REMARKS ON PROCESSING, CONSTRAINTS, AND THE LEXICON*

Thomas Wasow
Stanford University

Linguists have long recognized the desirability of embedding a theory of grammar within a theory of linguistic performance (see, e.g., Chomsky (1965:10-15)). It has been widely assumed by transformationalists that an adequate model of a language user would include as one component some sort of generative grammar. Yet transformational grammarians have devoted relatively little energy to the problem that Bresnan (in press) calls "the grammatical realization problem": "How *would* a reasonable model of language use incorporate a transformational grammar?" When this question has been raised, little support could be adduced for the hypothesis that the operations of transformational grammar play a part in speakers' or hearers' processing of sentences (see Fodor, *et al* (1974; chapter 5)). Instead of concerning themselves with questions of processing, transformationalists have concentrated their efforts (at least in the last decade or so) on the problem of constraining the power of their theory. The goal of much recent research has been to construct as restrictive a theory of grammar as possible, within the bounds set by the known diversity of human languages (see, e.g., Ross (1967), Chomsky (1973), Bresnan (1976), Emonds (1976), and Culicover and Wexler (1977) for examples of this type of research).

Computational linguists, on the other hand, have not explicitly concerned themselves very much with the problem of constraints (but see Woods (1973; 124-5) for an exception). Rather, their goal has been to find effective procedures for the parsing and processing of natural language. While this is implicitly a restriction to recursive languages, the computational literature has dealt more with questions of processing than with how to limit the class of available grammars or languages.

In previous papers (Osherson and Wasow (1976), Wasow (in press a, 1978)) I have argued for the legitimacy of the quest for constraints as a research strategy. I have argued that a theory that places limits on the class of possible languages makes significant empirical claims about human mental capacities, and can contribute to a solution to "the fundamental empirical problem of linguistics" (as Chomsky has called it) of how children are able to learn languages with such facility. I have tried to show that such psychological claims can be made, without making any assumptions about what role grammars play in performance. In short, I have argued that a theory of grammar can make significant contributions to psychology, independent of the answer to the grammatical realization problem.

Recent work by Joan Bresnan (in press) takes a very different position: she has suggested that transformationalists ought to pay more attention to the grammatical realization problem, and that considerations of processing suggest radical modifications in the theory of transformational grammar. Further, she argues that there is ample grammatical evidence

for these modifications. In this paper I will suggest some extensions of her proposals, and will explore some of their empirical consequences. Further, I will argue that her framework makes it possible to impose rather restrictive constraints on grammatical theory. Thus, I will argue that the grammatical realization problem and the problem of constraining transformational theory, while logically independent, are both addressed by Bresnan's proposals. If I am correct in this, then Bresnan's "realistic transformational grammar" represents a major convergence of the concerns of transformational and computational linguists.

My presentation will consist of three parts. First, I will briefly sketch Bresnan's framework. Second, I will suggest some extensions of her proposals and point out some consequences of these extensions. Third, I will propose how her framework can be constrained, and indicate certain desirable consequences of my proposals.

The primary innovation of Bresnan's framework is that it eliminates a large class of transformations in favor of an enriched conception of the lexicon. The grammar that results is one that Bresnan claims is far more realistic from a processing point of view than other versions of transformational grammar. She points out striking similarities between her proposals and recent computational and psycholinguistic work by Kaplan and Wanner, and she argues that Augmented Transition Networks can provide at least a partial answer to the grammatical realization problem within her framework.

