Prosodic Disambiguation In Silent Reading

Prosodic Disambiguation In Silent Reading

Janet Dean Fodor

Graduate Center, City University of New York

1. Implicit prosody

The research I present in this paper is the joint work of CUNY's sentence processing research group. Together, we have been investigating the *Implicit Prosody Hypothesis* (Fodor 1998, Quinn et al. 2000; see also the *Prosodic Constraint on Reanalysis* of Bader 1998).

(1) The Implicit Prosody Hypothesis (IPH): In silent reading, a default prosodic contour is projected onto the stimulus, and it may influence syntactic ambiguity resolution. Other things being equal, the parser favors the syntactic analysis associated with the most natural (default) prosodic contour for the construction.

We believe this is the key to some puzzles that have arisen in sentence processing theory.

As a hearer or reader processes a sentence, each word has to be attached into a representation of its syntactic structure. When there is structural ambiguity, a choice exists between different attachment sites for a word. Considerable consistency has been observed in the choices that are made, both across sentences and across speakers. Though non-structural factors may also influence attachment, there are some general structural trends. We have summarized these under the heading *Minimal Everything*: the parsing routines prefer a tree structure with the fewest nodes, the fewest chains, the least structural distance between adjacent words, and minimal revision in recovering from a garden path. But some inconsistencies in attachment preferences have been noted which are still in need of a definitive explanation.

Although these discrepant findings are from experiments on silent reading, we believe they are attributable to the prosodic properties of sentences. Perhaps for practical reasons, prosody has until recently been largely neglected in sentence comprehension studies (with a few notable exceptions); input has mostly been visual rather than auditory. Nowadays the tide is turning, with an increasing number of studies focussing on prosodic phenomena (see the special volumes of *Language and Cognitive Processes* 1996, and *Journal of Psycholinguistic Research* 1996). Of course, prosody is especially easy to overlook as a biasing factor in silent reading. And even once it has been thought of, its role is hard to prove. Prosodic patterns that are mentally projected by perceivers cannot be directly observed or measured, so their existence could only be established indirectly, by showing that they offer the most reasonable explanation for parsing phenomena we do observe. In this respect, experimental studies of implicit prosody are akin to experimental studies in the 1960's on the psychological reality of phrase structure.

The claim is that prosody is an important factor in parsing even when it is not really there. This may sound improbable but it has at least some initial plausibility. First, studies of real (overt) prosody in spontaneous speech or reading aloud have shown it to be largely congruent with syntactic structure but also sensitive to constituent length (Selkirk 2000), and this is true also of the parsing anomalies that need to be explained. Second, prosodic patterns are partly universal but partly language-specific, as are the ambiguity resolution preferences in parsing that need explaining. Finally, there is considerable evidence of phonological encoding in silent reading, even though prosody is not usually studied. With the help of Rayner and Pollatsek (1989) we found two previous studies that may relate to implicit prosody. One is an experiment by Kosslyn and Matt (1977) showing that silent reading of a passage is faster if readers are told it was written by a person they just heard speaking fast, than if they believe it was written by a person they just heard speaking slowly. This implies at least some kind of rhythmic or temporal encoding in silent reading. The other study of interest, by Slowiaczek and Clifton (1980), also did not specifically investigate silent prosody but invoked it as one possible explanation for the fact that when subjects spoke nonsense syllables (*colacolacola...*) while reading silently they were at a greater disadvantage (compared to speaking nonsense while listening to spoken input) for tasks requiring integration of content across phrases or sentences than in tasks involving individual word meanings.

Slowiaczek and Clifton regarded the function of "subvocal speech" as still a matter for speculation, but they didn't consider its existence to be in doubt. Their paper begins: "The subjective experience of hearing a voice inside one's head while reading seems nearly universal" (p.573). Rayner and Pollatsek write even more strongly "...we all hear an inner voice pronouncing the words that our eyes are traversing as we read" (p.443). On internally generated prosody in particular, Rayner and Pollatsek cite Brown (1970), claiming that "When you read a letter from someone you know very well, such as your mother, you often can hear her accent, or stress, or intonation pattern" (p.216). (In Roger Brown's case, it was Roman Jakobson's voice that was conjured up by reading his prose.) Rayner and Pollatsek add: "Also, when you read text such as this book, you do not hear your voice in a monotone (unless perhaps you always speak in a monotone). Rather, you

are aware of providing stress and intonation patterns to the words" (p.216). On the other hand, we have encountered a fair number of people who claim they do NOT hear their loved ones' voices in their letters, or hear their own voices when they read textbooks. These sceptics may be wrong, of course; introspection is a notoriously unreliable guide to what is going on in the mind.

If it is true that implicit prosody underlies a variety of ambiguity resolution phenomena, then obviously it is important to document its existence more convincingly. We have established a four-step procedure to test the IPH, which looks for parallels between (presumed) implicit prosody in reading silently, and explicit prosody in speech and reading aloud.

(2) To test the IPH:
[1] Find a factor F which can be manipulated in an experiment, and which measurably affects the OVERT prosody of a sentence.
[2] Show that the overt prosodic difference caused by F measurably influences an ambiguity resolution preference in parsing.
[3] Show (or claim?) that F does not affect parsing DIRECTLY.
[4] Include F in a silent reading task. Is ambiguity resolution affected by F as it is the listening task?

The idea is to see whether implicit prosody reliably mimics the parsing consequences of explicit prosody. We make the working assumption that if a prosodic pattern is projected onto a sentence in silent reading, it will be identical to the overt prosody for that sentence in a comparable context (i.e., same illocutionary force, focus structure, etc.). The trickiest part of the program is step [3], since it obliges us to consider all other possible ways in which F might bring about the observed results without ascribing a causal role to the prosody that F induces. But this is a familiar kind of challenge in psychology. Informally, it amounts to: we can see no other way to account for the data. More grandly it is known in the philosophy of science as *abduction* (Pierce 1931-58) or *inference to the best explanation* (Harman 1965).

In our work, the importance of step [3] has led us to rely on prosodic phrasing rather than prosodic focus effects. Focus has a semantic as well as a phonological aspect. Suppose it were discovered that a pronoun is preferentially interpreted as referring back to a focussed antecedent, in reading silently as well as in listening to speech. This might be due to the focal accent on the antecedent. But it might instead be due to the antecedent's prominent role in the information structure of the discourse (or even just to its greater physical salience). The strongest tests of the IPH thus employ a factor F which carries no semantic or pragmatic freight. In the experiments reported below, F is either constituent length, or language-specific prosodic rules for the placement of intermediate phrase boundaries, neither of which is confounded with non-prosodic factors.

