas 'meaning components' may be thought of as those entities which semantic features are set up to elucidate.

We shall not go into the arguments against such primitives here, except to note that Garrett's scepticism is supported by Johnson-Laird's (1983) report of a failure to find any effect of semantic-feature complexity in pairs of verbs such as *get/take*, *move/throw*, etc. We may also observe that semantic features have been criticised on theoretical–descriptive grounds within linguistics – e.g. Bolinger (1965) and Palmer (1981). While it is true that most psycholinguistic studies casting doubt on the validity of meaning components in lexical entries derive from comprehension tasks, rather than production, it is unlikely that representations of meanings are basically different between these two modes of performance.

An alternative way of capturing relations between words in semantics is the use of *meaning postulates* (Kempson 1977; Biggs 1982; Johnson-Laird 1983); as such, we should consider how far they may provide a suitable alternative for the representation of meaning in the mental lexicon. Meaning postulates may be thought of, in a preliminary fashion, as similar to the *redundancy rules* that are associated with semantic analysis in terms of features. Such rules specify that, for example, the semantic features of 'husband' may be reduced in the way illustrated in figure 7.4 (after the discussion in Kempson 1977: 88–92, 188–91). This achieves economy of representation in two ways: first, by putting certain semantic features into relationship with each other in a way that is stated once for all in the lexicon; and secondly, by allowing for a con-

'husband'

(a) Feature specification

[Human]
[Animate]
[Adult]
[Concrete]
[Married]
[Male]

(b) Redundancy rules + Feature specification

Redundancy rules (of general application)
[Human] x → [Animate] x
[Adult] x → [Animate] x
[Animate] x → [Concrete] x
[Married] x → [Adult] x
[Married] x → [Human] x

Resulting feature specification
[Married]
[Male]

Figure 7.4 Two feature specifications of 'husband', (a) without, (b) with redundancy rules. (Based on discussion in Kempson 1977: pp. 88–92.)

sequent reduction in the complexity of the semantic representation of individual lexical items.

The use of meaning postulates essentially extends this approach, to the point where *all* elements of semantic representation are stated in terms such as:

(A) x → (B) x

In principle, A and B may be semantic features [A], [B], just as in redundancy rules, but an alternative (in view of the difficulty with these primitives) is to have A and B representing the *lexical items* themselves of the language concerned. That is, the lexical items are treated as semantically primitive, and are set into relationship with each other by these rules. It is this type of meaning-postulate approach that we shall consider here.

Accordingly, the meaning of 'husband' may be specified as:

'husband' x → 'married' x
'married' x → 'adult' x
'adult' x →

At this point, our formalism proves unhelpful, for the word 'husband' must be put into relationship with 'male', but the word 'adult' is in relationship with both 'male' and 'female'. So we can state meaning postulates more conveniently as:

'husband' x → 'married' x
AND 'adult' x
AND 'male' x, etc.

As Kempson (1977) points out, meaning postulates provide for a weaker representation of word meaning than semantic features; they do not try to capture all aspects of a word's meaning, such as what defines a 'waitress' vs a 'woman', but simply seek to represent the full set of lexical relationships between these words and others in the lexicon. To the extent that these words have similar relationships with other words, they will be shown to have similar meanings; postulates do not go beyond this.

A further consequence of this approach is that the concept of a semantic entry is dispensed with entirely. This may appear to be a startling result, but it follows quite straightforwardly from what we have said. If we wish to point to the formalised semantic specification of a particular word, we can no longer point to some internal set of features; the semantic specification lies word-externally, in the relationships that the word in question contracts with the other words around it in the lexicon. Words, that is, derive their semantic value from the organisation of their semantic space.

7.3.2 *Lexical access via meaning representations*

If we now bring the discussion back to psycholinguistic issues, we can consider the implications of relying, as Garrett's model does, on message-level elements that are basically lexical, or 'word-sized'. Their corresponding entries in the mental lexicon are unanalysed, stored elements of lexical meaning; their meaningfulness derives from their being set in many-to-many relationships with each other. As such, the individual items in this lexicon have nothing corresponding to an internal meaning-specification. It is difficult to see, then, what information could possibly be contained about such items in the 'semantic-access' file of Forster's indirect model of lexical access (5.3.2).

It is as if we were to walk into a library with no catalogue system, but with a set of guidelines to the effect that books on topic A are shelved alongside books on topic B, and so on. In other words, we are led directly to the main stacks, and the preliminary stage of consulting the catalogue is eliminated. If this seems a rather unlikely, and inefficient, way of organising access to books in a library, we should not conclude that it is equally unsuitable as a model of human lexical access. After all, the analogy between a library and a mental lexicon is not perfect. If a message-level element having the properties of being human and female 'calls up' or 'activates' all such entries in the lexicon – 'woman', 'waitress', 'actress', etc. – this may naturally allow for such activated entries to interact with further properties of the message level, such as that the mental model involves a kitchen or restaurant event, in which case 'actress' is not further activated, but 'waitress' is.

In this way, the model is not only direct; it is also interactive and parallel in

processing mode, since it allows for different elements and properties of the message level to enter into and guide, or constrain, the search process, along a number of different pathways through the lexicon, and activating a number of lexical items as it proceeds.

7.3.3 *Semantic networks*

The sort of approach we have been considering here is frequently represented as a distinct type of theory, *semantic-network theory*; but Johnson-Laird (1983) is right to suggest that it is better thought of as a notation or formalism. As such, it derives from work (Collins and Quillian 1972) in artificial intelligence, where the meaning of a word is set up in terms of a network of 'is a' relationships, as in:

'husband' → is a → 'man' → is a → 'human' → is a → 'animal'...
→ is a → 'adult'...

Network formalisms have been proposed, on the boundaries of psychology, linguistics and artificial intelligence (e.g. Hinton 1981), which attempt to model human performance in terms of *network hierarchies* of lexical items, *network distances* between lexical items and *network transitions* associated with faster or slower links between points in the network. Modelling human performance in these terms is complicated (Collins and Loftus 1975; Anderson 1984); individual properties, such as distance, seem not to be reliable determinants of lexical access (as Johnson-Laird notes). Interactions between various network properties seem generally plausible, but are difficult to quantify and express as specific, testable hypotheses.

Johnson-Laird's own proposal (1983: 217) introduces another possible determinant of lexical access; he distinguishes between words that are basic, not defined in terms of any other words of the network, e.g. 'be', 'go', 'move', 'at', etc, and words that can be defined in terms of others, particularly the basic set, e.g. 'cost', 'weigh', 'fly', 'emigrate', 'behind', etc.

We shall not go into these issues further here. We shall return to the concept of a lexical network below, in considering some alternatives to the standard interpretation of Garrett's model.