I will now sketch very roughly what Bresnan's "realistic" transformational grammar is like. Rules like passive, dative, and raising rules, which are "structure-preserving" (in the sense that their outputs are structurally identical to independently required base-generated structures) and "local" (in the sense that the elements affected are always in the immediate environment of some governing lexical item, usually a verb), are eliminated from the transformational component and relegated to the lexicon. Lexical entries include, among other things, (strict) subcategorization frames and more abstract representations which Bresnan calls "functional structures" or "predicate argument structures". Subcategorization frames give the syntactic environments in which the lexical item may appear; these are expressed in terms of a basic set of grammatical relations, including "subject" and "object". These notions, while universal, are instantiated differently in different languages; for example, Bresnan takes essentially the structural definitions of "subject" and "object" proposed by Chomsky (1965; 71) as language-specific characterizations of these notions for English. Functional structures give a more abstract representation of the elements mentioned in the subcategorization frame, indicating what their "logical" relationships are. Thus, the

functional structure corresponds very roughly to the deep structure in the standard theory of transformational grammar; and the subcategorization frame corresponds even more roughly to the surface structure.

What the standard theory did with local structure-preserving transformations Bresnan can do in either of two ways. Relationships like active/passive are handled by positing two separate lexical entries for active and passive verb forms. The productivity of this relationship can be accounted for by means of a lexical redundancy rule, which would say, in effect, that corresponding to the typical transitive verb there is an intransitive verb which looks morphologically like the perfect form of the transitive, and whose subject plays the same logical role (i.e., in the functional structure) as the object of the transitive verb. Bresnan's other way of replacing local structure-preserving rules is illustrated most clearly with the raising rules. Raising to object position, for example, is used to capture the fact that the NP which is syntactically the object of one clause is logically not an argument of that clause at all, but a subject of the subordinate clause. Bresnan expresses this simply in terms of the relationship between the subcategorization frame and the functional structure; that is, the object of the main clause plays no role in the functional structure of that clause, but is "passed down" to play a role in the next clause down. In the interests of brevity I will not illustrate Bresnan's framework here. Rather, I will refer the interested reader to her paper, and go on to indicate my reasons for seeking to modify her proposals.

My primary motivation comes from some earlier work of mine (Wasow (1977)), which argued against the elimination of local, structure-preserving transformations. My argument was based on the observation that there are two similar but distinct classes of linguistic relationships whose differences can be expressed rather naturally as the differences between transformational rules and lexical redundancy rules. The clearest example of this is the English passive. It has often been suggested that some passive participles are adjectives and others verbs; I pointed out that adjectival passives and verbal passives differed in certain systematic ways. My central claim was that the surface subject of adjectival passives was always the deep direct object of the corresponding verb. For example, a passive participle which is demonstrably adjectival (e.g., because it is prefixed with *un-* or immediately follows *seem*) may not have as its surface subject the "logical" subject of a lower clause, the indirect object, or a chunk of an idiom: **John is unknown to be a communist*; **John seemed told the story*; **Advantage seemed taken of John*. A verbal passive, in contrast, could have as its subject any NP which could immediately follow the corresponding active verb: *John is known to be a communist*; *John was told the story*; *Advantage was taken of John*. This, I claimed, would follow from the hypothesis that adjectival passives are formed by a lexical redundancy rule, whereas verbal passives are transformationally derived, if lexical redundancy rules are "relational", in the sense that they are formulated in terms of grammatical relations such as subject and object, whereas transformations are "structural", i.e., they are operations on phrase structure tree.

It is evident that my earlier position is inconsistent with Bresnan's recent proposals. My extensions of her ideas, developed in collaboration with Ron Kaplan, are in part an attempt to capture within her framework the distinction my earlier paper sought to explicate in terms of the lexicon/transformation contrast. They are also motivated by the very interesting comments of Anderson (1977). Anderson suggests that I was mistaken in claiming that the operative factor in formulating rules like the adjectival passive rule was the deep grammatical relation of the surface subject. Rather, he argues, it is thematic relations like "theme", "agent", "goal",

and "source" (see Gruber (1965) and Jackendoff (1972)) which are crucial¹. Assuming Anderson to be correct, an obvious modification of Bresnan's system suggests itself, which would permit the distinctions of my earlier paper to be captured. Let us suppose that the functional structure in lexical entries is a specification of which thematic relations should be assigned to the elements mentioned in the subcategorization frame. Then we may distinguish two types of lexical rules: those that make reference to thematic relations and those that do not. The former would correspond to rules that my earlier paper called lexical, and the latter to those that I called transformations. This is the extension of Bresnan's framework that I wish to propose. I will illustrate by formulating the two passive rules and the dative rule and applying them to a fragment of the lexicon of English.

