
Does Performance Shape Competence?

Janet D. Fodor

Phil. Trans. R. Soc. Lond. B 1981 **295**, 285-295

doi: 10.1098/rstb.1981.0140

Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click [here](#)

To subscribe to *Phil. Trans. R. Soc. Lond. B* go to: <http://rstb.royalsocietypublishing.org/subscriptions>

Does performance shape competence?

BY JANET D. FODOR

*Department of Linguistics, University of Connecticut, U-145, Storrs,
Connecticut 06268, U.S.A.*

Are the peculiar design characteristics of human languages merely the result of evolutionary accidents that determined innate limitations on linguistic competence? Or do they stem from the fact that human languages are used for practical purposes, and that efficient use demands sentence structures that can readily be produced and parsed? The only way to decide this question is to look for functional explanations of language structure; the search should be informative, whether it ultimately succeeds or fails.

Some grammatical constraints have the effect of ruling out one structural analysis of a potentially ambiguous word string, thus disambiguating it and reducing uncertainty for the sentence parsing routines. So far, however, there has been no convincing account in these terms of the major 'island' constraints on transformational operations. It is *not* plausible, for example, that these are motivated by the strain of holding a transformational 'filler' constituent in memory until its associated 'gap' has been found, or by the difficulty of identifying gaps in certain sentential contexts.

I suggest that the parsing routines play a crucial but less specific role in determining the constraints. They call for *some* restrictions on filler and gap positions, but not for any *particular* restrictions. The actual pattern of permitted filler–gap relations stems from the expressive function of transformations. Only fillers and gaps that are central to expressive purposes are admitted; all others are excluded by the constraints.

Despite the difficulty of defining expressive function, this approach offers promising explanations for why different transformations, and similar transformations in different languages, are subject to different constraints. I conclude that the form of a language may indeed be shaped by its function, but only if 'function' is construed broadly enough to include communicative goals as well as the mechanics of the encoding and decoding operations.

Finally, there are indications that these functional pressures on language design can operate only within certain very narrow confines established by the nature of linguistic competence.

THE PROJECT

It would be interesting, to say the least, to be able to reconstruct the evolution of human language. Clearly we never will, but there are some subsidiary questions that may be more amenable to investigation. One of these concerns the relative influence of different components of the language faculty on the formal properties of the languages that we are capable of using. Do the design characteristics of human language facilitate language learning? Do they guarantee sentences that can be produced and understood rapidly and without error? Or do they merely reflect accidental, non-functional properties of the grammar representation centres of the human brain?

Whatever one's hunch about the right answer to these questions, the only way to set about finding out is to look for functional explanations of language structure. The search can be

[71]

informative whether it succeeds or fails. If, despite all our ingenuity, we keep coming up with nothing at all, we might reasonably conclude that functional considerations have had very little to do with the shaping of human language.

I need not emphasize the riskiness of this project. There could be significant functional influences that we are just too blind to see. There is also the opposite danger that we may be misled by some plausible-sounding pseudo-explanations. But it is surely worth a try. Quite apart from the fascination of speculating about why things are as they are, there are some more down-to-earth benefits to be hoped for. If we thoroughly understood the psychological mechanisms for language learning, and for sentence production and perception, we could begin to disentangle the demands that they make on language structure. In fact we know rather little about them at present, and so we can turn the enquiry upside down, and use the idea that languages are well designed for practical purposes as a source of hypotheses about how these mechanisms are structured.

Some interesting suggestions have been made recently by Wexler & Culicover and their colleagues (Wexler & Culicover 1980, and references therein) about the relation between language structure and language acquisition. I shall concentrate here on the sentence processing routines, and in particular on the routines for decoding the syntactic structure of sentences. I should have liked to include semantic comprehension processes and sentence production processes as well, but it is generally agreed that we know too little about these to be able to relate them to anything at all.

SUCCESSSES

One very obvious problem that the sentence processing routines face is ambiguity – both full ambiguity and temporary ambiguity that is resolved by subsequent words in the sentence. If, as seems likely, our brains lack the capacity to compute all possible analyses of an ambiguous word string, then we can only guess what the intended structure is; sometimes, inevitably, we shall guess wrongly. Ambiguity is not a problem in sentence production, since the producer knows which interpretation he intends and has only to find some well-formed realization of it. It is also not a problem within the formal grammar, which simply provides derivations for all sentences on all their interpretations, and is indifferent to whether two or more derivations happen to converge on the same surface form. So, whenever we come across grammatical devices that serve to reduce ambiguity, the suspicion is that they are there in the grammar to facilitate sentence comprehension.

