10

Parsing and Incremental Understanding During Reading

MARTIN J. PICKERING AND
MATTHEW J. TRAXLER

This chapter discusses a range of findings from our eye-tracking research in relation to current theories of parsing. First, we discuss findings that implicate extremely rapid semantic processing and findings that suggest that the degree of semantic processing affects the ease of recovery from misanalysis. We then provide an analysis of theories of initial parsing decisions that emphasises two dimensions: whether they can draw upon all available information or not; and whether they attempt to select or foreground the analysis that is most likely to be correct at the point of ambiguity. We next discuss a range of evidence from our laboratory that provides support for theories that do not attempt to select the most likely analysis. The evidence may support an account where at least some parsing decisions are driven by syntactic information alone. Alternatively, it may support an account in which the parser draws upon all available information, but where its goal is to select the most informative analysis, not the most likely one.

10.1 Immediate Effects of Plausibility

One of the clearest conclusions of research into language comprehension is that a great deal of processing occurs incrementally. Thus, lexical access, syntactic analysis, and semantic interpretation are not delayed, but begin as soon as the relevant word is heard or read (e.g., Marslen-Wilson, 1973; Just & Carpenter, 1980). The strongest of these findings is that the processor does not delay in providing a semantic interpretation for a sentence fragment as a whole: such

The order of authorship is arbitrary. The research discussed in this chapter was supported by ESRC Grant No. R000234542 and a British Academy Postdoctoral Fellowship (both awarded to the first author).

semantic processing of course requires that lexical and syntactic processing have already occurred.

Even more strikingly, this interpretation is immediately integrated with relevant background knowledge and information provided by discourse context. In Traxler and Pickering (1996b), we found that such semantic integration sometimes can have taken place by the time that the eye finishes its first fixation on a word. This finding suggests that semantic processing can occur as soon as about 250 ms after fixating a word. However, we cannot be certain that integration can occur within 250 ms, because some semantic processing of a word may occur during the previous fixation if it occurred towards the end of the previous word (e.g., Rayner & Pollatsek, 1989).

For instance, in one experiment, participants read sentences like (1a–d):

(1) a. That's the pistol with which the heartless killer shot the hapless man yesterday afternoon.
b. That's the garage with which the heartless killer shot the hapless man yesterday afternoon.
c. That's the pistol in which the heartless killer shot the hapless man yesterday afternoon.
d. That's the garage in which the heartless killer shot the hapless man yesterday afternoon.

The fragments up to the verb *shot* in sentences (1b, c) are likely to form the beginnings of implausible sentences, because it is impossible to shoot with a garage, or to shoot in a pistol. This contrasts with (1a, d), which are plausible. We found that subjects spent longer reading *shot* in (1b, c) than (1a, d) from the first fixation onwards. Our results cannot be due to low-level lexical effects (e.g., priming from *pistol* to *shot*), because of the crossed design. So readers must have assessed the plausibility of shooting with a garage or pistol, or shooting in a garage or pistol, as soon as they encountered the verb *shot*.

Note that the semantic anomalies in (1b, c) are fairly gross, in that the sentences describe impossible or absurd situations. In many cases, the anomaly corresponds to a selection-restriction violation. Our intention was to determine the effects of strong plausibility manipulations, and so we selected materials for which subjects rated the plausible sentences as highly plausible (5 or more on a 0–7 scale) and the implausible sentences as highly implausible (2 or less on such a scale). However, the implausible sentences are clearly grammatical. We conclude that some semantic processing begins very rapidly (though more elaborate semantic processing may, of course, be delayed, or never occur at all).

The implausible sentence fragments are not complete sentences, nor do they normally constitute possible complete sentences. For example, when the verb *shot* is encountered, it lacks an NP-object argument. Chater, Pickering, and Milward (1995) pointed out that such fragments do not have propositional interpretations according to standard formal semantics. They suggested that the processor incrementally constructs two semantic representations at different representational levels. The representation at the input level corresponds to the standard formal semantic analysis for a sentence fragment. At *shot*, this representation does not correspond to a proposition, because the verb *shot* is a two-place predicate, and it has only been combined with one semantic argument. Assuming that only propositions can serve as premises in inference, and that semantic integration involves inferences using premises drawn from both the sentence fragment and background knowledge, the input level cannot form the basis of incremental understanding. Hence Chater et al. proposed that the processor also constructs a knowledge-level representation, which is propositional; it is this representation that is integrated with general knowledge and can, for instance, explain the immediate effects of plausibility in the processing of sentence fragments.