My formalism is based on the assumption that the grammatical relations are given language-wide definitions in structural terms (at least in English) along the lines indicated by Bresnan, and that a verb's subcategorization frame merely indicates which relations it has, and what grammatical categories those relations are assigned to. (Thus, I differ from Bresnan in this respect, for she assumed that grammatical relations would be limited to NP's). I will adopt the following abbreviations: "SS" = (surface) subject; SO = (surface) object; "SO2" = (surface) second object; "1" = theme; "2" = agent; "3" = goal; "4" = complement. The rule forming verbal passive participles from the corresponding active lexical entries can now be formulated² quite simply as SS+SO. This is to be interpreted as follows: eliminate "SS" wherever it appears in the entry for the active verb (eliminating also any assignment it may have to a thematic relation) and change all occurrences of "SO" to "SS"³. The adjectival passive rule will differ from this in that it has an additional condition on it: if SO=1, then SS+SO. This condition insures that the SO is "local", in the sense that it bears a thematic relation to the verb. The dative rule⁴ also has a "localness" condition: if SO2=1, then SO+SO2. Let me illustrate these rules with a simple example, namely the verb *sell*. The basic lexical entry I posit for this verb includes the following information: SS=NP, SO=NP, SO2=NP; SS=2, SO=3, SO2=1. This, I claim, is among the information that must be included in a representation of *sell* in such uses as *They sold John two cars*. Applying the verbal passive rule to this entry, we get the following: SS=NP, SO2=NP; SS=3, SO2=1. This verb appears in examples like *John was sold two cars*. Since the original entry for *sell* did not meet the condition SO=1, the adjectival passive rule is not applicable; correspondingly, forms like **John was unsold two cars* are impossible. The condition for application of dative, SO2=1 is met, so we can derive an entry in which SS=NP, SO=NP; SS=2, SO=1. This corresponds to examples like *They sold two cars*. Notice that this last entry does satisfy the condition on the adjectival passive rule, so we can derive the following entry for an adjectival passive participle for *sell*: SS=NP; SS=1. This corresponds to examples like *Two cars were unsold*.

Let us now turn to some more complex examples. Specifically, I now want to look at several different verbs which share the same strict subcategorization frame, namely, SS=NP, SO=NP, SO2=VP. The verbs in question differ from one another along two dimensions, namely, the assignment of thematic relations, and control properties. What I mean by this latter phrase is quite simple: the understood subject of the VP in the SO2 position will be the SS in some cases and the SO in others. I will represent this in the functional structure by assigning a thematic relation not simply to SO2, but to SO2(SS) or SO2(SO), depending on the control properties⁵. My assignments of thematic relations are intended to reflect certain intuitions about the semantic roles of the various elements, but I cannot, in general, provide empirical arguments

for my assignments, other than the fact that they give me the right results. I do have an operational criterion for deciding whether to call the SO a 1 or a 3: when the verb in question could appear in a double object construction (i.e., immediately followed by two NP's), I called the SO a 3; otherwise, I called it a 1. Thus, in what follows, the assignments are correlated with the fact that *promise* and *tell* have double object forms (*I promised/told him nothing*), but *persuade* and *believe* do not (**I persuaded/believed him nothing*).