I shall give just two examples from English. A complementizer introducing a subordinate clause is generally optional in English, but it is mandatory if the clause appears at the beginning of its sentence. We can say *It's a shame that he's so cantankerous* and *It's a shame he's so cantankerous*; we can say *That he's so cantankerous is a shame*, with the subordinate clause in sentence initial position, but we cannot say **He's so cantankerous is a shame*.

It has been suggested (Bever 1970; Chomsky & Lasnik 1977) that the parsing routines for English will tend to assume that a sentence-initial clause is the main clause unless there is some specific indication to the contrary. If so, this provides a plausible explanation of what is otherwise a completely *ad hoc* restriction in the grammar of English. Interestingly, sentence-initial subordinate clauses in left-branching languages such as Turkish and Japanese are often not marked as subordinate. In these languages, unlike English with its predominantly right-branching structure, it is very common for sentences to begin with a subordinate clause. The

parsing routines for these languages would therefore not be tempted to jump to the conclusion that the first clause is the main clause, and so would not need any grammatical restriction to protect them from this potential garden path analysis.

Another no-ambiguity constraint governs the pairing of fillers and gaps in sentences that have been transformed. In a sentence such as *What do you store boxes in?*, there is an extra noun phrase at the beginning, the word *what*. I shall call this the filler. In order to determine its role in the sentence, it must be associated with a gap later in the sentence, at the position from which it was moved by the transformational rule. In this example the gap is after the preposition *in*, for *in* must normally be followed by a noun phrase. (We cannot say **You store boxes in*, but only *You store boxes in the attic*, etc.) In the sentence *What do you store in boxes?*, the word *what* is again a filler, but this time the gap is after the verb *store*. The word *boxes* is the filler in the so-called ‘Tough Movement’ constructions, *Boxes are easy to store things in* and *Boxes are easy to store in things*, and the gaps in these sentences are again after the preposition or after the verb. These two constructions can be combined, to give a sentence such as *What are boxes easy to store in?*, which has two fillers and two gaps. This sentence *ought* to be ambiguous, for we have seen that either *what* or *boxes* can be associated either with the gap after the verb or with the gap after the preposition. But it is not so. The sentence can only be interpreted as asking: *It is easy to store boxes in what?*; it does not also have the meaning: *It is easy to store what in boxes?*. Here, as in all similar examples, the two filler-gap linkages must be nested inside each other, not intersecting (see Fodor 1978).

What are boxes easy to store in
 *

That is, the grammar of English includes a constraint that blocks one of the potential derivations of a sentence with multiple fillers and gaps, and thereby disambiguates the sentence. This constraint does not simplify the grammar, but it does simplify sentence processing. Instead of having to guess, and risk guessing wrongly, about which filler was intended to go with which gap, the sentence parsing device can immediately and confidently link them in the right way.

There are other constraints like these in English, and as we begin to investigate other languages we are finding many more examples. The details may differ from language to language, because which constructions are potentially ambiguous depends on a variety of language specific facts. But the general drift is the same: the grammar is complicated by restrictions on otherwise quite generally applicable rules, and the complication is directed specifically towards cases where the rules threaten to create ambiguous constructions.

We can conclude, then, that performance does shape competence, at least to this extent. A grammar that contains no-ambiguity constraints is not superior, considered as a formal rule system, to one that lacks these constraints; indeed it is often inferior, since no-ambiguity constraints notoriously demand very rich formal mechanisms. Yet people apparently prefer these more complex grammars, and this is surely because people have to *use* their grammars as well as storing them in their heads.

FAILURES

This positive conclusion must be tempered in at least two ways. First, there are innumerable ambiguities in natural languages that are *not* excluded by grammatical constraints. I shall return to this problem below. Secondly, there are many other constraints that do not obviously facilitate sentence processing. In particular, there are some very general, possibly universal, constraints on the positions of fillers and gaps in transformed sentences that seem to bear little or no relation to the sorts of problems that these sentences pose for the processing routines.