2. An exception to the Reanalysis as Last Resort principle (RALR)

RALR (Fodor and Frazier 1980) is a sensible least-effort parsing principle, one aspect of the more general economy principle of minimizing revision. It says: Don't change the structure that has been built so far for a sentence unless subsequent input is incompatible with it. It was surprising, therefore, to find an apparent exception to RALR in Japanese: a first-pass left-branching structure which appears to flip spontaneously into right-branching as reading proceeds (Inoue and Fodor, 1995). This is illustrated in (3).

(3) kyokutanni shinsetsuna gakusei-no imooto
extremely kind student's sister

Intuitively, the structure that a reader builds for (3) starts out as indicated in (4a) (shown with English words, which have the same order as in Japanese). So by RALR we would expect the analysis to continue as in (4b). But in fact some native speakers judge that it becomes (4c), in which the originally formed phrase *kyokutanni shinsetsuna gakusei-no* has been broken up.

(4) a. [extremely kind student]'s
b. [[extremely kind student]'s [sister]]
c. [[extremely kind] [student's sister]]
d. [kind student]'s...
e. [[kind student]'s [sister]]
f. [[kind] [student's sister]]

However, with a simple modifier such as *shinsetsuna* in place of the complex *kyokutanni shinsetsuna*, the original left-branching structure (4d) does continue as (4e), not (4f), in accord with RALR. Thus the length or 'heaviness' of the modifier appears to affect its preferred position in the tree, and powerfully enough to trigger syntactic reanalysis. A similar phenomenon occurs in English; see Fodor (1998) on *the (recently) divorced bishop's daughter*.

Experimental confirmation of this effect in Japanese was provided by Hirose (1999). In a self-paced silent reading study, sentences were tested in 4 versions as in (5). Examples in English translation are given in (6). Note that the initial adverb in the short-modifier versions (i) and (ii) is a sentence-level adverbial included to balance the sentence length of the long-modifier versions (iii) and (iv).

(5) (i) [Adv] [Adj] $N1_{[GEN]}$ $N2_{[ACC]}$ $NP_{[NOM]}$ V (Adj compatible with N1, not N2)
(ii) (Adj compatible with N1 and N2)

(iii) [Adv Adj] $N1_{[GEN]}$ $N2_{[ACC]}$ $NP_{[NOM]}$ V (Adj compatible with N1, not N2)
(iv) (Adj compatible with N1 and N2)

(6) {Two weeks ago / Somewhat} {dimly-lit / unsophisticated} bar's hostess Satoru insulted.

No significant reading time penalty was found for version (i) compared with version (ii), i.e., the one-word modifier was indifferent to whether or not it made sense with N2. Evidently, it remained associated with N1. For the two-word modifier, N2 took longer to

read in version (iii) where it was incompatible with the modifier, than in version (iv), showing that the initial structure was sometimes readjusted so that N2 was the modified head. A related result for relative clause modifiers (also prenominal) in Japanese is reported by Kamide at al. (1998): the tendency to modify N2 was positively correlated with relative clause length.

One explanation for these effects of modifier length is that they arise at the syntax-phonology interface. The optimal length of a prosodic phrase (a major phrase in Selkirk's terms) is two minor (or accentual) phrases, not one or three (see the BinMin and BinMax constraints of Selkirk 2000). For a sequence of three minor phrases (such as Adj - N1 - N2) there is no way to avoid at least one violation of optimal length, so the grouping (4e) favored by RALR is as good as any other. But with four minor phrases (as in Adv - Adj - N1 - N2) the division 3 + 1 would be a double violation. A shift is thus motivated from (4b) to the violation-free division 2 + 2 as in (4c). This length-sensitive shift in the prosodic pattern was demonstrated for overt prosody in reading aloud by Kubozono (1993), who found that a left-branching sequence of four minor phrases "is prosodically neutralized with [a] symmetrically-branching structure" in which the four elements "are grouped into two subgroups of two" (p.221). Thus, Hirose's results suggest that implicit prosody in silent reading is subject to the same optimal phrase-length constraints as overt prosody is, and that these influence the resolution of syntactic ambiguities when overt prosody is absent.

Note that this explanation assumes that in reading ambiguous sentences, the syntactic structure is brought into line with the preferred prosodic structure. There is a general preference for prosody to be congruent with syntax, though there are known exceptions in which optimal length, avoidance of recursion, and other considerations outweigh the alignment constraints. To the extent that it does occur, syntax/prosody congruence is achieved in different ways in speaking, listening, and reading. In speaking, a meaning is chosen, its syntactic form is assigned, and the prosody is produced in conformity to the syntactic structure. In listening, the prosody is in the signal, and the syntactic structure computed by the hearer may conform to it when the syntax is not decisive, i.e., where there is syntactic ambiguity (see Schafer et al. 1996 and references there). In reading, both aloud and silently, syntactic ambiguity resolution may conform to prosody also, though it must be prosody projected onto the word string by the reader on the basis of a partial syntactic analysis.

Factor F in the case of the Japanese modifiers is the length or complexity of the modifier phrase. Kubozono's data show that F affects overt prosody; this is step 1. Step [2] would be a demonstration that attachment preferences in a listening experiment are sensitive to this overt prosody, showing increased N2 attachment when there is a break between an adjective and the adjacent noun. This was not established experimentally but was presumed with some confidence on the basis of informal judgments and experimental results for similar constructions. Step 3 is that there is no obvious relevant difference between long and short modifiers other than their effect on prosodic phrasing. Step 4 is Hirose's finding that F affects attachment preference in silent reading.

3. Left edge of Japanese relative clauses

Another parsing puzzle that needed to be solved concerns the preferred resolution of a positional ambiguity of the left edge of a relative clause (RC) in Japanese. Japanese RCs precede their head nouns, and have no overt relative pronoun or complementizer. So there is no unambiguous lexical/syntactic marking of the left edge of an RC. Case marking may suggest an argument structure discontinuity but this is not reliable. In speech, the beginning of an RC is typically accompanied by a prosodic break, but in written Japanese it is not normally signaled at all. The right edge of the RC is defined by the appearance of a noun (the head that the RC modifies) immediately after the verb. At this point the parser must decide how many of the constituents preceding the verb are in the RC. Those which are not in the RC will be in the matrix clause whose completion has yet to be encountered. In many cases, the first-pass parse will have structured all the constituents up to the verb into a single clause, which now needs to be broken up. Where to make the break is very often ambiguous.