Consider first *persuade*. The functional structure for this verb in examples like *They persuaded John to leave* would be SS=2, SO=1, SO2(SO)=4. The passive rule yields an entry whose functional structure is SS=1, SO2(SS)=4. Since SO=1 in the original entry, this passive may be either verbal or adjectival. Hence, we can get both *John was persuaded to leave* and *John seemed persuaded to leave*. On the other hand, the condition for application of dative is not met, and, accordingly, we cannot get **They persuaded to leave*. Transformational studies going back to Rosenbaum (1967) have pointed out numerous differences between the behavior of *persuade* and that of *believe*. The standard analysis of these differences has involved the claim that the surface object of *believe* was raised from the subject position of the complement. The system proposed here can mimic that analysis by assigning to *believe* a functional structure in which the SO bears no thematic relation⁶: SS=2, SO2(SO)=1. These are the assignments for examples like *I believe John to be at home*. The verbal passive rule will apply, yielding the functional structure SO2(SS)=1, for examples like *John is believed to be at home*. Since neither the condition on the adjectival passive rule nor that on the dative rule is met, we can predict the non-occurrence of examples like **John seems believed to be at home* and **I believe to be at home*. The next verb I wish to consider is *tell*, which standard transformational accounts would not distinguish in any relevant way from *persuade*. For reasons noted above, I assign *tell* the functional structure SS=2, SO=3, SO2(SO)=1, as in examples like *We told John to bring the beer*. Applying the verbal passive rule we get SS=3, SO2(SS)=1, covering examples like *John was told to bring the beer*. The condition on the adjectival passive rule is not satisfied, so we cannot derive **John seemed told to bring the beer*. Notice now that the condition for applying the dative rule is met. Applying the rule results in the following functional structure: SS=2, SO(1)=1; this structure is ill-formed, since there is no controller. Accordingly, examples like **We told to bring the beer* are impossible. Finally, consider *promise* in examples like *I promised John to mow the lawn*. *Promise* is exactly like *tell*, except that the controller is the subject, not the object, i.e., the functional structure is SS=2, SO=3, SO2(SS)=1. If we try to apply either passive rule, we will get the following functional structure: SS=3, SO2(1)=1. This is ill-formed for the same reason that the dative of *tell* was, namely, lack of a controller. The corresponding examples are also impossible: **John was promised to mow the lawn* or **John seemed promised to mow the lawn*. Dative, however, can apply, yielding an entry whose functional structure is SS=2, SO(SS)=1. This corresponds to examples like *I promised to mow the lawn*.

I hope that this fragment of the lexicon suffices to show that my proposed modification of Bresnan's system permits an elegant and natural account of a number of syntactic distinctions, including some which have not been discussed in the literature, to my knowledge. One nice feature that I would like to emphasize is that my proposals provide a rather straightforward account of Visser's (1973: 2118) observation: "A passive transform is only possible when the complement relates to the immediately preceding (pro)noun." In my terminology, passive will be impossible when the active has a complement controlled by the SS, as in the case of *promise*,

for passivization will always lead to an uncontrolled complement. Thus, to take another standard example of Visser's generalization, we can account for the distinction between *strike* and *regard* much as we accounted for the difference between *promise* and *tell*. Both will have the following subcategorization frame: SS=NP, SO=NP, SO2=AP. Their functional structures will include the assignments SS=2 and SO=1; they will differ in that *regard* will have SO2(SO)=4, while *strike* has SO2(SS)=4. These assignments are for examples like *John regards/strikes Mary as pompous*. If we apply passive to *regard* we get SS=1, SO2(SS)=4, as in *Mary is regarded as pompous*. Applying passive to *strike* we get SS=1, SO2(1)=4, which is ill-formed, as is **Mary is struck as pompous*. Notice, incidentally, that this example illustrates that, in the system I advocate here, constituents other than VP's can serve as predicates and be subject to control.