The most obvious problem for the parsing routines is that a transformationally moved constituent does not appear in a position appropriate to its semantic role in the sentence. For example, in the sentence *Who did you say that Tom was planning to invite — to the party?*, the word *who* represents the object of the verb *invite* and yet appears far from that verb in the word sequence. It is natural to suppose that the parser copes with a sentence like this by placing the word *who* into some sort of temporary memory store, and processing the rest of the sentence in the usual way until the verb *invite* is encountered and discovered to have no object; then the *who* can be retrieved from storage, and processed as the object of the verb. We might also reasonably assume that it is a strain for the parser to store an unassigned constituent while simultaneously processing other parts of the sentence, and that the strain will be greater the longer the filler has to be stored (see Wanner & Maratsos 1978).

The sorts of constraints that would alleviate this problem would obviously be constraints that conspire to keep a filler close to its gap. This is not, however, what the observed constraints do. Some of them are completely indifferent to the distance between filler and gap. For example, most languages prohibit a gap inside a relative clause if the filler is outside it. We cannot say **Who was Tom planning to invite the girl [that went to school with —]?*, where the gap associated with *who* is inside the relative clause that modifies *the girl*. But it is equally ungrammatical to say **Who was the girl [that went to school with —] planning to invite Tom?*, where the gap is also inside a relative clause but is now only eight, rather than twelve, words away from its filler. Other constraints actually increase the distance between a filler and its gap. In English, a gap inside a subject clause cannot be associated with a filler outside it. We cannot say **Who was [that Tom had invited —] announced at the planning session?*. But if the subject clause is extraposed to the end of the sentence, the gap is perfectly acceptable. We can say *Who was it announced at the planning session [that Tom had invited —]?*, even though the gap is now eleven, rather than only five, words away from its filler. The fact that transformational rules require the parser to associate fillers and gaps over long stretches of intervening material thus does not offer any explanation for why the transformational rules are restricted as they are.

A rather different parsing problem looks at first sight to be more promising. We are assuming that the parser stores a filler constituent and then looks for its gap in the remainder of the sentence. But gaps are not always easy to find. They are not overt items of the input but are merely boundaries between items, and sometimes they are indistinguishable from normal word boundaries. In the sentence *Which book were you reading — to the children?*, the gap follows the verb *reading*, and represents its direct object. The verb *reading* can be used either transitively or intransitively; we can say *John was reading the book* or *John was reading*. And there is no way to tell, from the first few words of the transformed sentence, which of these two verbs it contains. If *reading* is transitive, its object must be the phrase that has been gapped; if *reading* is intransitive, the position after it is not a gap position. Only the end of the sentence indicates which is the

case. For example, the sentence *Which book were you reading to the children from — ?* starts in exactly the same way as the other, but here the verb must be intransitive since the gap clearly follows the preposition *from*.

To protect the parser from this uncertainty, the constraints on transformations would have to ensure that potential gap positions never coincide with the positions of optional constituents. They might prohibit a gap if the gapped constituent is optional. In this case the sentence *Which book were you reading to the children from — ?* would be grammatical but the sentence *Which book were you reading — to the children?* would be ungrammatical. Or they might prohibit the omission of optional constituents in a sentence containing a gap. Then the sentence *Which book were you reading — to the children?* would be grammatical, but the sentence *Which book were you reading to the children from — ?* would be ungrammatical. Either way, the parser would be able to tell, when it reached the position after *reading*, whether or not this was the true gap.

Once again, the observed constraints differ sharply from the predictions. There is no sign of any grammatical conspiracy against just those gaps which create uncertainty for the parser. Both of the sentences we have been considering are permitted by the grammar, despite the problems they pose. And other sentences are excluded, even though their gaps are quite unambiguous. In English, for example, there can be no gap in the subject position of a *for-to* complement clause; we cannot say **Which book were you waiting for — to be sent from the library?*. Yet a gap in this position would be immediately recognizable, since the subject is not an optional constituent. Similarly, we cannot say **Who did you speak to — and Tom?*, even though this gap, between a preposition and a conjoined noun phrase, could not possibly be construed as anything other than a gap.

The fact that gap identification problems do not correlate with the observed pattern of constraints is particularly disappointing because there is an interesting story to be told about why gaps differ in their distribution from resumptive pronouns. Pronouns are not subject to the constraints that we have been considering, even when they have the same syntactic role as gaps. We can say *This chair is too wobbly for anyone to stand on —*, where *this chair* is a filler associated with a gap after the preposition; or we can say *This chair is too wobbly for anyone to stand on it*, with a pronoun in place of the gap. The *for-to* constraint rules out the gap construction ** This chair is too wobbly for — to be used in the classroom*, but it does not rule out the pronoun construction *This chair is too wobbly for it to be used in the classroom*. From the point of view of sentence processing, this difference makes sense. Unlike a gap, a pronoun does appear as an overt element of the input word string, and its position in the sentence is perfectly clear. The absence of constraints on pronouns thus encourages the idea that the constraints on gaps exist because gaps are hard for the parser to find. Why, then, are they not sensitive to the clear differences in detectability between different gap positions?