Mazuka and Itoh (1995) gave three examples of this ambiguity, and proposed that the preferred dividing point is between the subject and the verb phrase of the clause created during the first-pass parse. This is plausible, since this is the position of the major syntactic break in the clause, a natural splitting point; the revision routines might as well take advantage of the work of the first-pass parser. However, Hirose (1999) asked native speakers to rate the difficulty of these sentences and found that they did not always align the left edge of the RC with the first-pass subject-predicate division; instead, the preferred analysis appeared to vary with the length of the subject (one word or two words). So she created a set of examples in which subject length was manipulated systematically between a single noun, and a coordination of two nouns. The structure of her examples is outlined in (7). The word string is ambiguous between the structures (7a) and (7b). The subject of the sentence was either short (e.g., the name *Morisita*) or long (e.g., the coordinated names *Hosokawa-to Morisita*, 'Hosokawa and Morisita').

(7) Subj Obj[ACC] Adv Verb] head-Noun[DAT] Adv main-Verb
e.g., Subj new medicine truly trusted friends finally met

a. Subj $_{RC}$[e_i Obj Adv Verb] Noun$_i$ Adv Verb (takes DAT only)
'Subj finally met the friends who truly trusted the new medicine.'

b. Subj Obj $_{RC}$[*pro* e_i Adv Verb] Noun$_i$ Adv Verb (takes ACC and DAT)
'Subj finally showed the new medicine to the friends whom he truly trusted.'

In structure (7a) the original first-pass clause [Subj Obj Adv Verb] has been divided after the subject; Hirose called this subject reanalysis (SR). In (7b) the division falls after the subject and the object; this is subject-object reanalysis (SOR).

SR is what Mazuka and Itoh proposed as the preferred reanalysis. But in a sentence completion task, Hirose found that SR was preferred only for the long subjects (73%). For the short subjects, there was no significant preference (only 55% SR responses, not different from chance). The difference between the two cases was significant ($p<.05$). Why should subject length have this effect? A likely explanation (based on Hirose's discussion though for brevity some details differ here) is as follows. For a long subject, the first-pass syntax and prosody would concur on a break after the subject, and the revision routines would naturally adopt this as the left edge of the RC. But for a short subject, a more natural prosodic break would be after the object (or later), because the short subject is not optimal as a major phrase by itself. The prosodic boundary after the object would fight against the strong syntactic boundary after the subject, and the revision routines would be torn between the two.

Subject length is the factor F here. For step 1, Hirose found in a reading-aloud experiment that subject length affected the overt prosody of these sentences during the first-pass parse (i.e., before the reader had encountered the head noun which would trigger reanalysis to a relative clause). Pause duration, and reset of F0 downstep (i.e., a rise or plateau in an otherwise stepwise-declining fundamental frequency), were taken as indices of a prosodic break. As predicted, they were stronger after the long subject than after the short one. For step 2, it was confirmed in a subsequent listening experiment (Hirose, personal communication) that the location of an overt prosodic break after the subject or object influences the parser's choice of SR or SOR structure as predicted. When the prosodic break reinforced the major syntactic break (i.e., a prosodic break after the subject only), there was a 79% preference for the SR analysis. When the prosody was neutral (a break after both the subject and the object, or after neither), there was only a 60% SR preference. When the prosody conflicted with the syntax (a prosodic break after the object only), there was no bias (52% SR choice). In other words: although the first-pass subject-predicate break was a powerful attractor for the RC boundary, the placement of a prosodic break could increase or decrease this tendency, and when it conflicted with the syntactic break it was strong enough to eliminate the syntactic bias entirely. Steps [1], [2] and [4] are thus accomplished. Step [3] rests on the absence of any obvious explanation of how subject-length could affect ambiguity resolution in this way, other than via its effect on prosody. The length effect is not explained, for example, by any tendency to move the minimum number of words/constituents up into the main clause, as Mazuka and Itoh suggested, nor by a tendency to retain the minimum number in the relative clause. Also, Hirose (2000) used double names such as *Morisita Sinjiroo*, instead of conjoined names as in the original study, to eliminate the syntactic coordination structure as a possible cause.

This left-edge-of-RC ambiguity differs from the Japanese modifier-scope ambiguity discussed in section 2, in that the reanalysis is not itself triggered by a prosodic problem. Rather, reanalysis is syntactically triggered by the head noun signaling the existence of a RC; the role of the prosody is in selecting between possible reanalyses. The fact that dividing the clauses between the first-pass subject and VP is more strongly favored when there is a prosodic break there suggests that the parser's general preference for minimal revision includes prosodic as well as syntactic and semantic revisions: other things being

equal, the reanalysis routines favor a revision that requires the least change in the first-pass (implicit) prosody. Bader (1998) also demonstrated that first-pass implicit prosody can affect reanalysis. His work on German dative/possessive ambiguities showed that processing load is higher when a syntactic revision involves a concomitant revision of implicit prosody than when it does not.[1] By contrast, the phenomena I turn to now involve first-pass parsing; there is no garden path, and no reanalysis. The occurrence of prosodic effects in these cases thus implies that some prosody assignment can take place quite early in the parse, presumably fed by rather shallow syntactic processing, and can then influence other syntactic decisions. However, the length of constituents is part of what determines prosodic phrasing, and the eventual length of a constituent may not be reliably assessable when its left edge is being processed and attached into the tree. Some look-ahead may be available, or the parser may make guesses based on canonical phrase length. But if not, it seems that even a first-pass prosodic influence must sometimes lag a little behind syntactic parsing, and alter an attachment preference that was starting to take hold, rather than biasing the syntactic decision in advance.

4. Cross-linguistic differences in RC attachment: Predictions

A major puzzle for parsing theory was raised by Cuetos and Mitchell (1988). A preference to keep adjacent words close together in the tree structure is entailed (for right-branching languages) by the principle of *Late Closure (LC)*, or *Right Association*, or *Local Association*, or *Recency*. It is not important for present purposes which formulation of the locality principle is correct, or indeed whether they differ at all. What does matter is that Cuetos and Mitchell noted that this locality tendency appears not to be universal. For RC attachment to a complex NP with two competing noun hosts, as in (8), LC is the preferred structure for English but not for Spanish. In English the tendency is for the RC to modify the closer (lower) noun: *actress* in (8a). In Spanish the RC is more often taken to modify the earlier (higher) noun: *criada* in (8b). The percentages shown in (8) are only approximate, and the cross-linguistic differences are not large, but they are fairly stable across a number of experiments.

(8) a. Someone shot the servant of the actress who was on the balcony. 60% LC
b. Alguien disparó contra la criada de la actriz que estaba en el balcón. 40% LC

This discovery was disturbing, because all other evidence was compatible with the hypothesis that parser operations are innate and universal, differing only in that the parser applies different grammars to different languages. In particular, LC appears to be true for all other syntactic constructions, not only in English but also in Spanish (see Igoa, 1996).