7.4 The internal structure of the sentence level

Thus far, we have been concerned with the message level and the ways in which it might be mapped onto lexical and syntactic aspects of the language system. This brings us to consider the nature of the sentence level in more detail, and Garrett argues that it is possible to delineate its internal structure by looking closely at movement errors (exchanges and shifts) and replacement errors (substitutions and blends), of the sort that we reviewed in chapter 3.

7.4.1 *The sentence sublevels*

Garrett proposes the terms *functional* and *positional levels* for sublevels within the sentence level. They are distinguished by the nature of the elements involved, their qualitative inter-relationships and the domains over which they operate.

The functional level

Movement errors at this level typically involve word elements, of corresponding grammatical categories (noun and noun, verb and verb, etc.), playing similar roles in phrases (modifier and modifier, head and head, etc.), and the movement domain traverses a phrase-constituent boundary. We may note that this has the secondary effect of yielding rather extensive domains of movement, compared to those of the positional level:

(6) As you *reap*, Roger, so shall you *sow*.

The positional level

This involves subword elements, syllables, segment clusters, singleton segments and, occasionally, subsegmental features (see part I, chapter 3, section 3.4.1) with no grammatical-category or phrasal-role correspondence, but frequently with similar phonological roles (onset and onset, coda and coda, etc.) frequently within the same constituent phrase (which may better be characterised in terms of phonological criteria – a phonological phrase), and hence over rather shorter distances than word movements at the functional level:

(7) *sh*ow *sn*ovelling

In an earlier study, Garrett (1980a) provided a tabular demonstration of some of these distinctions for exchange errors, as in table 7.3. From this it will be seen that word exchanges are set apart from both stranding and sound exchanges, with most (not all) stranding exchanges belonging to what is here referred to as the positional level. There is claimed to be no effect of subword segmental structure in functional-level word exchanges, and a strong segmental effect in positional sound exchanges.

The picture thus far may be set out as in figure 7.5.

7.4.2 *The internal structure of the syntactic component*

It remains for us to ask how the mapping from the functional level to the positional level might take place, via some syntactic component that the

Table 7.3 *Constraints on exchange movement errors – phrasal membership and grammatical category*

Exchange error	Phrasal membership		Grammatical category[b]	
	Within[a]	Between	Same	Different
Word (N = 200)	0.19	0.81	0.85	0.15
Stranding (N = 100)	0.70	0.30	0.43	0.57
Sound (N = 200)	0.87	0.13	0.39	0.61

[a] Internal to a simple NP or VP, where VP is taken to include main verb plus obligatory phrasal constituents such as direct object NP.
[b] All types of exchanges, word, stranding and sound, are almost entirely confined to the major lexical categories Noun, Verb, Adjective, Preposition.
From Garrett 1980a: tables I, II, p. 189.

Garrett model does not specify. Lapointe (1985) has addressed this issue directly, in the context of a study of language pathology, specifically, agrammatism in Broca's aphasia (see ch. 8, sections 8.2.3 and 8.4.2 for further information on these concepts). His conclusions are not restricted to the field of language pathology, however, and it will serve our purpose here very well to consider them.

Lapointe is concerned just with verb phrases in language production, and therefore his discussion provides an illustration from this area of how the syntactic component might work. He calls his model a *syntactic processor*, and discusses it in terms of (a) what sort of input it receives from the functional level; (b) the means by which positional-level frames are stored in the system; and (c) the nature of the output from the processor.

The input to the syntactic processor

Concerning the input from the functional level, much currently remains unclear. It may be thought of as specifying underlying grammatical relations, or as basically semantic in nature, or both (depending on one's view of how distinct these concepts might be). Lapointe assumes, for the sake of convenience, that input to his syntactic processor will include information such as, e.g.,

(. . . indicative, active, durative, present, sing-3 . . .)

where the dots indicate information from the functional level about elements preceding and following the verb phrase in the utterance.

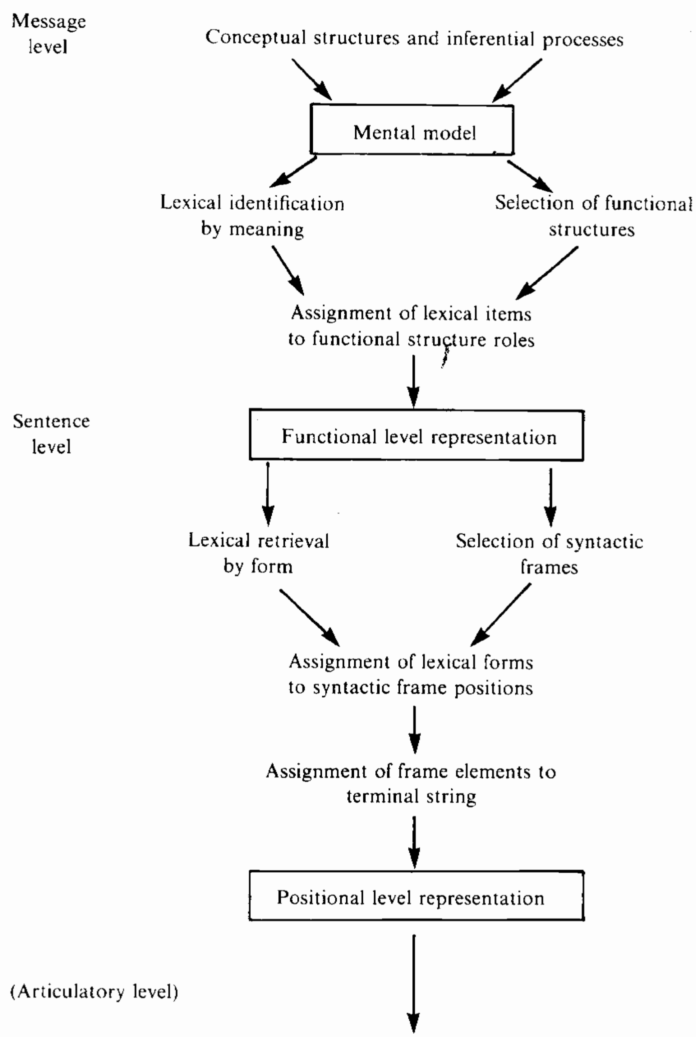

Figure 7.5 Schematic layout of the message and sentence levels in the Garrett model. (Based on Garrett 1982: pp. 67–8.)

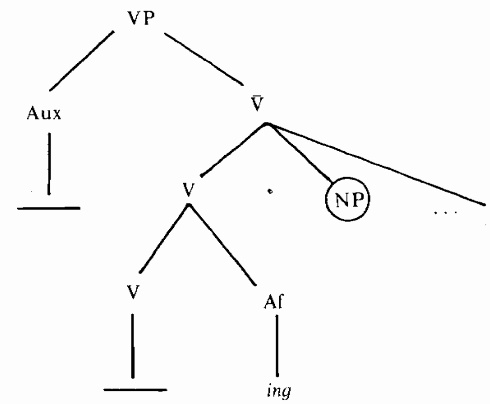

Figure 7.6 A verb-based positional frame fragment. (From Lapointe 1985: example (30), p. 130.)