This concludes my suggestions for modifying Bresnan's framework. I hope I have succeeded in indicating how a grammar which makes extensive use of the lexicon in place of syntactic transformations can handle an array of syntactic facts in a satisfying manner. Next, I wish to argue that a system of the sort outlined here can be effectively constrained in reasonable and interesting ways. Intuitively, it seems quite plausible that such a system would be easy to constrain, for by drastically reducing the role of transformations, it opens the way for reductions in the power of transformations. A number of candidate constraints on transformations come to mind. For example, within Bresnan's framework one might plausibly argue that no transformation can create new grammatical relations (e.g., there will be no "subject-creating" transformations like passive or raising to subject), or that no transformation can change the words in the sentence morphologically (e.g., there will be no nominalization, agreement, or case-marking transformations--cf. Brame (1978)). Various ways in which lexical rules might be constrained also come to mind; most immediately, it seems to me that many of the "laws" of relational grammar proposed by Postal and Perlmutter in recent years could be translated straightforwardly into the kind of framework discussed here. In this paper, however, I would like to consider the consequences of a constraint on transformations modeled on the Freezing Principle of Culicover and Wexler (1977). My proposal depends on distinguishing two classes of transformations: root transformations (Emonds (1976)), and what I will call unbounded rules. Root transformations are rules like English subject-auxiliary inversion in questions, which apply only to main clauses; unbounded rules are transformations (e.g., *wh*-movement) which involve a crucial variable, i.e., they move something over a variable or they delete something under identity with something on the other side of a variable⁷ (see the contributions by Chomsky, Bach, Bresnan, and Partee in Culicover, *et al* (1977) for discussion of whether unbounded rules are truly unbounded). The constraint I wish to propose, which I will call the **interaction constraint** is the following: once a rule of one of these classes has applied to a given structure, no further rule of the same type may apply to that structure. More specifically, when a transformation applies, the smallest constituent containing all of the affected elements becomes **frozen**, in the sense that no further transformations of the same type may analyze it. This means, in effect, that there will be no interactions among root transformations, nor among unbounded transformations (though a root transformation may interact with an unbounded rule, as in the case of English *wh*-questions). I believe that there are several desirable consequences of prohibiting such interactions.

First of all, let me mention a somewhat conjectural reason for advocating the interaction constraint. As noted above, a very similar proposal emerged from the learnability studies of Wexler, Culicover, and Hamburger; they were able to prove

that a class of grammars in which nodes were frozen under similar conditions was learnable by a fairly simple learning device. Hence, it seems plausible to conjecture that the interaction constraint might be useful in devising a learnability proof for some version of Bresnan's theory. In any event, it seems that the interaction constraint would make the language-learner's task easier by limiting the extent to which surface structures could deviate from base forms (see Coker & Crain (in preparation)).

Second, there is empirical support for the interaction constraint. Emonds (1972: 38-40) shows that only one root preposing transformation can apply per sentence. Since the smallest structure containing initial position in a root sentence is the whole sentence, Emonds's observation is an immediate consequence of the interaction constraint. Similarly, many of the ways in which unbounded transformations are prohibited from interacting are familiar. For example, the fact that elements in relative clauses are inaccessible to unbounded transformations has been extensively discussed in the literature (e.g., Ross (1967), Chomsky (1973), to cite only two accounts). This fact follows from the interaction constraint, since an unbounded transformation is involved in the formation of relative clauses. Hence, examples like **Who do you know a man who saw?* or **John is taller than I know a man who is* are excluded by the interaction constraint. The fact that comparative clauses and embedded questions are also "islands" has been less widely discussed in the literature, but is also a consequence of the interaction constraint. Thus, such examples as **Who is John louder than Mary persuaded to be?* or **Who does John wonder when Bill will see?* are excluded because they involve *wh*-movement extracting material from clauses in which *wh*-movement or comparative deletion has taken place. Likewise, comparative clauses are impervious to further applications of comparative deletion: **John was kind to more people than he liked Bill more than I liked* (where this would mean, if grammatical, that the number of people John was kind to exceeded the number of people liked better by Bill than by me). In short, the interaction constraint seems to make the right predictions about a substantial array of data.