[1] The factor Bader manipulated was sentence focus, induced by a focus particle and realized in speech by a focal accent. Revision of the first-pass prosody was thus not independent of revision of the focus structure of the sentence, making it difficult to be certain that implicit prosody was the cause of the observed parsing difficulty. However, Bader's experiment 3 substituted long focus particles for the short particles of experiment 1, and reanalysis was then easier. Bader argues that particle length affects only the prosodic contour. So step 3 is satisfied here.

Since this unsettling finding of cross-language variation in preferred RC-attachment, more languages have been put to the test. Some behave like English and some like Spanish, but the groupings resemble no sensible linguistic classification. (For details see Mitchell and Brysbaert 1998, Ehrlich et al. 1999, and Lovrić et al. 2000).

(9)

EARLY CLOSURE TENDENCY		LATE CLOSURE TENDENCY	
Afrikaans	German	Brazilian Portuguese	Norwegian
Croatian	Italian?	Egyptian Arabic	Romanian
Dutch	Russian	English (American)?	Swedish
French	Spanish	English (British)	

As the question marks indicate, a few cases are unresolved; I have made my best guess as to which list they belong in. Note that Swedish and Norwegian contrast with German and Dutch, and Romanian contrasts with Spanish, so the typology is not a simple one, such as Germanic/Romance. Some interesting explanations have been proposed, all quite promising but all disconfirmed by one language or another. These explanations appeal to word order variability, alternative forms for expressing genitive meanings, the anaphoric role of relative pronouns, as well as mere arbitrary parameterization. For a comprehensive survey, see Mitchell and Brysbaert (1998). A prosodic explanation is not disconfirmed by current data — though perhaps this is only because of our very fragmentary knowledge at present of the prosodic characteristics of different languages. The IPH holds out the hope that as we learn more about prosody, the classification in (9) will fall into place. It should be stressed, however, that all of these results in (9) are from silent reading experiments, so if the explanation does turn out to be prosodic it must involve implicit prosody.

The one point on which languages apparently agree is that for short RCs there is a stronger tendency toward low attachment than there is for mid-to-long RCs. This is so for all languages that have been tested so far, listed in (10).

(10)

Croatian	(Lovrić et al. 2000)
English	(Fernández and Bradley 2000)
French	(Pynte and Colonna 2000)
German	(Walter et al. 1999, though only for subject NP)
Spanish	(Fernández forthcoming)

So here we find yet another length effect, like those Hirose observed in Japanese. One might speculate that ALL constituent-length effects on attachment preferences are the result of prosodic phrasing, explicit or implicit. Frazier and Fodor (1978) attributed the chunking of input strings to limitations on short term working memory, as was the custom in those times. Examples of chunking, sometimes at odds with the syntactic phrasing, were the empirical motivation for the two-stage organization of the Sausage Machine parser.

Nowadays it seems more likely that the Sausage Machine's Preliminary Phrase Packager is the prosodic component of the grammar at work.

For F = RC-length, a prosodic explanation is not hard to conjure up. Our experiments (see below) confirm what prosodic theory predicts: that there is less probability of a prosodic break before a short RC such as *who cried,* than before a long one such as *who cried all through the night*. Although a prosodic boundary before a clause is common, an RC that consists of just one minor phrase cannot comfortably constitute a major phrase by itself, so it will tend to group with preceding words. Consider the implications for the complex NP structure in (8). If the whole complex NP including the RC can be one major phrase, no break is needed within it at all. So even if (11a) is more natural for a long RC, (11b) is possible for a short one. Other possibilities include (11c), with a break between the nouns, though this will depend in part on whether N1 is grouped with preceding material (e.g., a verb), since N1 is likely to be too light to stand alone.

(11) Prosodic phrasing and RC-length

a. [N1 of N2] [RC]	especially if RC is long
b. [N1 of N2 RC]	if RC is short (one minor phrase)
c. ...N1] [of N2 RC]	

In speech, differences in prosodic phrasing such as these are known to affect the RC-attachment preferences of hearers. Maynell (1999) has shown that more high attachment responses are given by listeners when a pause and prepausal lengthening create a prosodic break after N2, as in (11a). The IPH predicts that this will be so in silent reading also. Thus, the RC-attachment phenomenon provides steps [1] and [2] of the program for testing the IPH: RC-length affects overt prosodic phrasing, and overt prosodic phrasing affects RC-attachment. Step [3] is as for other length effects: the lack of any reason to suppose that constituent length affects attachment other than via prosody.[2] And step [4] is the set of facts in (10), indicating an apparently universal trend, in silent reading, for a lowering of a language's attachment preference when the RC is short.

Now consider F = language-specific prosodic rules. The aim is to account for the odd classification of languages in (9). As background for the cross-linguistic differences, I will tentatively assume as universal an interface principle that calls for a break in the prosody wherever there is a substantial discontinuity in the syntactic tree. This is related to the Align XP constraint of Selkirk (2000), but I take it to be graded, with a stronger tendency to break (or stronger acoustic markers of a break) when more XP brackets pile up at the same position (see Cooper and Paccia-Cooper 1980, and discussion in Schafer et al. 1996). Also, I take its directionality to be tied to that of the syntax: Align_R XP for

[2] Explanations have been offered in the literature for the influence of constituent length on preferred constituent order, but these have no obvious application to the attachment ambiguities of present concern. To the best of my knowledge only Pynte and Colonna (2000) offer a non-phonological account of the RC-length effects. This is an interesting proposal in terms of the timing of parsing decisions. I think it may adapt to left-branching languages less well than the IPH does, but evaluation must await more detailed implementation of both approaches.

right-branching languages or constructions, and $Align_L$ XP for left-branching. For English and French (right-branching), this means that a prosodic break is likely between a word that is at a lower tree position and an immediately following word at a higher position (i.e., with several right brackets between them). A perceiver might naturally construe such a break as a SIGNAL to attach the next word high, as in Maynell's experiments. Of course, it is not a RELIABLE signal if the break could have had some other cause (e.g., a length constraint, another alignment constraint in the grammar). In other words: there can be ambiguity for perceivers concerning the source of a prosodic break, and perceivers may sometimes resolve the ambiguity the wrong way. In particular, since left-edge and right-edge prosodic boundaries are not distinct (each break marks the end of one prosodic phrase and the start of another), a break that is left-aligned with the syntax may be misinterpreted as right-aligned, and vice versa. To explain the cross-language parsing differences, therefore, it could be hypothesized that the languages with a high RC-attachment preference are those whose interface constraints favor a prosodic break before an RC, a break which could be misconstrued as motivated by a syntactic discontinuity.