The positional frames

Lapointe considers the positional frames to have the sort of structure represented in figure 7.6. He refers to this type of structure as a *fragment*: it represents the maximal phrase structure associated with a particular head category (here, it is V), both higher than the V-node (up to a VP node) and lower, down to the *stem* + *affix* morpheme structure; and it also includes slots, defined as empty spaces under certain nodes, where the head lexical element (a suitably specified verb from the lexicon), and dependent function-words (such as auxiliary verbs) may be inserted, and circled symbols (e.g. the NP in this example), representing positions where other constituents (having their own internal structure) may be attached.

The fragment and function-word stores

As far as storage is concerned, Lapointe envisages two distinct types of stores: one, for lexical head categories (an N-store, as well as a V-store, and others), and another type for 'dependent function' (or grammatical) word elements (e.g. determiners, auxiliary verbs, and so on). This distinction is well attested in both descriptive linguistics and in performance data from both normal and abnormal language use, and, in spite of difficulties associated with it (see 3.3.2), it is natural to embody it, within this type of model, in terms of different types of store. Thus, the basic function of the processor is to access fragments from these stores and to combine them in grammatically appropriate ways.

Table 7.4 *A partial V-fragment store for English*

English V fragment store			
V	Aux V + ing	Aux V + ed″	Aux being V + ed″ . . .
V + s	Aux been V + ing	Aux been V + ed″	Aux been being V – ed″
V + ed			
Aux V + ed′			

From Lapointe 1985: table 6, p. 132.

The internal organisation of these fragment stores is obviously an important issue. Lapointe argues, on the basis of a dimension of morphosemantic complexity (the details of which need not detain us here) that a partial V-fragment store for English might look like the arrangement of table 7.4. The least complex phrase-structure type is located in the leftmost column, with increasingly complex structures arranged in columns to the right of this; and the least complex forms within each structural type are located at the top of the columns.

The organisation of the auxiliary-fragment store is more unsettled, but Lapointe suggests an arrangement with *be* forms in the leftmost column (least complex), then a column with *have* forms, then *do* forms and finally (most complex) the modal auxiliaries.

The operation of the syntactic processor

Turning now to the operations performed by the syntactic processor, we may first refer to the diagram in figure 7.7. This shows the syntactic processor to have three subcomponents: a *control* mechanism, a store *locator* and a stem *inserter*, as well as three types of store: the two fragment stores mentioned above, for fragments and function words, as well as an *address index* containing the addresses of cells in the fragment and function-word stores, where specific information may be found.

The control mechanism receives input from the functional level, activates the address index to find the location of the required cells in the fragment and function-word stores and passes this information to the locator.

The locator may be thought of as a set of *read/copy devices* (see our discussion of Shattuck-Hufnagel's (1983) scan-copier model at the level of sound structure, in section 4.4.1), one for each fragment store (N-store, V-store, etc.), and each function-word store (determiner store, auxiliary-verb store,

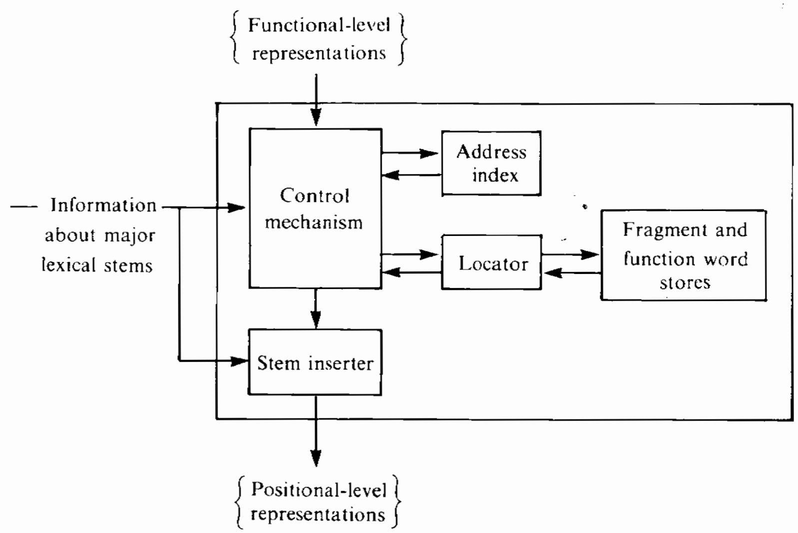

Figure 7.7 Schematic diagram of the syntactic processor. (From Lapointe 1985: 133.)

etc.). Each read/copy device is constrained in the way it can move through its store, along tracks defined by the rows and columns; in this way, it is possible to operationalise the concept of 'complex form' by making it more distant from – hence, less easily accessed by – the read/copy device, which is assumed to be 'at rest' in the top left-hand corner of each store. It is also assumed that the read/copy device can only access one cell's information at a time, between returns to its resting place.

Once the locator has activated the relevant set of read/copy devices and has received input from their operation, its task is completed by handing this information back to the control mechanism, in the form of fragments and function words. The control mechanism's task then is to combine these in grammatically appropriate sequence, and to pass the result on to the stem inserter.

The stem inserter, finally, inserts the phonological representations of the lexical items that have been activated in the lexicon (a parallel process to the one we have just been following through the syntactic processor) into the appropriate positions defined in the output of the control mechanism. It may be thought of as a copying mechanism, and here we make contact with our discussion in chapter 4 (Shattuck-Hufnagel's scan copier is a type of mechanism that performs this function).

7.4.3 ***Movement errors***

Exchange errors

These are assumed to arise as representations at a higher level are mapped onto representations at the next lower level (refer back to fig. 7.5). They are a consequence of a failing to preserve ordering relations between elements across the interlevel boundaries. Thus, word exchanges arise from mismappings of message-level elements onto functional-level elements, and sound exchanges from mismappings of functional-level elements onto positional-level elements. Stranding errors, or most of them, belong to this latter area, too.

Since exchanges of all types almost exclusively involve major lexical categories (see note (b) in table 7.3, above) we may draw a distinction, valid for each of these two interlevel mappings, between frame elements (at message, functional or positional levels), which do not exchange, and lexical elements, which do.

As we have seen (in 3.4), Shattuck-Hufnagel (1983) envisages a segmental and stress-specified skeletal frame for each lexical item, in order to account for the distinction between those sound elements that move and those that do not: movable elements are articulatory specifications, more usually as complexes than single features, while stable elements are syllable patterns, and stress patterns, constituting the lexical frame.

We may assume a similar distinction in respect of the assignment of major lexical items to phrasal-planning frames, at the functional level. Between these types of exchange, we may consider stranding errors to arise early on in the mapping from the functional level to the positional level, where word forms as minimal grammatical elements (e.g. *fool* in *foolish*) are retrieved and copied into erroneous positions in the positional frame. This is prior to their specification as phonological units (e.g. the syllables /fu:/ and /lɪʃ/), from which sound exchanges may arise.