Finally, I would like to suggest that the interaction constraint serves not only to restrict the class of grammars made available by linguistic theory, but also to limit the class of languages generable by the available grammars (see Wasow (in press a) for discussion of this distinction). I will not attempt any formal demonstration of this conclusion here, but will sketch briefly why I believe it to be the case. Peters and Ritchie (1973) prove that the language generated by a transformational grammar is recursive if it is possible, on the basis of a surface string, to effectively compute a maximum size of a deep structure from which that string could be derived. The interaction constraint, together with the standard condition on recoverability of deletions (see Peters and Ritchie (1973)), limit the extent to which deletions may shrink a structure. To show why this is the case, it will be useful to invent some terminology: let us call A a parent of B if B can be derived from A by a single application of one transformation. A parent's parent will be called a grandparent, and so on. Now consider a string of length n . Because of the recoverability condition, its parent cannot be longer than $2n$ (measuring length in terms of number of terminal symbols). Likewise, its grandparent cannot be longer than $4n$. However, if the grandparent were the full $4n$ long, then the parent would be frozen by the interaction constraint, and the original string would be underivable. In fact, each (length n) half of the parent must have a parent of length no more than $2n-1$, if we are to avoid blocking the derivation by the interaction constraint. Thus, the maximum size of a grandparent is $4n-2$. By similar reasoning it is not hard to see that the maximum size of any ancestor $m+1$ generations removed is $2^m(2n-m)$. Since this number becomes zero when $m=2n$, there is an

effective upper bound on the size of any ancestor. Hence, the interaction constraint, together with the standard condition on recoverability of deletions, limits the class of languages which can be generated to a subclass of the recursive sets⁸. This provides yet another point of convergence with computational concerns, since, as noted above, a language must be recursive in order to be effectively processed.

I have sketched a version of transformational grammar which seems to hold considerable promise. There are a number of problems with this approach which I am aware of and undoubtedly many more I am blissfully ignorant of. What I have presented here was intended, more than anything else, as an indication of a program of research, and I have hence felt free to ignore many important issues. The primary point I wish to make is that the study of language appears to have progressed to a point where the concerns of the transformationalist and the concerns of the computational linguist need not conflict, and indeed may be addressed by a single theory.

* I wish to express my gratitude to Adrian Akmajian, Joan Bresnan, and especially Ron Kaplan for very stimulating discussions of some of the material in this paper. They are, of course, absolved of any responsibility for its shortcomings. I am also very grateful to the Xerox Corporation for making its resources, human and electronic, available to me in the preparation of this paper. Some of the research reported on here was begun under a Summer Stipend from the National Endowment for the Humanities.

Footnotes

1. No rigorous definition of these notions has ever been offered in the literature, and certain problems with the way they have been used have been pointed out (e.g., Hirst and Brame (1976)). I do not wish to commit myself to all of the claims which have been made in the literature about these notions, and my notation below is intended to reflect this. I do, however, believe that those who have discussed thematic relations are onto something important.
2. Obviously, there is more to forming passives than this; for example, I ignore morphology.
3. Those familiar with Postal and Perlmutter's version of relational grammar will recognize the resemblance of last sentence to the Relational Annihilation Law. Notice, by the way, that my passive rules say nothing about the *by* phrase. I am assuming, with Bresnan (in press), that there is an independent rule assigning agent status to the objects of some *by* phrases. This rule would operate not only in passives, but also in examples like *The symphony was by Beethoven*.
4. Notice that I am formulating the dative rule "backwards", that is, with the double object construction as the input. My rule says nothing about the prepositions *to* and *for* because I assume that the functional role of their objects will be covered by separate rules, as is the case with *by*. Examples like *John's call was to Mary* and *This present is for you* lend credence to my assumption.
5. This is to be understood as saying that the SO2 will be treated as a predicate, with its own assignments of thematic relations, and with the element in parentheses treated as if it were the SS of that predicate.
6. Jane Robinson has suggested to me that it might be more appropriate semantically to treat the subject of believe as a 3. This would be perfectly compatible with my analysis.
7. My treatment here ignores anaphora rules like VP deletion and sluicing. I am assuming that these rules are not transformations, but a separate category of rules, subject to their own unique conditions (see Wasow (in press b) for discussion).
8. As given, my argument does not take into account root transformations or specified deletions (see Wasow (in press a)). It is quite trivial, however, to extend the argument to cover these cases.