Grammars do differ with respect to whether they stipulate a prosodic break at the onset of certain types of syntactic constituents. Perhaps a more general theory is possible, but for present purposes we may think of these stipulations as language-specific instantiations of an $Align_L$ constraint, specified for a particular syntactic category, such as RC. The grammar of English appears to lack strict constraints of this sort. (Possibly it is flexible in this regard because its word order is not flexible; this is a common claim.) An English sentence may or may not break prosodically at the left edge of a CP or an IP or a VP, and so forth, depending on what is optimal on the basis of other factors. Thus, English is fairly free to place breaks where length considerations recommend them, and/or to use them to disambiguate tree structure in accord with configurationally sensitive interface rules. Languages such as French and Croatian, as discussed below, appear to be more rigid, and require a break at the left edge of certain types of constituents, such as RCs or PPs. This may also be true of other high-attaching languages such as Spanish, German, Dutch, etc.

To summarize: language-specific interface constraints may encourage or discourage the prosodic separation of an RC from the preceding noun. Given a universal propensity to interpret such separation as a sign of high attachment, this leads to different predictions for syntactic ambiguity resolution, as shown in (12). Prosodic structure (12a) would favor early closure, with attachment to the whole NP headed by N1. The phrasing (12c) would favor late closure, i.e., attachment to N2 only. It is unclear at present what the prediction should be for the grouping in (12b) which does not divide up the complex NP at all (more likely for short RCs than for long RCs, as noted in (11)). The result might be no attachment bias at all, or else the neutral prosody might allow a pure syntactic locality effect to show through, giving a preference for N2 attachment (late closure).

(12) Prosodic phrasing and RC-attachment preferences

a. [N1 of N2] [RC]	predicts early closure preference (N1-attachment)
b. [N1 of N2 RC]	neutral, or late closure

c. ...N1] [of N2 RC] predicts late closure preference (N2-attachment)

Another observation which relates ambiguity resolution of RC-attachment to overt prosody is the finding of Schafer et al. (1996) that if one of the two nouns has a focal accent, the RC is attracted to that noun. It may (or may not) be proper to extrapolate from focal accent to any differential prominence between the two nouns. If it is, then this would predict late closure in languages in which N2 is more prominent than N1, and early closure in languages in which N1 is more prominent than N2. This is a point that we are pursuing in other work but I will not report on it here. It can create methodological complications at step [3], as noted above. Thus, the discussion below concerns only prosodic phrase boundaries. It looks for evidence that these mediate RC-attachment in silent reading, as predicted by the IPH. Note that to the extent that attachment preferences can be linked to documented differences in prosodic phrasing across languages, the old Sausage Machine hypothesis that they are due to working memory limitations becomes less tenable.

5. Cross-linguistic differences in RC attachment: Methods and results

We have data relevant to the IPH from English, French and Croatian. For step [1], our experiments use unambiguous RC constructions, with N1 or N2 attachment disambiguated by agreement marking on the RC verb. Subjects read these sentences to themselves to understand them, and then read them aloud for recording. To identify prosodic breaks we measure F0 at the mid-point of the vowel in the stressed syllable of N1, of N2, and of the RC verb (which is the first prominence in the RC). We also measure any pauses and/or pre-pausal lengthening at the end of N1, and of N2. A major phrase boundary is typically associated with a pause or syllable lengthening before it, and/or a reset of F0 after it (though phonetic detail differs across languages). Thus we can establish which of the patterns in (12) is most characteristic of the RC construction in the language being tested, possibly depending on the length of the RC. Step [2] is a listening study to establish attachment preferences for syntactically ambiguous RCs with different overt prosodic phrasings. Step [4] is a written questionnaire with ambiguous sentences, each followed by a comprehension question to establish where the reader has attached the RC.

5.1. English and French

The experiments on English and French are from the dissertation in progress of Deirdre Quinn. Since this work is on-going the data are not final, but they are of interest. We have yet to complete step [2] for these languages, i.e., testing listeners' attachment preferences in spoken sentences where overt prosody is the independent variable. It is important to the logic of the argument to check that the distribution of overt prosodic breaks has the expected effect on hearer's parsing preferences, but in practice we have tended to save this part of the project to last, in hope that it can be safely assumed on the basis of Maynell's data (though that is admittedly only for English). The attachment facts are anticipated in tables (9) and (10) above: English generally favors low attachment, while French prefers high attachment though less so for short RCs than for long RCs. Quinn et al. (2000) reported acoustic data for four selected speakers of American English and three of

European French. There are now also data from eight undergraduate speakers of American English and eight of Canadian French. A questionnaire study jointly conducted with Roberto Almeida has established that the attachment facts for Canadian French are in line with those for European French (see Pynte and Colonna 2000 and references there). The questionnaire showed high attachment for long RCs and no bias for short ones: 69.7% and 50.4% N1 attachment respectively ($p<.005$). The attachment facts are collated in (13) to facilitate comparison with the prosodic results summarized in (14).

(13)

	ENGLISH	FRENCH
LONG-RC	low attachment (LC)	high attachment (EC)
SHORT-RC	lower attachment (LC)	no attachment bias

(14)

	ENGLISH	FRENCH
LONG-RC	forced high: break	forced high: break
	forced low: no break	forced low: break
SHORT-RC	forced high: no break	forced high: no break
	forced low: no break	forced low: no break

Consider the results for short RCs first. There is no reliable indication in the F0 and pause/duration data, of a prosodic break between N2 and a short RC. This is the case for both languages, and regardless of whether RC-attachment was forced high or low by number agreement. This is as expected in view of the unnaturalness of a single phonological word as a major phrase by itself. It also offers the most plausible explanation of why ambiguity resolution preferences shift toward lower attachment for short RCs in both languages. The prosodic contour for short RCs lacks the prosodic boundary that perceivers tend to interpret as a sign of high RC-attachment. Given a preference for syntax-prosody congruence, as observed above, that contour is more naturally interpreted as a low-attachment contour. (See section 6 for further discussion.)