Shifts

Interestingly, the distinct category of shift errors can fit into this account of the error mechanism for exchanges. Shifts tend to involve inappropriate ordering of adjacent elements relating to major and non-major lexical items:

(8) (a) unless you got somethin' to bétter do
(b) easy enoughly

Assuming that non-major items are part of the planning frame, Garrett suggests that 'shift errors are understandable as a consequence of the processes

which determine the siting of such elements in the (lexically interpreted) terminal string of the positional representation' (1982: 51).

Taken together with the observation (Cutler 1980) that shifts typically show stresses that move with their intended sites (as in example (8b)), this tends to support the view that shifts are best interpreted as movements of lexical items *vis-à-vis* certain stable frame-elements (see 3.4.1).

Thus we extend the statement of principle regarding frame elements generally from 'they do not exchange' to the stronger claim 'they do not move'.

7.4.4 ***Lexical factors***

We have distinguished between major lexical items, which are subject to exchange errors in the formulation of the functional-level representation, and frame elements, including non-major words and affixes, which form part of the positional-level frame. We shall now examine a little more closely the nature of errors involving major words. One of the striking observations to be made is that nearly all such errors appear to fall into either meaning-related or form-related types, and hardly ever into both at once.

Meaning-related errors

Substitution errors tend to involve antonyms (*hot/cold*; *love/hate*) and cohyponyms (*wash/brush your hair*), while blends are more common with synonyms, as in *dinner is ret* (*ready/set*), or, more occasionally, hyponym/superordinate pairs, as in:

(9) They have more protein than *meef* (meat/beef)

However, the last two types may be difficult to distinguish, since, in context, hyponym/superordinate pairs may be synonymous (*bitch/dog*).

These errors are presumably to be located in mapping the message-level representation onto the functional level. Consistently with his scepticism regarding the role of semantic features at the message level, Garrett observes cases of substitution where the meaning relationship is associative or inferential rather than decompositional, as in example (10) (Garrett 1984: 56):

(10) I just put it in the oven at very low *speed*
(the idea to be communicated was that the ham had to cook slowly)

Form-related errors

Substitution errors involving form similarities tend to show initial-position effects, as in:

(11) (a) because I've got an *appartment* now (appointment)
(b) they haven't been *married* . . . uh, measured . . .

Table 7.5 *Levels and processes involved in major lexical-class errors*

	Processes
Message level → functional level	
word substitution, blends	Lexical selection on basis of message-level properties, insertion into functional-level representation; meaning relevant, form irrelevant
word exchanges	Lexical assignment by functional role to positions in functional-level representation; grammatical category preserved, meaning, form irrelevant
Functional level → positional level	
word-substitutions, sound exchanges, stranding errors	Selection of word forms; meaning irrelevant, grammatical category irrelevant; position in word (e.g. initial, stressed syllable) relevant

Based on discussions in Garrett 1980a: 206–17.

and one may also observe some stressed-syllable effects (which may play a role in each of these examples). Unfortunately, the issue of computing word-final similarities between elements in a substitution is confused by the limited range of variation in inflectional and derivational suffixes in English. We have noted initial-position and stressed-syllable effects before (in 3.4.1, on sound exchanges and tip-of-the-tongue phenomena, and in 5.5, in the cohort model of word recognition).

Lexical retrieval

We may start to put together some of these observations concerning substitution and exchange errors involving major lexical items, within the framework of figure 7.5, as set out in table 7.5. The mapping of message-level onto functional level representations may be thought of as initiated through meaning relationships (as we have noted, involving more than purely semantic parameters). Word substitutions and word blends involving meaning similarity have their source here. Form characteristics are irrelevant, if by these we mean morphophonological forms. However, functional roles are also implicated in this selection process, so it would seem reasonable to expect some correspondence of grammatical class among such errors, inasmuch as there is some relation between grammatical and meaning categories. Thus, for example, *hot/cold* are more natural antonyms than *heat/cold*, at least in those contexts where *cold* is used as an adjective (*it's cold/hot in here*), since *heat/cold* in such a context would represent different grammatical classes. In so far as

grammatical class, reflecting functional role, represents abstract form, this may therefore be relevant.

Word exchanges seem to derive from a stage where initial, meaning-based, selection has been made, and processes of assignment of lexical items to sites in the functional-level frames are operative. Garrett suggests that 'sentence-level processes, once set in train, are neither monitored for nor couched in terms of meaning parameters' (1982: 57). Hence, grammatical category is preserved, but meaning, as well as morphological form, is irrelevant.

Word substitutions involving formal similarity, sound exchanges and stranding errors are all characterised by their independence of meaning and grammatical class, and by the influence of form. The word substitutions of this type arise from that phase of the word-retrieval process in which word forms are selected for subsequent insertion into the positional-level planning frame. Sound exchanges and strandings derive from processes which assign segmental specifications to sites in the positional-level frame. There is thus a parallel between these errors at the positional level and word exchanges at the functional level.

Closed-class items

Thus far we have concentrated mainly on major-class lexical items, and have merely noted that non-major words and grammatical affixes, or closed-class items, seem to be specified as part of the frame, and not to be subject to movement. Unlike major-, or open-class, items, there is an intimate link between closed-class elements and the phrasal configurations in which they occur. Garrett suggests that we regard them as specified, not by lexical retrieval, but rather by 'the (unknown) processes which select phrasal frames' (1982: 61) at the positional level, under the influence of functional-level constraints.

Segmental errors are very rare in these items, which may indicate that they are not specified for their segmental phonological structure until some point after open-class items have been so specified. This situation would then also account naturally for the form of the closed-class item {indefinite article} → *an* in:

(12) an _anguage *l*acquisition problem

and of the closed-class item {past} → /id/ in:

(13) Well, I *wait*ed him to *warn*.

The status of prepositions

A particularly interesting situation is found with prepositions. These behave like major-class items (noun, verb, adjective) in respect of their occurring in exchange errors, at the message-to-functional-level mapping: but, like frame elements, at the positional level they do not get involved in sound-exchange errors. Garrett addresses this situation by treating prepositions in a systematically ambivalent way: as lexical elements at the functional level which become 'demoted', through a cliticisation process, to the status of frame elements at the positional level. The notion of a clitic is basically a phonological one, so Garrett argues that the positional level be regarded primarily as phonological in nature. In this sense, 'phonological' does not exclude syntactic factors such as constituent boundaries and categories: 'phonological phrasing' and 'syntactic phrasing' coincide, to some important degree.

7.4.5 *Summary*

Garrett's model, as far as concerns us here, may be regarded as essentially embodying the distinction between two levels: the functional, consisting of abstract syntax and meaning-specified lexical items, vs the positional, consisting of phrasal groups of a syntactico-phonological kind, affixed elements and form-specified lexical items.