I cannot review here all the other suggestions that have been made about how sentence processing problems might have shaped the constraints. I have argued elsewhere (Fodor 1981) that none of them succeeds in detail. The only useful point that emerges is that sentence parsing would be assisted by the existence of *some* fairly stringent limits on where gaps can appear. Clearly, the ideal situation would be one in which the parser, having detected a filler constituent, could tell immediately where its gap must be, without having to search for it at all. This would be so if each filler were associated by the grammar with a unique gap position. For some rules, such as Equi Noun Phrase Deletion and Subject Raising, this is so. In the sentence *Mother expects — to persuade you to talk to Susan* there is a filler noun phrase, *Mother*,

associated with a gap at the subject position of the subordinate clause; no other gap position is permitted. The grammar does not allow **Mother expects you to persuade — to talk to Susan*, or **Mother expects you to persuade Susan to talk to —* (meaning, respectively, that Mother expects you to persuade Mother to talk to Susan, and that Mother expects you to persuade Susan to talk to Mother). Other transformational rules do not have unique gap positions. A fronted question word, as we have seen, can be associated with a subject gap, an object gap, a gap in a prepositional phrase, a gap in a subordinate clause, and so on. Even here, however, the constraints cut back quite sharply the number of different gap positions that the parser must consider, and so the idea that the constraints are designed to offer some *general* assistance in sentence parsing still has some plausibility. What has to be abandoned is the hope of relating the particular profile of observed constraints to specific details of the gap detection routines.

THE EXPRESSIVE FUNCTION OF TRANSFORMATIONS

If our goal is to find some performance motivation for the constraints, we cannot shrug off their details as unimportant, for they are by no means haphazard. Similar patterns are found in language after language. They may perhaps be attributable to the language learning routines, or to the sentence production routines. But if not, the obvious conclusion to draw would be that they stem from some quirk in the innately determined character of linguistic competence.

I think this conclusion would be unwarranted, because it fails to take into account another facet of language, one that does not fit squarely on either side of the competence–performance distinction but bridges the two. What I have in mind is the fact that transformational rules serve a purpose in natural languages; they extend the range of expressive distinctions that we can signal in our utterances. We have made enormous progress in linguistics by concentrating on the formal characterization of grammatical rules, and ignoring the question of their expressive role. But when we examine the distribution of fillers and gaps it becomes quite clear that expressive considerations are the determining factor. Specifically: the constraints exclude filler–gap associations that do relatively little expressive work, and tolerate those that are most valuable for expressive purposes.

A natural language has a phrase structure base, overlaid with transformations. The phrase structure system is ideal for the expression of predicate–argument relations. Very roughly, a clause consists of a verb associated with a set of noun phrases that represent its semantic arguments; together, these elements specify who did what to whom. In addition to these predicate–argument relations, and largely independent of them, we find other semantic and pragmatic phenomena. The scope of questions and topicalizations can vary. The questions word in *John knows who Mary noticed that you were talking to —* has the same role in the predicate–argument structure as it does in *John knows that Mary noticed who you were talking to —*, but the scope of the question is quite different. There are also variable-binding phenomena. For example, *Barbara wants — to win* means that Barbara wants Barbara to win, But *Everyone wants — to win* doesn't mean that everyone wants everyone to win; instead, its meaning is what logicians would represent by the formula $(\forall x) (x \text{ wants } x \text{ to win})$, in which the subject of *to win* is a variable bound by a quantifier.

Movement and deletion transformations provide a means for marking these additional semantic and pragmatic properties in the syntactic structure of a sentence. The deletion of a noun phrase that is to be interpreted as a bound variable effectively distinguishes it from other

noun phrases that can have independent reference. The movement of a questioned noun phrase permits question scope and predicate–argument relations to be signalled simultaneously. The moved noun phrase has, in effect, two positions in the sentence. In *John knows who Mary noticed that you were talking to* — , the surface position of *who* indicates that the question comprises the whole of the complement of *knows*, and the gap position indicates that the focus of the question is the object of *were talking to*.