For long RCs the results differ across languages. For English the acoustic measurements indicate the presence of a prosodic break only for forced-high attachment of a long RC. For French there is evidence of a break preceding long RCs regardless of height of attachment. Let us consider English first. It appears, as anticipated, that English grammar does not impose an Align_L RC constraint; the phrasing is controlled only by length considerations and the general principle of congruence with the syntax. Thus the prosodic contour reflects the syntactic structure. For high attachment, where N2 and the string-adjacent relative pronoun are structurally non-adjacent, a prosodic discontinuity accompanies the discontinuity in the syntactic tree. For low attachment, where N2 and the relative pronoun are structurally close in the tree, they are undivided prosodically. The mild preference for low attachment in English ambiguously-attached RCs can be accounted for on the assumption that, though the grammar leaves it free, the default is for there to be no prosodic break in English before an RC, even a long one. (This is for

restrictive RCs. For non-restrictive RCs, not discussed here, a break is much more likely.) This tendency not to separate an RC from the preceding noun in English agrees with informal observation, though more substantial evidence must be gathered. An intriguing indication that no-break is the English default is that naive subjects tend to produce a less distinctive contour for high-attachment examples than careful speakers do.[3] For French, by contrast, the acoustic data suggest that the grammar more or less insists on a prosodic break before a long RC, regardless of the position of the RC in the syntactic tree. Since this is the prosodic pattern that is generally associated, by syntax-prosody congruence, with high attachment, then the fact that this prosodic contour is dictated by French grammar can explain the preference in French for high attachment of ambiguous long RCs.

Here, then, we observe a phonological difference between the two languages which can explain readers' differing ambiguity resolution preferences. For French prosodic breaks, at least in the construction under consideration here, the syntactic category of the following constituent matters more than the position of that constituent in the tree, while the reverse is true for English. This could be captured formally by constraint ranking, as in Optimality Theory: Align_L RC ranks high in French but low in English. Optimal length constraints such as BinMin and BinMax are operative in both languages, as is the graded Align_R XP constraint which reflects structural discontinuity. Linguistic facts such as these may have several interesting psycholinguistic implications. For instance, prosody projection in reading should be easier where the determinants of prosodic contours are local and easily recognizable facts such as the edge of a particular syntactic phrase type that is lexically well-marked, like the left edge of an RC in a language with non-deletable clause-initial relative pronouns. By contrast, prosodic sensitivity to constituent length means, as noted above, that some look-ahead is needed (though perhaps not always available) for readers to estimate how long a constituent is going to be. The problem is multiplied wherever prosodic phrasing is configuration-based, since readers can't be sure where the prosodic boundaries belong without knowing how the structure of the sentence is going to unfold. But note the trade-off with listening. Languages with fixed, category-based alignment with syntax may be easier to read aloud, but at the sacrifice of configurational disambiguation for hearers: they lack prosodic means of signaling a higher versus a lower attachment site for the same syntactic category.

5.2. Croatian

The work on Croatian is by Nenad Lovrić and the details will be presented in his forthcoming dissertation. I summarize here the findings reported by Lovrić at al. (2000; 2001). There are two constructions to be considered in this language. Croatian generally uses a prepositionless structure N1 - N2[GEN] - RC. Lovrić's experimental data show that

[3] While both our English-speaking subject groups showed an F0 reset at the beginning of a high-attaching long RC only, the effect was sharper for the original four speakers than for the undergraduate subjects. Also, the former group but not the latter had a distinctively higher F0 on N2 for N2-attachment than for N1-attachment, which may reflect the Schafer at al. finding noted above, that the RC tends to modify the more prominent noun. This prominence effect would seem to be a more fragile phenomenon than the prosodic phrasing effect, if our group data are representative of more careful versus less careful speech.

this construction behaves much like the French N1 - *de* N2 - RC with respect to both prosodic phrasing and preferred attachment. That is: for long RCs (two or more prosodic words), a prosodic break is more or less obligatory before the RC, and RC-attachment tends to be high (73.6% N1); for short RCs (1 prosodic word), a break is less likely and there is a correspondingly smaller proportion of high attachment responses (48.6% N1). But some dialects of Croatian, particularly the Zagreb dialect, also have the prepositional construction N1 - *od* N2[GEN] - RC. The forms with and without preposition *od* have the same meaning (though in other contexts *od* can mean 'from', like German *von*). The form with *od* is more colloquial, but within a fairly wide range of conversational speech the two are in free variation. Thus, as step [3] of the research program requires, there is no semantic or pragmatic contribution of *od*. It is arguable (though of course also disputable) that there is also no significant syntactic difference between the forms with and without *od*, i.e., that they differ only with respect to whether the preposition is phonologically overt or null. (See den Dikken et al. 1999 for discussion of this syntactic issue in several languages, though not Croatian.)

Of interest for present purposes is that *od* favors a prosodic break at the left edge of the PP. (This is probably not a fact about *od*, but a fact about Croatian prepositional phrases in general; the latter is suggested by the discussion of the prosodic consequences of the proclitic status of prepositions, by Godjevac 2000.) To break before *od* and then again before the RC would violate the optimal length constraints by making N2 into a major phrase by itself. Therefore an indirect effect of *od*'s presence would be to reduce the chance of a prosodic break between N2 and the RC. In sum: the break induced by *od*, at the beginning of the *od*-PP, would not only separate N1 from N2, but also push N2 into a grouping with the RC, even a long RC which would otherwise have been separated off by the prosodic break preceding it that Croatian normally requires. For an *od*-construction, therefore, (15b) is a more natural phrasing than (15a) without *od*. This is relevant to parsing because the phrasing in (15a) is of the kind that encourages high RC-attachment, while the phrasing in (15b) is associated with low RC-attachment. Lovrić's data for step [2] confirm that hearers do attach high when there is an overt break before the RC, and low when there is a break before (*od*) N2 only. The IPH prediction for ambiguity resolution in silent reading in Croatian is thus that the French-type RC-attachment pattern without *od* should give way to low, or at least lower, RC-attachment with *od* present.

(15)	a.	[N1 N2] [RC]	(especially if RC is long)	favors N1-attachment
	b.	...N1] [*od* N2 RC]	(especially if RC is short)	favors N2-attachment

The results of Lovrić's experiments support these predictions concerning the phonological and syntactic effects of *od*. Only pause/pre-pausal lengthening data were used to estimate prosodic break positions, because Croatian is a tone language so its F0 contours are subject to a variety of other influences. But the duration data are very clear. There was, as expected, significant lengthening of N1 when *od* is present (mean difference 113 msec, $p<.01$), indicating a boundary at the left of the *od*-PP. There was no reliable interaction with either the length or the forced high/low attachment of the RC; see (16). This is not surprising since the RC has not yet been processed at this point in the sentence.