Further, the model claims that, by the criteria of

(a) grammatical class,
(b) phonological similarity, and
(c) domain of movement,

naturally occurring speech errors can be assigned to one or other of these levels. Errors that (a) respect grammatical class, (b) ignore phonological similarity, and (c) operate over relatively large domains, can be located at the functional level; those that (a) ignore grammatical class, (b) respect phonological factors, and (c) operate over relatively small domains, can be located at the positional level.

This situation reduces to two simple claims:

1. the positional level is blind to functional-level information;
2. the functional level is blind to positional-level information.

A model like this, having such independence or autonomy between levels, lends itself naturally to interpretation as serial in operation. According to this interpretation, both the functional- and positional-level representations are real-time constructs – that is, they have sequentially ordered constituents along the time dimension, and cannot communicate with each other because, for a given domain of processing, the functional level must be completed prior to the initiation of the positional level.

7.5 Serial versus parallel interpretations

A serious challenge to the serial-processing view, however, has been made in a study by Dell and Reich (1981). Working within the same framework as Garrett, they wish to recognise more complex flows of information between these levels than the serial model allows.

They took the claims of the serial model and operationalised them as follows.

1. Sound errors (i.e. of the positional level) should show no lexical bias. That is, in errors such as

(14) *f*itch *p*ork (pitch fork)

the fact that the exchange yields one non-word, *fitch*, and one real word, *pork*, is totally coincidental. The result could just as easily, and irrelevantly, been two real words or two non-words. Notice that what we may call the lexicality effect (see also ch. 5, section 5.2.3) here is taken as deriving from outside the positional level, specifically, from the functional level.

2. Word errors (i.e. of the functional level) should show no phonological bias. That is, in errors such as

(15) no-one is *taking* you into *talking* . . . (a nap)

the fact that two phonologically similar word forms are involved is purely by chance. This follows from the interpretation of the positional level as phonological, and autonomous from the functional level.

Dell and Reich set out to test each of these versions of the claims embodied in the serial model.

7.5.1 *Is there a lexical bias in sound errors?*

Garrett himself addressed this issue. He made an estimate of how often sound exchanges could be expected to create words by chance, by sampling word pairs from published interview data and exchanging their initial sounds. This suggested that words could be created by chance in this way about 33 per cent of the time.

Looking at Fromkin's (1973) corpus of errors and the MIT data up to that point in time, Garrett found word outcomes from initial-sound exchanges running at 40 per cent and 38 per cent respectively – not sufficiently above estimated chance level to provide evidence for a lexical bias.

Another point Garrett made was that, since real words are characterised by some very striking differences in their frequency of occurrence (*fife* is very much less frequent than *wife*), the lexicality effect should reveal itself also in terms of a frequency effect (see also ch. 5, section 5.2.3). In other words, if for example *pork* arising from an error such as *fitch pork* is actually the lexical item *pork* as in *I like apple sauce with my roast pork*, then the production of *pork* in error contexts ought to match the known frequency characteristics for the real word. More generally, most of the apparent-word outcomes from sound errors should consist of the more frequently occurring words in the language.

However, there seemed to be no relation between frequency of word outcomes from sound errors and frequency of words in the language, in the data Garrett examined: more than 60 per cent of the apparent-word outcomes had fewer than twenty occurrences in the Kučera and Francis (1967) word-frequency lists. This suggested that *pork* from *fitch pork* is not the real word *pork*, but a non-word that is homophonous with it.

However, this still leaves the issue as to whether there is what we may call a lexical-form effect, according to which error outcomes are more likely to converge on, or be homophonous with, actual words in the language. Does the existence of certain actual word-forms in the language 'pull' the errors into certain phonological patterns?

Garrett's finding that apparent-word outcomes occur at 38–40 per cent of the time in the data might seem to rule this out, in view of the estimate of chance level at around 33 per cent. But, in considering chance levels of homophony, we must take into account the phonological nature of what Dell and Reich call the *source* words: e.g. *red* has many phonological neighbours, and hence can easily slip into the apparent words *head*, *bed*, *fed*, *shed*, etc., while *pipe* has many fewer neighbours (if we just consider outcomes from initial-segment exchanges). Unless we take this factor into account, our estimates for chance levels, for particular word pairs, are going to be wrong.

Concentrating on examples of this type, word-initial consonant errors, before a following vowel, Dell and Reich found 363 instances in approximately 4,000 naturally occurring speech errors in their corpus at Toronto University (the Toronto corpus). Of these, 196 were exchanges, ninety-nine were anticipations (including some, e.g. *leading ... uh, reading list*, that were potential exchanges which the speaker caught in time), and sixty-eight perseverations.

Estimating chance level

The framework for discussion is as shown in table 7.6. In these terms, we can say that the first outcome-string provides the crucial test of the lexical-form effect, as far as exchanges and anticipations are concerned.

Table 7.6 *The framework for the analysis of exchange errors*

Intended words		Outcome string		
1st	2nd	1st	2nd	
pitch	fork	fitch	pork	exchange
		fitch	fork	anticipation
		pitch	pork	perseveration

Based on discussion in Dell and Reich 1981: 616.

This is the initial, and arguably the initiating, part of these errors, and the remainder, the second outcome-string, may be regarded as a residual or default element.

We have seen how Shattuck-Hufnagel's (1983) scan-copier model also embodies this view. In the case of an exchange, the scan-copier monitor, after the initial error yielding *fitch*, places the unassigned /p/ segment in the gap left by /f/ (in the schema for the second word). In anticipations, it ensures a second copying of the appropriate segment, /f/, in that slot.

But in perseverations, it is the second outcome-string that should be tested for lexical-form effects, since here the status of the first outcome-string is guaranteed (no error, hence a real word).

For each of the first two types of error, exchanges (n = 196) and anticipations (n = 99), a five-step procedure was carried out to estimate the appropriate chance level of word-form creation. This procedure is illustrated in table 7.7, for an illustration corpus of exchange errors (n = 4), from Dell and Reich (1981). The criterion for lexicality of the outcome string was whether it appeared in *Webster's seventh collegiate dictionary* either as an entry, a grammatical form of such an entry or a proper name from one of the appendices. The procedure was followed through for both first and second outcome-strings.

Dell and Reich point to two advantages in using such a procedure. First, because it uses the phonological properties of the 'sound-slipping' words themselves in the estimate, the result is sensitive to the average length, phonological structure, grammatical class and other properties of the words involved. Secondly, because all the slips studied involved initial prevocalic consonants, all the strings created in the matrices of steps 4 and 5 of table 7.7 were phonetically possible. So the estimate is not skewed at the outset by being based on combinatorially impossible instances.

