

Grammatical Competence and Parsing Performance

Bradley L. Pritchett

(Carnegie Mellon University)

Chicago: The University of Chicago Press, 1992, xii + 192 pp.
Hardbound, ISBN 0-226-68441-5, \$55.00;
Paperbound, ISBN 0-226-68442-3, \$19.95

Reviewed by
Gerry Altmann
University of Sussex

Much psycholinguistic experimentation in recent years has been concerned with the nature of the human processes that give rise to the difficulties associated with sentences such as *The author read the review smiled* and *Whilst he was writing the review flowed onto the page*. The problem, of course, is that language is ambiguous, and these examples illustrate that in the face of local ambiguity, the human processor chooses to pursue just one analysis, even if that analysis subsequently turns out to be false. The fact that there exist such *consistent* preferences (consistent across examples and across different people) suggests that the decision about which analysis to pursue is not arbitrary. Studying such examples may thus tell us something about the nature of human sentence processing, and even about the relationship between grammar and process, a topic of interest to psycholinguists, linguists, and computational linguists.

Bradley Pritchett's book takes the reader through a wide range of examples of so-called "garden path" sentences, and reminds us that there are two questions that any theory of sentence processing must address: What is the cause of the various parsing preferences? And what determines the ease of re-analysis? The second question is posed in the face of materials such as those above that require conscious effort, and others that require no *conscious* re-analysis but do nonetheless require some syntactic restructuring (compare *He knew the book well* and *He knew the book was interesting*, which differ in terms of the phrasal position to which *the book* is assigned). In the brief summary that follows, I shall gloss some of the main claims purely for the sake of brevity.

Two principles are proposed, one to explain why certain analyses are chosen over others, and another to explain why certain forms of syntactic reanalysis are problematic. *Theta attachment* proposes, in effect, that at any moment the processor attempts to map as many thematic roles onto as many arguments as possible. Pritchett is clear, however, that the semantic content of the roles is largely irrelevant: what matters is the number of roles and their configurational position. He then proposes the *theta reanalysis constraint*, which states, again in effect, that the processor cannot reassign a role to an argument if the new role comes from outside the *theta domain* of the original role. Over the course of a very rigorous examination of how these principles account for a large variety of parsing preferences and associated reanalysis phenomena, the two principles evolve into *generalized theta attachment* ("every principle of the syntax attempts to be maximally satisfied at every point during processing") and the *on-line locality constraint* ("the target position . . . assumed by a constituent must be *governed* or *dominated* by its source position . . . , otherwise attachment is impossible"). Both principles are cast in the framework of Government-Binding theory.

The arguments are persuasive, and the work deserves to command as much interest as previous work by John Kimball, Lyn Frazier, and others. Anyone with an interest in sentence processing and syntactic ambiguity and the relationship between grammatical knowledge (competence) and processing (performance) should buy this book. I myself think it is one of the better books on my shelves. But that is not to say that I believe that Pritchett's account is exactly right. In particular, there are a number of empirical issues that the account raises that deserve exploration, as well as meta-theoretic issues that may themselves arouse some controversy.

A definite plus with Pritchett's work is that his theory is sufficiently detailed to allow predictions to be made about precisely when processing difficulty will be experienced. With regard to empirical validation, this is an important feature of the project. Whether or not every reader will agree with Pritchett's analysis of which examples are problematic and which are not is another matter. With regard to the consistency of the theory, it is absolutely vital that a consensus is achieved; otherwise the theory is severely compromised. This is one area where future psycholinguistic experimentation may be particularly useful.