Other experiments show that locally ambiguous fragments are incrementally understood. In a second experiment, Traxler and Pickering (1996b) contrasted sentences like (2a, b), which were embedded in short contexts:

(2) a. We like the book that the author wrote unceasingly and with great dedication about while waiting for a contract.
b. We like the city that the author wrote unceasingly and with great dedication about while waiting for a contract.

Although both sentences are globally plausible (as measured by pre-tests), the plausibility of the sentences differ on a misanalysis where *the book* or *the city* serves as the NP object of *wrote*. At *wrote*, (2a) is plausible, because an author is likely to write a book, but (2b) is implausible, because an author is not likely to write a city. We found that subjects took longer to read (2b) than (2a) from the first fixation after encountering *wrote*. Again, we suggest that the source of the plausibility effect is the integration of the knowledge-level representation with background knowledge.

Interestingly, both experiments found first-fixation plausibility effects on a (finite) verb. In other studies, we have found rapid plausibility effects resulting from the processing of nouns, but never on first fixation. For instance, Pickering and Traxler (1998a) had subjects read subordinate clause ambiguities like (3) and complement clause ambiguities like (4), together with disambiguated

control sentences:

(3) a. As the woman edited the magazine about fishing amused all the reporters.
b. As the woman sailed the magazine about fishing amused all the reporters.

(4) a. The criminal confessed his sins which upset kids harmed too many people.
b. The criminal confessed his gang which upset kids harmed too many people.

The misanalyses in (3a) and (4a) are plausible, and the misanalyses in (3b) and (4b) are implausible. Subjects made more regressive eye-movements from *the magazine* in (3b) than in (3a) before the eye went past this region. Subjects made more regressive eye-movements from *which upset kids* in (4b) than in (4a), though there were no differences between *his sins* in (4a) and *his gang* in (4b). These experiments again show rapid effects of plausibility in ambiguous sentences. No such differences occurred in the disambiguated control sentences, indicating that the initial (mis)analysis was the locus of the plausibility effect. We have shown similar effects in other experiments for which the misanalysis is unlikely to be correct (see below).

Pickering and Traxler (1998a) found comparable effects in sentences where the manipulation of plausibility would only have an effect if subjects integrated material from a previous sentence:

(5) a. The janitor polished bronze statues of the old maths professor that the principal hated and the dean of the art school. While the janitor was polishing the professor that the principal hated reviewed the spring term teaching schedule.
b. The janitor polished bronze statues for the old maths professor that the principal hated and the dean of the art school. While the janitor was polishing the professor that the principal hated reviewed the spring term teaching schedule.

In (5a), the misanalysis of the target sentence is plausible, because *the professor that the principal hated* can refer to the statue. But in (5b), the misanalysis is implausible, because the same phrase can only refer to an actual professor. Participants spent longer reading this phrase during first pass (but not on first fixation) in (5b) than (5a). But in control sentences disambiguated by a comma after *polishing*, the manipulation of context had no effect. Hence, we found that participants made rapid use of plausibility information that involved integrating material across sentences.

All of these experiments provide clear evidence for incremental understanding. We have three explanations for why we have only found first-fixation plausibility effects on the verbs in Traxler and Pickering (1996b), besides chance. First, the processor may be more likely to perform rapid semantic processing on verbs than on other words, because they are the central elements of sentences. Second, there may be more alternative analyses available for nouns than for verbs. For example, *The woman sailed the magazine* might be the beginning of the plausible *The woman sailed the magazine owner's yacht*, where *magazine* is not the head noun of the object of *sailed*. The existence of this alternative might slow semantic processing of *magazine* slightly. However, it does not prevent rapid attachment of *magazine* as the head noun of the object of *sailed*, as the regressions effect demonstrates. Finally, both first-fixation effects occurred at points where the processor formed an unbounded dependency. Perhaps the processor is particularly alert when it performs this fairly unusual syntactic operation.