Table 7.7 *Calculating chance expectations for lexical bias in sound errors*

Sample 'corpus' of four exchanges

pitch fork → fitch pork
Lawrence and Warden → Wawrence and Larden
postal code → coastal pode
chin tickled → tin chickled

Step 1 Create two lists

list one	list two
pitch	fork
Lawrence	Warden
postal	code
chin	tickled

Step 2 Determine the proportion of each initial phoneme in each list

list one	list two
/p/ = 0.50	/f/ = 0.25
/l/ = 0.25	/w/ = 0.25
/č/ = 0.25	/k/ = 0.25
	/t/ = 0.25

Step 3 Strip each word of its initial consonant

list one	list two
-itch	-ork
-awrence	-arden
-ostal	-ode
-in	-ickled

Step 4 Combine stems from list one with initial consonants from list two; determine if the resultant strings are words

	/f/ 0.25	/w/ 0.25	/k/ 0.25	/t/ 0.25	row sum
-itch	0	1	1	0	0.50
-awrence	0	1	0	1	0.50
-ostal	0	0	1	0	0.25
-in	1	1	1	1	1.00

0 = non-word average of row sums = 0.56
1 = word

The probability that the first outcome of an exchange will create a word is 0.56

Step 5 Combine stems from list two with initial consonants from list one: determine if the resultant strings are words

	/p/ 0.50	/l/ 0.25	/č/ 0.25	row sum
-ork	1	0	0	0.50
-arden	0	0	0	0.00
-ode	0	1	0	0.25
-ickled	1	0	0	0.50

0 = non-word, 1 = word average of row sums = 0.31

The probability that the second outcome of an exchange will create a word is 0.31

From Dell and Reich 1981: table 2, p. 618.

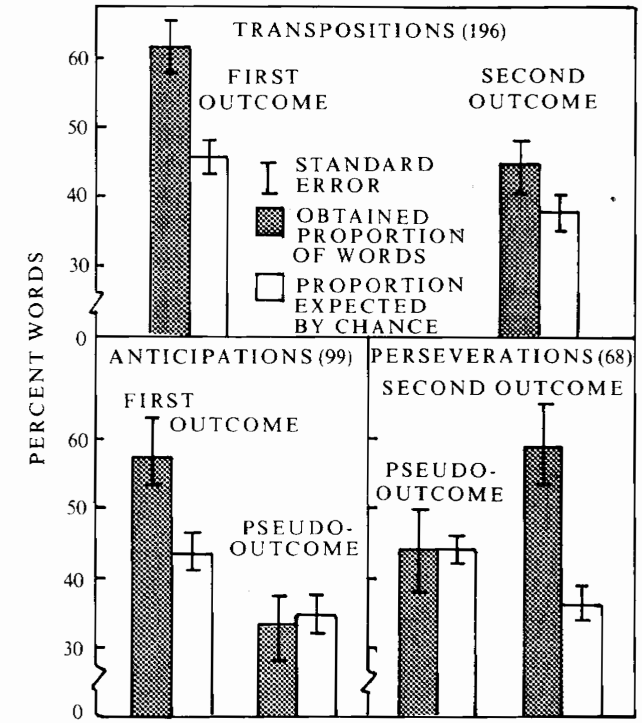

Figure 7.8 Proportions of word outcomes in initial consonant errors and chance estimates. (From Dell and Reich 1981: fig. 1, p. 619.)

The results are set out in figure 7.8. They indicate a strong lexical-form effect in the first outcome-string (note the gap between the range of the standard error for expected vs obtained proportions of words), but not in the second outcome-string. This is true for both exchange errors (transpositions) and anticipation errors. Concerning the perseveration errors, it was felt that their low number (n = 68) rendered the standard calculation procedure unreliable. Instead, the mean of the expectations derived from exchanges and anticipations was used (45 per cent and 44 per cent respectively for first outcome-strings, and 35 per cent and 36 per cent for second outcome-strings).

Dell and Reich conclude that there is strong evidence here that:

1. the mechanism of exchanges and anticipations is to be sought in the first outcome-string, and of perseverations in the second; and
2. this mechanism preferentially gives rise to word forms.

The answer to the question, 'Is there a lexical bias in sound errors?', would seem to be:

1. no, so far as we can tell, if by 'lexical bias' we mean that *pork* in *fitch pork* is the same item as the real word *pork*; but
2. yes, if by lexical bias we refer to the tendency for such errors to slip into patterns that are created in the language by existing phonological word-forms of the stock of lexical items.

7.5.2 *Is there a phonological bias in word errors?*

We now turn to the possible role of phonological (i.e. positional-level) factors in word errors. Following the outline of Garrett's model, we can deal with this under two headings: identification of words as complexes of meaning properties; and assignment of words as meaning complexes to roles in the functional-level structure.

Phonological factors in word-identification processes

How far might semantic substitution errors, e.g. *wash* for *brush* in the context ____ *your hair*, be contingent upon dimensions of phonological similarity, such as C(C)V/ʃ/? We have seen that Garrett's model (fig. 7.5) distinguishes meaning-based word-substitution errors as deriving from message-to-functional-level mappings, and form-based word-substitution errors as deriving from functional-to-positional-level mappings. If these are discrete, serially ordered stages, then word-substitution errors will be either of one type or the other. Apparent phonological similarities in target/outcome semantic substitutions will be there by chance only; so, obviously, we must again address the issue of obtaining reliable estimates of chance levels.

Dell and Reich (1981) identified 289 word substitutions, and 63 word blends, in the Toronto corpus, all involving content words. They determined the phonological similarity of the target/outcome pairs in each case by considering the identity of individual phonemes in first, second, third and fourth positions, and comparing them with chance estimates computed from the same corpus of substitutions and blends. These estimates, as before, are derived from all possible pairings of target words with outcome words in the corpus, and by marking the resulting percentage of identical phoneme matchings for each position in the word individually. The results show that *for word substitution and blends in general* there is a clear and sequential effect of phonological similarity: the actually occurring errors tended to share identical phonemes with their targets, especially towards the beginning of the word (see fig. 7.9).

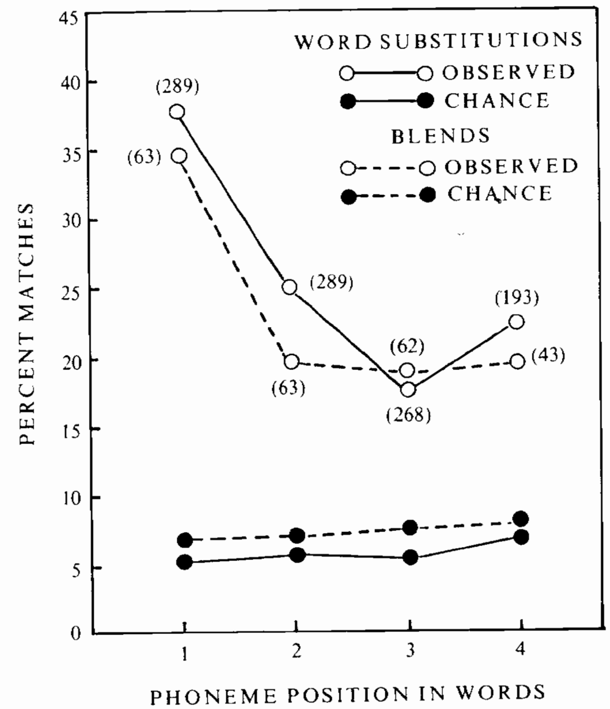

Figure 7.9 The phonological similarity between the interacting words in word-substitution errors and blends. Numbers in parentheses are the number of cases on which each percentage is based. (From Dell and Reich 1981: fig. 2, p. 623.)