If we had been given the task of designing a language for the expression of the sorts of things that people typically want to express, we could have been very proud of ourselves if we had thought up such a clever scheme. Of course, there are other formal devices that would have been equally useful, and indeed there are languages that handle the matter very differently from English. There are the so-called free word order languages, which use inflectional morphemes to mark predicate–argument relations and are thus at liberty to order phrases in accord with their scope or discourse role without the need for specific transformational rules to move a constituent from one fixed position to another. And there are languages that use special particles to mark the scope of a question, rather than moving the questioned phrase, and languages that use resumptive pronouns where English uses gaps, and so on. Filler–gap languages like English represent just one good solution to the demands of expressiveness. Incidentally, it is rather clear that the parsing routines for these other kinds of language will differ quite considerably from those for English. The common thread that unites all human languages is not how their sentences are processed but the fact that they all, in their different ways (and, admittedly, to different degrees), have the means for making the same sorts of expressive distinctions. My claim that the constraints on transformations are shaped by expressive factors is just a specific instance of this more general observation.

PREDICTING THE DETAILS OF THE CONSTRAINTS

My proposal is that the prohibited gap positions in a language are the ones that speakers would make use of relatively rarely even if it were acceptable to use them. Under pressure from the performance mechanisms, the grammar is required to relinquish *some* gap positions, and it chooses to relinquish those that do the least expressive work. This proposal will remain vague in the absence of some measure of the expressive usefulness of different gap positions, but it is clear, in principle at least, how to set about pinning it down. What we need is to find some language L that is structurally similar to English but which makes use of pronouns where English has gaps. We could then examine a large corpus of conversation and writing in language L, and determine the most frequent positions for these pronouns. The expectation is that the acceptable gap positions in English will coincide with the positions at which the pronouns in language L tend to cluster. I have not conducted this experiment, but I am prepared to make some guesses about what we would find.

We know that different transformations are subject to different constraints, and I shall start by considering the relatively unconstrained rule that fronts question words in English. There is nothing about the semantics of questions that favours the questioning of one constituent of a sentence rather than another. A speaker may wish to remedy his ignorance about any aspect of an event, represented by any constituent in the sentence that describes it. So the focus of a question can be expected to vary widely, and it would be suicidal from an expressive point of view if a filler–gap language were to insist on a single gap position in all questions. On the other hand, it would not be surprising if some positions were made use of rather rarely. In

particular, it seems likely that the focus of a question will usually be a major constituent of the sentence, and will not often be a phrase that is embedded within a modifying constituent. That is, in language L we are more likely to encounter sentences of the form *Which girl did you meet her at the party at which John turned somersaults in the conservatory?*, than sentences of the form *Which girl did you meet John at the party at which she turned somersaults in the conservatory?* (Remember that I am not suggesting that the latter would never occur, but only that it would be relatively uncommon.) If this is so, it helps to explain why *wh*-gaps in English are permitted in a wide variety of positions, but not in modifying clauses such as relative clauses and adverbial clauses.

The Equi Noun Phrase Deletion transformation, as we have noted, is even more heavily restricted than *wh*-movement. Equi NP Deletion is a lexically governed rule, applying only in constructions that contain verbs like *try* or *want* or *persuade*. These verbs express a relation between an individual and a property or action. For example, *Barbara tried to impress John* expresses a relation between someone who tries, Barbara, and the action of impressing John. I would submit that more often than not, the person who tries will be the agent of the action intended to result from the trying. That is, sentences of the form *Barbara tried for herself to impress John* will occur more commonly in language L than sentences of the form *Barbara tried for John to impress her* or *Barbara tried for Bill to hope that she had impressed John*. Agents are realized syntactically as subjects in the unmarked case, and so the pronouns in these constructions of language L can be expected to cluster at the subject position of the complement clause. Therefore for these constructions, in contrast to questions, it would not be at all unreasonable for a filler-gap language to limit its gaps to subject position.

Similar considerations apply to the gap position for verbs like *persuade* and *promise*, but here there is a choice of filler. In *Barbara persuaded John to go home* and *Barbara promised John to go home*, the filler could in principle be either the subject of the main clause or its object. In fact the grammar insists that the filler is the subject, *Barbara*, in the *promise* sentence, and the object, *John*, in the *persuade* sentence. In view of the semantics of *promise* and *persuade*, this difference is not at all surprising. When Barbara promises John, she commits herself to doing something; when Barbara persuades John, she gets him to do something. That is, the agent of the embedded clause will typically be coreferential with the subject of *promise* and with the object of *persuade*, and so these will be the most commonly needed fillers when the embedded agent is gapped.