I suspect, however, that Pritchett does not view such experimentation favorably. In one of his footnotes (n. 59), he comments on the current debate in the literature between theorists who argue for a purely syntactic basis for parsing preferences, and those (such as myself) who argue for a pragmatic/contextual basis. It is true that much of the debate has centered on how to interpret the particular methodologies involved, and Pritchett dismisses the possibility that there are nonsyntactic elements at work by saying that the debate "too often degenerates into methodological squabbling devoid of any real theoretical interest. Rather than participate explicitly, I will simply continue to assume, based on the sort of evidence reviewed here in chapters 1 and 2, that parsing strategies are basically syntactic." Well, that's certainly one way to cope with the debate—the evidence reviewed in the early chapters does *not* include any of the evidence that favors the contextual approach to parsing preferences.

The distinction between a syntax-only explanation of the preferences exhibited by the sentence processor and a context-based explanation is extremely important. The fact that controversy surrounds whether the data support one or another position is immaterial to whether the theoretical issue is an important one. I am, naturally, biased. Moreover, I accept that only recently has new evidence been provided that really can tip the scales in favor of one account or another.

I raise the issue here because it highlights an important aspect of Pritchett's goal, namely to understand better the relationship between grammar and process. At first glance, theta attachment, as I glossed it earlier, looks as if it adds an air of "ecological validity" to the literature on parsing preferences. It makes sense, in attempting to establish the traditional who-did-what-to-whom, to discharge as many roles as soon as possible. Consequently, theta attachment appears to achieve some degree of explanatory adequacy. Even in its most generalized form, which makes appeal to the maximal satisfaction of syntactic principles, it provides some explanatory satisfaction. From a personal viewpoint, thinking in terms of thematic roles is useful because it allows exactly the kind of interaction that Pritchett rejects, namely contextual.

For instance, in an increasingly lively debate, various colleagues (Mark Steedman, Stephen Crain, and Alan Garnham) and I have argued that in appropriate contexts, the difficulty associated with *The professor told the student that he was worried about to work harder* can be eliminated. In the absence of an explicit context, the *that*-clause is preferentially interpreted as a complement clause instead of a relative clause, and this leads to a garden path on *to work*. Our claim is that if there are two or more students in the context, one of whom the professor is worried about, then the processor will fail

to identify a unique referent for *the student*, and will therefore treat the *that*-clause as a nominal modifier. This would fit in well with the view that the processor attempts to discharge roles as soon as possible, because the recipients of roles are not referring expressions themselves, but their *referents*. If you can't identify the referent, it makes sense to do everything you can to identify it so that you can then assign a role to it. Pritchett's position is thus compatible with the notion that referential context can influence initial parsing decisions.

Pritchett insists that in fact his account is *not* to do with thematic roles, but rather with their configurational (cf. "syntactic") properties. Whereas I seek an explanation for parsing preferences in terms of what syntax is *for*, Pritchett seeks an explanation in terms of syntax *per se*. I admit that he does so convincingly. I would like to have seen some more explanation, however, of *why* reanalysis requires that the new position of a constituent must be governed or dominated by the original position. What principle of *processing* forces this? Pritchett's answer is that what is required is not a principle of processing, but a principle of *grammar*, and that this (i.e., the government/dominance fact) is it. But why?

In summary, Pritchett's work is topical and radical, and will, I suspect, be influential. I do not think he's right to place the burden of explanation on grammar *per se*. First, because we are not machines for building phrase markers, there is a view that grammar serves a purpose, and that the nature of the purpose constrains the operations of the machine. Second, because I could, in principle, build a machine that *could* reanalyze constituents in violation of his reanalysis constraint (the on-line locality constraint). That fact alone tells me that his account is a *description* of events, and not an *explanation* of them. This said, the description itself *is* an achievement, and I wish *I'd* thought of it. . .

Gerry Altmann, a lecturer in Experimental Psychology at the University of Sussex, is a managing editor of the journal *Language and Cognitive Processes*. His research on syntactic ambiguity resolution focuses on the influence of context on initial parsing decisions, and he has published a variety of psycholinguistic studies on this topic. Altmann's address is: Cognitive and Computing Sciences, University of Sussex, Brighton BN1 9QH, U.K.; e-mail: gerry@epvax.sussex.ac.uk.