For N2 duration, many factors were in play. There was a main effect of RC-length (N2 was longer before a long RC than a short RC; mean difference 110 msec, $p<.01$). There was also a main effect of attachment (longer N2 before an RC forced high than an RC forced low; mean difference 40 msec, $p<.03$). These effects of RC-length and attachment occurred both with and without *od*. N2 lengthening implies a boundary following N2, preceding the RC. So here we are seeing the usual increase in the probability of a prosodic boundary preceding a long RC, and preceding a high-attached RC. It is of interest that these subjects (ten college students) did distinguish prosodically between high and low RC-attachment, though sensitivity to syntactic configuration (attachment height) was small compared with the effect of RC-length.

(16)

N1 lengthening (msec) due to *od*	Forced-high attachment	Forced-low attachment
Long RC	98	118
Short RC	107	130

N2 lengthening (msec) due to *od*	Forced-high attachment	Forced-low attachment
Long RC	-45	-82
Short RC	-1	-34

Now we can consider how N2 duration was affected by the presence or absence of *od*, as shown in (16). As expected, N2 became shorter (mean 40 msec, $p<.01$), indicating a reduced tendency for a prosodic break between N2 and RC when *od* is present, presumably because the break preceding *od*-N2 inhibits another break immediately after N2, as argued above. This effect of *od* on N2 showed interactions with the length and attachment-height of the RC, as seen in (16). The N2-shortening due to *od* was less preceding short RCs than long ones. This is probably just a floor effect since N2 was already very short preceding short RCs even without *od*, i.e., few breaks before short RCs in any case. N2-shortening was less for forced-high attachment than for forced-low. This is as expected also, since high attachment requires maintaining a break between N2 and the RC even in the face of contrary pressure due to *od*. Thus, all of the overt prosodic consequences of *od* are explicable, given the one assumption that *od* triggers a break before it, which tends to suppress the usual break before an RC in Croatian.

The data for syntactic ambiguity resolution in silent reading are shown in (17). As predicted, the presence of *od* lowered the RC-attachment preference ($p<.01$). Thus it appears once again that manipulation of a property (presence or absence of *od*) which alters prosody can affect syntactic processing even when no overt prosody is present. With all four steps of the program in place, we can conclude that the effect on syntactic processing is mediated by implicit prosody imposed on the sentence by the reader.

(17)

% high attachment	No preposition	Preposition *od*	Effect of *od*
Long RC	69.8	56.3	13.5
Short RC	44.8	26.3	18.5

6. Summary, conclusions, and a question

The data reviewed here substantiate the role of implicit prosody in sentence processing, which has been overlooked until very recently in both experiments and theory development. We have seen it at work here in three different constructions and four languages, so it appears to be a quite widespread phenomenon. The methodological program set out in (2) is important because it makes it possible to argue that the correlations observed between implicit prosody and syntactic parsing preferences are more than just correlations: implicit prosody is shown to be a cause, not merely an effect, of the syntactic preferences that readers exhibit.

The fact that readers' ambiguity resolution tendencies can be altered by manipulating purely phonological properties of the stimulus is worth knowing for practical purposes. Phonological factors need to be properly controlled even in reading tasks, to prevent the intrusion of prosodic artifacts into experiments on other aspects of sentence processing. For parsing theory, recognizing the role of implicit prosody is important because it makes it possible to defend the idea of a fully universal innate human parsing mechanism. Apparent differences in parsing strategies from one language to another may be attributable to differences in their grammars, once it is acknowledged that grammars contain prosodic components, with principles that are similar but not identical across languages. That ALL cases of apparent non-universality of the parsing routines might be the actions of a universal parser applying non-universal prosodic grammatical principles (Fodor 1998) is an open-ended speculation which can't be proved, but at very least the IPH is an explanatory tool that we may wield when exceptions to universality threaten.

To provide additional evidence for the IPH we can seek out more language-specific peculiarities like the prosodic effects of optional *od* in Croatian, and see whether they have equally specific effects on parsing in silent reading. The more distinctive the prosodic pattern, the more precisely its effect on reading could be confirmed or disconfirmed. Perhaps also we can find a set of criteria to sort subjects along a scale of how strongly they mentally represent prosody in silent reading. The IPH makes the clear prediction that readers with strong inner prosody should show robust parsing consequences of constituent length and alignment rules, while people who are insensitive to implicit prosody should be immune to them.

Finally, implicit prosody has some curious aspects which we must hope are explicable though they have a slightly paradoxical twist to them. Most interesting is why prosodic ambiguities are resolved in the way they are. Why does the inner ear misinterpret the inner voice? For instance, it has been claimed here that a contour with the universal characteristics of high-attachment is projected by French speakers onto all (long) RCs,

including those whose meaning entails low attachment. It seems proper to conclude, therefore, that this "high-attachment" prosody is actually ambiguous in French between high attachment and low attachment. But in that case, when a French reader encounters an ambiguously-attached RC and assigns prosody to it by the left alignment rule for RCs, why should that reader be more inclined to interpret the RC as attaching high? The IPH explanation for the cross-language attachment data assumes that readers do misinterpret the causes of their own inner prosody in just this way. A prosodic boundary imposed in accord with the rule that long RCs must be preceded by a break is sometimes misidentified as being present to mark a discontinuity in the tree structure (high attachment). Likewise, the absence of a break before a short RC is dictated by optimal phrase-length principles but is often construed as signaling local (low) attachment.

Since this happens in listening too, it is not peculiar to the IPH. It is an interesting phenomenon that we need to understand in any case. It is inevitable that a listener (via real ears or the inner ear) will sometimes guess wrong about the cause of an ambiguous prosodic break. But a random guess would give an even split (e.g., a 50% probability of hearing a pre-RC break as a signal of high attachment), whereas the evidence we have reviewed indicates the existence of a skewing factor favoring a configurational interpretation of prosodic contours. What could it be? It is imaginable that alignment constraints are ranked differently for speakers and perceivers. Since that seems an unwelcome addition to phonological theory, we may seek a more performance-oriented approach. Suppose tree configuration is universally the 'ideal' basis for prosodic phrasing, but for practical reasons producers prefer category-based rules that can be applied on a local basis, as observed above. The practical problems don't apply to perceivers, so they interpret ambiguous contours by configurational principles. Other explanations are imaginable too, not all mutually exclusive. Perhaps the bias to configurational interpretation is an instance of *Minimal Everything*: it might be less work to apply a general principle relating prosodic discontinuities to syntactic discontinuities than to apply a language-specific rule. Or perhaps the parser gives precedence to configuration over category because configuration is on the parser's mind: its job is to figure out how the parts of the sentence relate to each other. For that purpose, a configurational interpretation has the advantage of offering the perceiver a basis for syntactic ambiguity resolution (right or wrong), while a category-specific rule does not.