References

- Anderson, S. (1977) "Comments on the Paper by Wasow", in Culicover, *et al* (1977).
 Brame, M. (1978) "The Base Hypothesis and the Spelling Prohibition". *Linguistic Analysis* 4.1.
 Bresnan, J. (1976) "On the Form and Functioning of Transformations". *Linguistic Inquiry* 7.1.

Bresnan, J. (in press) "A Realistic Transformational Grammar", in M. Halle, J. Bresnan, and G. Miller (eds), *Linguistic Theory and Psychological Reality*. MIT Press, Cambridge, Massachusetts.

Chomsky, N. (1965) *Aspects of the Theory of Syntax*. MIT Press, Cambridge, Massachusetts.

Chomsky, N. (1973) "Conditions on Transformations", in S. Anderson and P. Kiparsky (eds), *A Festschrift for Morris Halle*. Holt, Rinehart, and Winston, New York.

Coker, P. and S. Crair (in preparation) "Linguistic Processing: The Grammatical Basis of Sentence Interpretation". Claremont Graduate School, Claremont, California.

Culicover, P. and K. Wexler (1977) "Some Syntactic Implications of a Theory of Language Learnability", in Culicover, *et al* (1977).

Culicover, P., T. Wasow, and A. Akmajian, eds (1977) *Formal Syntax*. Academic Press, New York.

Emonds, J. (1972) "A Reformulation of Certain Syntactic Transformations", in S. Peters (ed), *Goals of Linguistic Theory*. Prentice-Hall, Englewood Cliffs, N.J.

Emonds, J. (1976) *A Transformational Approach to English Syntax: Root, Structure-Preserving and Local Transformations*. Academic Press, New York.

Fodor, J. A., T. Bever, and M. Garrett (1974) *The Psychology of Language*. McGraw-Hill, New York.

Gruber, J. (1965) *Studies in Lexical Relations*. MIT dissertation.

Hust, J. and M. Brame (1976) "Jackendoff on Interpretive Semantics". *Linguistic Analysis* 2.3.

Jackendoff, R. (1972) *Semantic Interpretation in Generative Grammar*. MIT Press, Cambridge, Massachusetts.

Osherson, D. and T. Wasow (1976) "Task Specificity and Species Specificity in the Study of Language: A Methodological Note". *Cognition* 4.

Peters, S. and R. Ritchie (1973) "On the Generative Power of Transformational Grammars". *Information Sciences* 6.

Rosenbaum, P. (1967) *The Grammar of English Predicate Complement Constructions*. MIT Press, Cambridge, Massachusetts.

Ross, J. (1967) *Constraints on Variables in Syntax*. MIT dissertation.

Wasow, T. (1977) "Transformations and the Lexicon", in Culicover, *et al* (1977).

Wasow, T. (1978) "Some Thoughts on Mental Representation and Transformational Grammar". Paper delivered at MIT Sloan Foundation Workshop on Mental Representation.

Wasow, T. (in press a) "On Constraining the Class of Transformational Languages". *Synthese*.

Wasow, T. (in press b) *Anaphora in Generative Grammar*. Story-Scientia, Ghent.

Woods, W. (1973) "An Experimental Parsing System for Transition Network Grammar", in R. Rustin (ed), *Natural Language Processing*. Algorithmics Press, New York.