Our findings are, of course, compatible with the range of evidence for incremental understanding during the processing of both spoken and written language (e.g. Boland, Tanenhaus, Garnsey, & Carlson, 1995; Garnsey, Tanenhaus, & Chapman, 1989; Garrod, Freudenthal, & Boyle, 1994; Marslen-Wilson, 1973, 1975; Marslen-Wilson, Tyler, & Koster, 1994; Swinney, 1979). The main conclusion of interest is the extreme rapidity with which such effects can occur during normal reading.

10.2 Plausibility of Misanalysis Affects Reanalysis

At the point of disambiguation, a locally ambiguous sentence becomes unambiguous. If a reader has adopted the wrong analysis, then this analysis will have to be abandoned at this point at the latest. Recovery from some forms of misanalysis is harder than others (e.g. Ferreira & Henderson, 1991; Warner & Glass, 1987). In a number of studies, we have shown that semantic characteristics of the misanalysis affect recovery. In particular, recovery is harder if the misanalysis is plausible than if it is implausible.

Consider (3) again, from Pickering and Traxler (1998a):

(3) a. As the woman edited the magazine about fishing amused all the reporters.
b. As the woman sailed the magazine about fishing amused all the reporters.

We have already shown that subjects initially treated the NP *the magazine about fishing* as the object of *edited* or *sailed*, and that they incrementally understood the fragment under this analysis. Now consider what might happen next. In (3a),

readers should strongly commit to this interpretation and should successfully integrate it with general knowledge. But in (3b), readers should commit less strongly to this interpretation. They might abandon the misanalysis even before reaching the point of disambiguation, and adopt the correct analysis. They might retain the implausible misanalysis, but commit themselves to this interpretation only weakly. Finally, their strategy might vary between these two alternatives.

Pickering and Traxler (1998a) argued that readers commit more strongly to an analysis if its interpretation is plausible than if it is implausible. Participants spent longer reading, and produced more regressions from, the disambiguating region from the verb *amused* onwards, in (3a) than in (3b), though first-pass effects on the verb itself were weak. Since they spent longer reading the ambiguous NP in (3b) than in (3a), we found a characteristic crossover pattern, with more disruption occurring in sentences with implausible misanalyses during the ambiguous region, but more disruption occurring in sentences with plausible misanalyses after disambiguation.

We found similar effects with sentences like (4), repeated here:

(4) a. The criminal confessed his sins which upset kids harmed too many people.
b. The criminal confessed his gang which upset kids harmed too many people.

Participants made more regressions after the disambiguating verb *harmed* in the initially plausible (4a) than in the initially implausible (4b), and they read the verb for longer in (4a) than (4b) in one experiment. We also found similar effects with (5a, b), discussed earlier. Here, the effects on the disambiguating verb were significant from the first fixation onwards.

These findings provide strong evidence that subjects commit more strongly to an analysis if that analysis is plausible than if it is implausible. It is less clear whether implausible sentences trigger immediate reanalysis or not. In sentences like (4), processing difficulty after disambiguation for sentences with implausible misanalyses (i.e., (4b)) was greater than processing difficulty for unambiguous control sentences. This provides some evidence that participants do not always abandon the misanalyses before the point of disambiguation. In any case, these experiments strongly suggest that the degree of semantic commitment to an analysis affects the process of reanalysis.

These effects of plausibility on analysis and reanalysis provide good evidence about how the parser selects an analysis. For instance, the finding of disruption on *the magazine* in (3b) compared with (3a) indicates that the parser selected the transitive analysis under which *the magazine* is the object of the subordinate verb during very early processing. This conclusion is

reinforced by the finding of disruption in (3a) compared with (3b) after the point of disambiguation. Hence, we can conclude that readers adopted the transitive analysis on at least a high proportion of occasions whilst reading these sentences.

Since most or all verbs in this study were preferentially transitive, the result is not in itself particularly surprising, and would be predicted by current theories (e.g., Ford, Bresnan, & Kaplan, 1982; Frazier, 1979; Trueswell, Tanenhaus, & Kello, 1993). The interesting point is that the conclusion does not depend on the comparison of different construction types or different uses of punctuation between conditions. This removes a possible confound that occurs in many studies of this construction (e.g., Clifton, 1993; Ferreira & Henderson, 1991; Frazier & Rayner, 1982; Mitchell, 1987; Warner & Glass, 1987). Similar conclusions can be drawn about the complement clause ambiguities in (4).