Thus there is an abundance of explanatory notions to be sifted through in further research. The reason for sketching some of them here is to make sure there is no inherent contradiction in the idea of a neutral (default) prosody assignment leading a perceiver to make a non-neutral guess about the structure that gave rise to that prosody. Years of work has gone into studying the parsing principles or 'strategies' that bias the parser's choice of which syntactic structure to build when a word string is ambiguous, but we are only just beginning to study the comparable question for prosodic ambiguities.

References

Bader, M. 1998. Prosodic influences on reading syntactically ambiguous sentences. In *Reanalysis in sentence processing*, ed. J. D. Fodor and F. Ferreira, 1-46. Dordrecht: Kluwer Academic Publishers.

Brown, R. 1970. Psychology and reading. In *Basic studies on reading*, ed. H. Levin and H. P. Williams, New York: Basic Books.

Cooper, W. E. and J. M. Paccia-Cooper. 1980. *Syntax and Speech.* Cambridge MA: Harvard University Press.

Cuetos, F. and D. C. Mitchell. 1988. Cross-linguistic differences in parsing: Restrictions on the use of the Late Closure strategy in Spanish. *Cognition* 30:73-105.

den Dikken, M., E. Stenshoel, and M. Vinereanu. 1999. Possessed! Paper presented at the CUNY Linguistics Program Conference, New York.

Ehrlich, K., E. Fernández, J. D. Fodor, E. Stenshoel, and M. Vinereanu. 1999. Low attachment of relative clauses: New data from Swedish, Norwegian and Romanian. Poster presented at the 12th Annual CUNY Conference, New York.

Fernández, E. M. forthcoming. *Bilingual sentence processing: Relative clause attachment in English and Spanish.* Amsterdam: John Benjamins.

Fernández, E. M. and D. Bradley. 1999. Length effects in the attachment of relative clauses in English. Poster presented at the 12th Annual CUNY Conference on Human Sentence Processing, New York.

Fodor, J. D. 1998. Learning to parse? *Journal of Psycholinguistic Research* 27:285-319.

Fodor, J. D. and L. Frazier. 1980. Is the human sentence parsing mechanism an ATN? *Cognition* 8:417-459.

Frazier, L. and J. D. Fodor. 1978. The Sausage Machine: A new two-stage parsing model. *Cognition* 6:291-326.

Godjevac, S. 2000. An autosegmental/metrical analysis of Serbo-Croatian intonation. *Ohio State University Working Papers in Linguistics*, 54:79-142. Columbus, OH.

Harman, G. H. 1965. The inference to the best explanation. *Philosophical Review* 74.1:88-95.

Hirose, Y. 1999. Resolving reanalysis ambiguity in Japanese relative clauses. Doctoral dissertation, City University of New York.

Hirose, Y. 2000. The role of constituent length in resolving reanalysis ambiguity. In *Proceedings of the First Tokyo Conference on Psycholinguistics*, 55-74. Tokyo: Hitsuji Syobo Publishing Ltd.

Igoa, J. M. 1995. Parsing decisions and the construal hypothesis: Attachment preferences in primary phrases in Spanish. Paper presented at the Second Symposium on Psycholinguistics, Tarragona, Spain.

Inoue, A. and J. D. Fodor. 1995. Information-paced parsing of Japanese. In *Japanese sentence processing*, ed. R. Mazuka and N. Nagai. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kosslyn, S. M. and A. C. Matt. 1977. If you speak slowly, do people read your prose slowly? Person-particular speech recoding during reading. *Bulletin of the Psychonomic Society* 9.4: 250-252.

Kubozono, H. 1993. *The organization of Japanese prosody.* Tokyo: Kurosio Publishers.

Lovrić, N. In preparation. Implicit prosody in silent reading: Relative clause attachment in Croatian. Doctoral dissertation, City University of New York.

Lovrić, N., D. Bradley and J. D. Fodor. 2000. RC attachment in Croatian with and without preposition. Poster presented at the AMLaP Conference, Leiden.

Lovrić, N., D. Bradley and J. D. Fodor. 2001. Silent prosody resolves syntactic ambiguities: Evidence from Croatian. Paper presented at the SUNY/CUNY/NYU Conference, Stonybrook. NY.

Maynell, L. A. 1999. Effect of pitch accent placement on resolving relative clause ambiguity in English. Poster presented at the 12th Annual CUNY Conference on Human Sentence Processing, New York.

Mazuka, R. and K. Itoh. 1995. Can Japanese speakers be led down the garden path? In *Japanese sentence processing*, ed. R. Mazuka and N. Nagai. Hillsdale, NJ: Lawrence Erlbaum Associates.

Mitchell, D. C. and M. Brysbaert. 1998. Challenges to recent theories of language differences in parsing: Evidence from Dutch. In *Sentence processing: A crosslinguistic perspective*, ed. D. Hillert, 313-335. San Diego: Academic Press.

Peirce, C. S. 1931-1958. *Collected papers*, ed. C. Hartshorne and P. Weiss. Cambridge, MA: Harvard University Press.

Pynte, J. and S. Colonna. 2000. Decoupling syntactic parsing from visual inspection: The case of relative clause attachment in French. In *Reading as a perceptual process*, ed. A. Kennedy, R. Radach, D. Heller and J. Pynte, 529-547. Oxford: Elsevier.

Quinn, D. In preparation. Implicit prosody in silent reading: French and English relative clause attachment. Doctoral dissertation, City University of New York.

Quinn, D., H. Abdelghany, and J. D. Fodor. 2000. More evidence of implicit prosody in silent reading: French, English and Arabic relative clauses. Poster presented at the 13th Annual CUNY Conference on Human Sentence Processing, La Jolla, CA.

Rayner, K. and A. Pollatsek. 1989. *The psychology of reading*. Englewood Cliffs, NJ: Prentice Hall.

Schafer, A., J. Carter, C. Clifton, Jr. and L. Frazier. 1996. Focus in relative clause construal. *Language and Cognitive Processes* 11:135-163.

Selkirk, E. 2000. The interactions of constraints on prosodic phrasing. In *Prosody: Theory and experiment*, ed. M. Horne, Dordrecht: Kluwer Academic.

Slowiaczek, M. L. and C. Clifton, Jr. 1980. Subvocalization and reading for meaning. *Journal of Verbal Learning and Verbal Behavior* 19: 573-582.

Walter, M., C. Clifton, L. Frazier, B. Hemforth, L. Konieczny, and H. Seelig. 1999. Prosodic and syntactic effects on relative clause attachment in German and English. Poster presented at the AMLaP Conference, Edinburgh.

Ph.D. Program in Linguistics
Graduate Center, CUNY
365 Fifth Avenue
New York, NY 10016

jfodor@gc.cuny.edu