Now, the question arises as to how far this is true of the subgroup of substitutions and blends that are meaning-based. Accordingly, Dell and Reich identified those cases where a meaning relationship such as antonymy, cohyponymy or superordinate–hyponym could be discerned. In this process, morphologically and associatively linked words such as *optician/optometrist*, and *Hungarian rhapsody/restaurant*, were excluded, to avoid possible unwarranted phonological effects. This yielded 130 'semantic' errors and 159 'non-semantic' ones (see fig. 7.10). It appears that, while the phonological effect is less marked on the semantic group, it is still greater than would be expected by chance, and hence runs counter to the serial-processing model. It is also greater than the percentage of phoneme matches found on a set of 464 near-synonyms (Whitten, Suter and Frank 1979), so the possibility that this result might be influenced by a coincidence of semantic and phonological dimensions in the language itself must be ruled out.

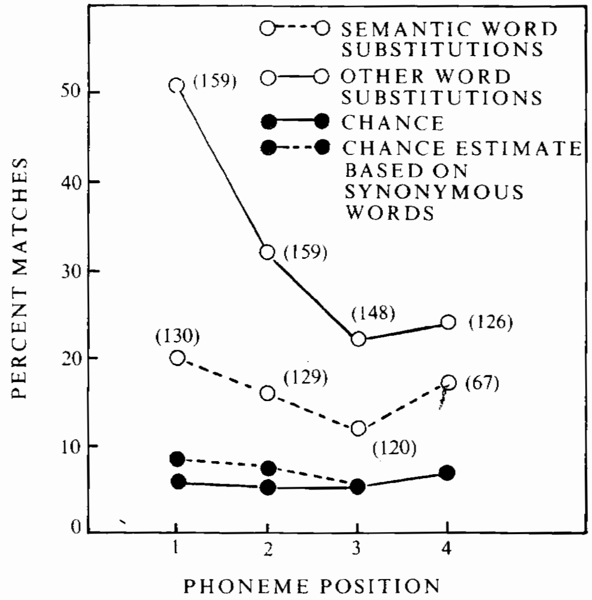

Figure 7.10 The phonological similarity between the interacting words in semantic and other word-substitution errors. Numbers in parentheses are the number of cases on which each percentage is based. (From Dell and Reich 1981: fig. 3, p. 624.)

Dell and Reich make the point, also, that nearly all the phonological matches in the semantic group were from target/outcome pairs that showed only a partial similarity: so the overall phonological effect is partial and pervasive (contributed to by all the items to some degree), rather than strong and intermittent (as would be the case if just a few items had shown near-identity as between target/outcome). This is exactly the wrong sort of result for a serial model of the Garrett type.

Phonological factors in word-assignment processes

Dell and Reich (1981) also looked at 155 content-word misordering errors from the Toronto corpus, and divided them into functional-level errors (n = 81) and others (n = 74), on the criteria of shared grammatical class and error domains of at least two intervening words (see 7.5 above). Phonological similarity of words participating in each error was assessed as before. The functional-level type errors showed a greater degree of phonological relationships than would be expected by chance, and the functional-level and other types of error were not different from each other in this respect. Once

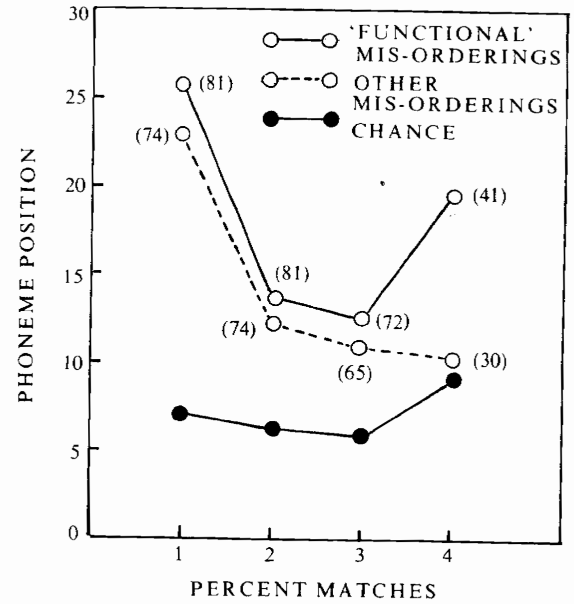

Figure 7.11 The phonological similarity between the interacting words in 'functional'-word mis-ordering errors and other word and morpheme mis-ordering errors. (From Dell and Reich 1981: fig. 4, p. 626.)

again, the functional level appears to be significantly influenced by positional-level phenomena (see fig. 7.11).

7.5.3 *Interpretation of the evidence*

The Dell and Reich study therefore makes three observations:

1. sound errors (involving prevocalic initial consonants in content words) tend to yield word forms;
2. meaning-based word-identification errors (substitutions) are open to phonological influences;
3. meaning-based word-assignment errors (exchanges) are open to phonological influences.

We must now ask how this information is going to be used in our understanding of production processes. We shall consider a number of possibilities, in turn.

1. Are the findings relevant? The Toronto corpus was collected, quite rapidly, by students of psycholinguistics at the University of Toronto. They were

instructed to carry a notebook at all times and to write down every little slip they heard over a one-month period, together with its context. This was repeated for five one-month periods.

Does this method lead to a lack of comparability between the Toronto corpus and other corpora? It is difficult to see how it might. Using a large number of collectors would reduce the possibility of a collector-perceptual bias (some individuals may be more sensitive to certain types of error), and the possibility of a sampling bias (relying on a few individuals to collect the data tends to narrow the range of conversational partners, topics and settings). But it is not clear that these advantages would lead to a data-base sufficiently different from the established corpora at UCLA and MIT, particularly since the Dell and Reich study is limited to easily detectable and central types of error. Possibly their decision to look just at prevocalic consonants in initial position of content words is a distorting factor; but they performed the same analysis on Fromkin's (1973) corpus and, although only sixty errors of the same type could be found there, they report essentially the same results. We have to conclude that the results are relevant.