The grammatical constraints thus crystallize pre-existing tendencies in the use of the language. They may thereby render unsayable certain things that we might, occasionally, want to say. In this particular case, however, the loss of expressive capacity is minimal, for there are ways of getting around the constraints if we really need to. The Passive transformation allows non-agents to be realized as subjects. So, though we cannot say **Barbara tried for John to impress* —, we can say *Barbara tried to be impressed by John*.

The existence of alternative ways of encoding a message is important to many other restrictions on gaps. Keenan & Comrie (1977), in their study of universal trends in the positioning of gaps in relative clauses, noted an odd counterexample in certain Malayo-Polynesian languages. These languages permit gaps in subject position, and in indirect object and other oblique noun phrase positions, but not in direct object position even though direct objects are more readily gapped than oblique noun phrases in other languages. What Keenan & Comrie observed was that these exceptional languages have a particularly well developed passive construction. There is no need for the grammar to permit direct objects to be gapped, because a direct object can very naturally be converted by passivization into a subject, and

then gapped in that position. This otherwise very curious restriction thus makes excellent sense on the assumption that the grammar is under pressure to give up any gap position that does not carry a significant expressive burden.

I shall give just one more example. A complement clause in English that functions as the object of the main verb can contain a gap, and so can a complement clause that functions as the subject but has been extraposed to the end of the sentence. We can say *This last question, John thinks that people will misinterpret —*, and *This last question, it's absolutely certain that people will misinterpret —*. But as we have noted already, a subject complement that has not been moved cannot contain a gap. We cannot say **This last question, that people will misinterpret — is absolutely certain*. This Sentential Subject Constraint resists a semantic explanation, for extraposition of a subject clause leaves the meaning of the sentence unchanged while completely changing the acceptability of gaps. Nevertheless, the constraint is understandable in terms of the preservation of expressive power.

English is a predominantly right-branching language. Relatively little recursion is permitted on left branches, and even where it is permitted speakers tend not to exploit it. An extraposed subject complement, like an object complement, branches rightward, and is generally judged to be more natural than an unmoved subject clause, which branches leftward. The difference in naturalness increases as the subject clause becomes more complex. Compare, for example, *It is absolutely certain that when John reads through the exam he will become convinced that people will misinterpret the last question*, and *That when John reads through the exam he will become convinced that people will misinterpret the last question is absolutely certain*. Thus left-branching sentential subjects will, on the whole, occur less frequently and be less complex when they do occur, than right-branching extraposed subjects. It follows that there will be rather few phrases within left-branching subjects that are even candidates for being gapped, though there will be many within right-branching subjects. If there is pressure for the grammar to prohibit gaps in either of these constructions, the choice is therefore obvious. And the synonymy of the two constructions renders the prohibition almost completely painless, since what is unsayable in one form will be sayable in the other.

In support of this explanation, we may note that there is no Sentential Subject Constraint in predominantly left-branching languages. Subject phrases in these languages conform to the prevailing branching pattern and can be of any degree of complexity. They therefore contain many potentially useful gap positions, so that a prohibition on gaps within them would severely limit what can be said, especially as the alternative extraposed form is either highly marked or non-existent in these languages.

CONCLUSION

I have tried to show, though admittedly with only a handful of examples, that the constraints on fillers and gaps are not arbitrary, but also do not admit of a performance explanation in the usual narrow sense. The sentence parsing device (and perhaps the other performance mechanisms too) may call for *some* constraints, but the shape of the constraints is the result of a compromise between these pressures and the pressure towards maximum freedom of expression.

One practical consequence of this for our research methodology is that we cannot, after all, hope to learn very much about the details of the performance mechanisms by looking at the details of language structure. The one thing that we have learned about the performance

mechanisms is that they are quite robust. They are not so finely tuned that they will jam unless the input conforms to some very narrow specifications. All they need, apparently, in order to operate effectively, is an approximate upper bound on the amount of uncertainty they have to contend with.