2. Do we abandon the Garrett model? The demonstration of interactions between processing levels in Garrett's model is not, by itself, evidence against these levels. Dell and Reich distinguish between *constraints* (e.g. shared grammatical class, phonological-structure properties, domain of error, etc.) on the basis of which the levels are established, and *probabilistic tendencies* for interactions to take place between these levels. The evidence thus far is that we keep the model but abandon the strictly serial interpretation of the information flow within it.

3. How do we interpret the lexical-form effect? There are two parts to the answer. The first has to do with the claim that the effect has to be sought in the initiating part of the sound-based error, the source of the exchange, anticipation or perseveration. To go back to the *fitch pork* example, we would have to conclude that, whatever the formal similarities between the second element here and the real word *pork*, there is no lexical-form effect operative here, since the error is triggered by the first element, *fitch*. But in exchanges like *coastal pode*, we should conclude that the first element, *coastal*, shows a lexical-form effect.

The second part of the answer is that we are talking here about lexical *form*, rather than simply about *lexical* effects: this is because there seems to be little or no evidence that it is the real word *coastal*, as a form–meaning complex, that turns up in the first-element position here (complete with associations of white cliffs, mewing seagulls and the sound of waves on rocks, perhaps).

To the extent that this is true, we may not be dealing here with interactions between the functional and positional levels as such. Rather, we should perhaps think in terms of links between different factors in the generation of the positional level: (a) the existence of word forms in the lexicon; and (b) the filling of phonological segments in the articulatory memory, or buffer store, that controls the articulatory output of speech.

4. How do we interpret the phonological effect? This might well be thought of as a word-form effect as well, since the phonological dimension that is observed in semantic errors may derive from the stored phonological properties of the language, as found in the word forms in the lexicon.

This effect is truly one that links the functional and positional levels. A strict serial approach effectively looks on 'the lexicon' as two distinct lexicons, the semantic lexicon, activated by message-level factors, and consisting of word meanings; and a phonological lexicon (a better term perhaps is a word-form lexicon), consisting of stored word-strings.

The evidence from Dell and Reich's study is more consistent with a single lexicon, with connections between stored word-meanings and stored word-forms. Activation would occur in parallel at both the meaning and the sound-structure levels in such a lexicon. It would be a type of network memory-store (see section 7.3.3 above), with links also to non-linguistic, encyclopaedic knowledge. As Dell and Reich envisage the situation,

> retrieval processes in this lexical network occur by spreading activation with each activated node sending a proportion of its activation to all nodes connecting to it . . . In the mapping from some unspecified representation to the functional representation it is assumed that semantic nodes in the lexicon are activated. This activation then spreads throughout the lexicon. The functional representation is built up as grammatical rules select and order word constituents . . . However, the words that are available for selection (those that are highly activated) will have been influenced by nonfunctional factors as a result of spreading activation. In particular, words that are phonologically related to intended words will have become activated because activation spreads to them via phoneme nodes shared with intended words. (1981: 627–8)

The sort of network they are assuming is depicted in figure 7.12.

7.6 Conclusions

In this chapter we have reviewed, quite selectively, some current approaches and suggestions regarding the components of a production model. It must be emphasised that there is no settled view on these matters, and that much research remains to be done. The treatment of models of meaning, the

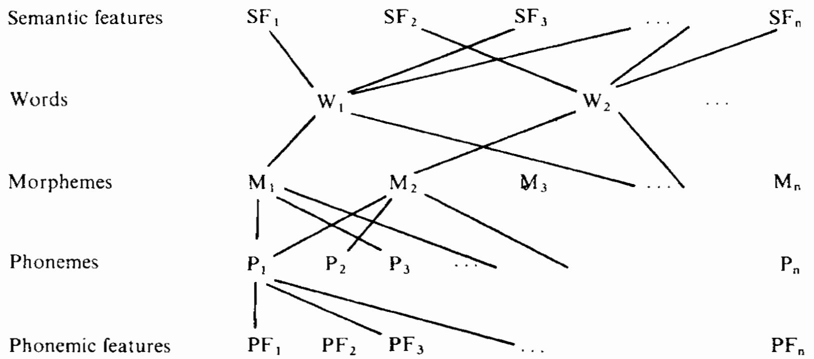

Figure 7.12 A network for the semantic and form properties of words. (Based on discussion in Dell and Reich 1981: pp. 627–8.)

lexicon and syntax here are offered as glimpses of how things might be organised, but obviously much of this area is controversial and open to further developments.

To tie things together a bit, let us return briefly to the layout envisaged in figure 4.1. Questions have to be raised both about the relations between the various components, and their internal structure. Concerning relations, we may pick out for mention here those between: (a) the message level and the lower levels, within both the syntax and the lexicon hierarchies; (b) the upper and lower levels within each of these hierarchies; and (c) the two major hierarchies themselves, of the lexicon and the syntax. We have outlined both serial and parallel positions in respect of (a) and (b), and it is not difficult to envisage how they might be extended to (c): given that the message level controls both lexical access and syntactic structuring, it could be that either the one is dependent upon the other, or that the two processing hierarchies interact, with lexical decisions affecting syntactic choices, and vice versa.

We may also recall the point made earlier regarding the level of signal production, namely that it characteristically reflects coarticulatory processing. Within the abstract language system, discrete linearisation of elements (phonological segments, lexical items and phrasal and clausal categories) is of the essence: but we have seen that not only does this not rule out the possibility of parallel, interactive processing, but that, indeed, certain effects appear to make this likely. What is still not clear is the level or levels at which linearisation takes place (given some non-linear concept of message structure, such as the mental model). One possibility is that 'given' topics tend to precede 'new' ones by virtue of their familiarity to the speaker and hence their relatively

easier lexical access; such a view would be consistent with the lexical hierarchy taking the lead, but interacting with the syntactic (Bock 1982).

More detailed working-out of interactive accounts of language production may be found in Stemberger (1985) and Dell (1986). In addition, it is naturally the case that message structures may be encoded both through the language-production hierarchy and through other communicative means, such as gesture: McNeill (1985) provides a view of the inter-relationship between these encoding hierarchies which emphasises their ability to be regulated by a unitary system of linearisation and rhythmic timing. Aspects of message-level structuring are discussed within the Parallel Distributed Processing (PDP) framework in Rumelhart, Smolensky, McClelland and Hinton (1986), from whom we may take the following observation with which to close not just this chapter, but also our main account of the issues in modelling normal language-processing (chs. 4 to 7):

> We believe that processes that happen very quickly – say less than 0.25 to 0.5 seconds – occur essentially in parallel and should be described in terms of parallel models. Processes that take longer, we believe, have a serial component and can more readily be described in terms of sequential information-processing models ... We would caution, however, that when one chooses a formalism such as production systems and attempts to use it ... to describe the conscious sequential processes that occur at this slow time scale, it is important not to fall into the trap of assuming that the microstructure of these sequential processes should also be described in the same terms. (pp. 56–7)