I shall end by restoring linguistic competence to a more central position in the ‘negotiations’ that determine how languages should be structured. Earlier, I noted that grammars contain constraints to exclude some ambiguities in sentences, but not all. And in some cases there are no expressive asymmetries that discriminate between the ambiguities that are tolerated and those which are not. It is possible that certain ambiguities are more disruptive of sentence processing than others, and that the grammar concerns itself only with the most serious cases. But as a general principle this is not at all well supported. For example, multiply centre-embedded sentences, such as *The snow the match the girl lit touched melted*, are appallingly difficult to process, and so are sentences of the form *The horse raced past the barn fell*. Why, then, does the grammar of English not exclude these constructions? The answer, I believe, is that the constraints that would be needed to exclude them – without simultaneously excluding a host of perfectly manageable and useful sentences as well – are unnatural and complex from a grammatical point of view. They require the grammar to be able to count, or to attend to complicated interactions between syntactic rules and the phonological forms of certain lexical items. And this the grammar will not do. In other words, the grammar representation system – linguistic competence – acts as a filter on potential constraints motivated by functional considerations. However much the parser might benefit from some constraint, and however little damage it would do to the expressive capacity of the language, that constraint will be rejected if it cannot readily be encoded into the grammar.

This factor of encodability in the grammar is also important, I would argue, to the constraints on fillers and gaps that we have been considering. The difference between Equi Noun Phrase Deletion, with its unique gap position, and *wh*-movement, with its wide range of gap positions, does correlate with a difference in their expressive functions. But we might wonder whether this is sufficient to account for the rather sharp contrast between the two kinds of rule, which is found in many different languages. There are rules whose gaps may appear anywhere at all except where prohibited by the general ‘island’ constraints; and there are rules whose gaps may appear only as the subject of the clause immediately subordinate to the matrix clause; but there appear to be no rules whose gaps are uniquely restricted to any other position in the subordinate clause.

In a generalized phrase structure grammar of the kind that Gazdar (this symposium) has proposed, this division of rule types follows from simplicity considerations. A clause that contains a noun phrase gap is categorized as S/NP. The /NP annotation on the S node is passed down through the nodes of the phrase maker to the lexical level, where it is realized as a gap. It can be prevented by the rules from passing down into a relative clause, or a sentential subject, or other syntactic ‘island’. But to bring it to rest at one specific predetermined position in a sentence would demand very costly constraints in the grammar. There is just one exception, and this is where the unique target position is the highest subject position within the S/NP constituent. In this case only, the gapped clause consists of a complete phrase; it is exactly equivalent to a $\bar{V}P$. Therefore, instead of stipulating that some construction requires an S/NP, the grammar can stipulate that it requires a $\bar{V}P$. The unique gap position is thereby guaranteed without adding any complexity to the grammar.

Though this is not the place to explore the issue, I would just note that this simplicity explanation does not hold within an explicitly transformational grammar. What I should like to be able to show is that *all* the observed constraints on fillers and gaps can be encoded directly and simply into context-free phrase structure rules, given only the slash annotation to mark nodes that dominate gaps. If this proves to be possible, it would strongly support this particular theory of grammars. But more to the present point, it would also indicate that, whatever functional pressures a human language is subject to, they must operate within certain very narrow confines established by the nature of linguistic competence.

REFERENCES (Fodor)

- Bever, T. G. 1970 The influence of speech performance on linguistic structure. In *Advances in psycholinguistics* (ed. G. B. Flores d'Arcais & J. M. Levelt), pp. 4–30. Amsterdam: North-Holland.
- Chomsky, N. & Lasnik, H. 1977 Filters and control. *Ling. Inquiry* 8, 425–504.
- Fodor, J. D. 1978 Parsing strategies and constraints on transformations. *Ling. Inquiry* 9, 427–473.
- Fodor, J. D. 1981 *Parsing, constraints, and the freedom of expression*. (In preparation; to be published by Bradford Books, Montgomery, Vermont.)
- Keenan, E. L. & Comrie, B. 1977 Noun phrase accessibility and universal grammar. *Ling. Inquiry* 8, 63–99.
- Wanner, E. & Maratsos, M. 1978 An ATN approach to comprehension. In *Linguistic theory and psychological reality* (ed. M. Halle, B. Bresnan & G. A. Miller), pp. 119–161. Cambridge, Massachusetts: M.I.T. Press.
- Wexler, K. & Culicover, P. W. 1980 *Formal principles of language acquisition*. Cambridge, Massachusetts: M.I.T. Press.