Hence, we can employ the manipulation of plausibility to investigate the processing of sentences that may shed light on more controversial aspects of the parsing process. In particular, we can use it to determine how the parser makes an initial choice of analysis at a point of ambiguity. We next outline some relevant aspects of current controversies and then interpret a number of our experiments in relation to these controversies.

10.3 Models of Parsing

Current theories of language comprehension pay great attention to the process of syntactic ambiguity resolution. Let us briefly outline a taxonomy of possible models of this process, and remove those models that are not serious contenders. First, incremental models can be contrasted with delay models. In an incremental model, each new word is syntactically and semantically integrated with the context as it is encountered. The evidence for incremental understanding, as well as the mere existence of garden paths, demonstrates that the parser constructs analyses in an incremental manner. We cannot rule out the possibility that the parser employs some delay under some circumstances, but the most straightforward assumption is that it does not.

Incremental models can be serial, with only one analysis being constructed and interpreted, or parallel, with more than one analysis being constructed at once. We can rule out unranked parallel models, where the parser treats all different analyses as equally important, because such models would not predict garden-path effects. Hence, the only reasonable possibilities are serial models and ranked-parallel models. In both cases, the parser makes a single analysis most prominent, choosing it in a serial model, foregrounding it in a parallel model (see Pickering, in press, for further discussion).

We now focus on the question of how the parser selects this analysis. (We use 'selects' for both serial and ranked parallel models.) Rather than discuss the predictions of different parsing models in detail, we ask a general question: does the parser select the analysis that is most likely to be correct, given the information available at the point of ambiguity? Such an account would appear to make for a well-adapted parser, assuming that reanalysis involves some cost, because selecting the most likely analysis would minimise the need for reanalysis. Additionally, if the person immediately acted upon the interpretation of the most likely analysis in some way, then the chances of that person acting in an inappropriate manner would be minimised. We call this a likelihood account. Later, however, we shall suggest that likelihood may not be as adaptive as it initially appears.

But likelihood certainly has an intuitive merit, so it is perhaps surprising that no one has explicitly proposed it. This may, in part, be because most researchers in sentence processing have concentrated on algorithmic or mechanistic models rather than on models that emphasise the higher-level goals of the system (Anderson, 1991; Marr, 1982). However, some models make predictions closely related to likelihood. One class of account proposes that the parser selects the analysis that corresponds to the most frequent subcategorisation frame of the verb (Ford et al., 1982; see J.D. Fodor, 1978, with reference to unbounded dependencies). Subcategorisation frame preferences are not the only factor involved in the computation of likelihood, but they are clearly a very important component, at least for ambiguities involving the attachment of arguments. Hence, such models are roughly compatible with the goal of maximising the likelihood that the selected analysis will ultimately be correct.

In addition, parallel models that assume that parsing preferences are based on the simultaneous interaction of multiple constraints approximate to likelihood models (MacDonald, 1994; MacDonald, Pearlmutter, & Seidenberg, 1994; Trueswell, Tanenhaus, & Garnsey, 1994; Trueswell et al., 1993; cf. Taraban & McClelland, 1988; Tyler & Marslen-Wilson, 1977). These constraints relate to any properties of the encountered sentence that may influence its continuation, including subcategorisation preferences, other syntactic cues, the meaning of the fragment, the nature of the discourse context, and prosody or punctuation. Hence, such models are unrestricted, in that all potentially relevant sources of information can be employed during initial processing (Pickering, in press; Traxler & Pickering, 1996a). Data about these sources of information can be obtained from corpus counts or from production tasks where subjects complete sentences from the point of ambiguity onwards. For instance, a few subjects might complete *The man realised* using a noun phrase (e.g., *his goals*), but more using an embedded sentence (e.g., *his goals were unattainable*). From this,

such models would assume that the parser would foreground the embedded-sentence analysis (e.g., Trueswell et al., 1993). This is as predicted by likelihood accounts. However, recent constraint-based models sometimes assume that the parser pays attention to broader classes in making decisions (e.g., Juliano & Tanenhaus, 1994; see Mitchell, Cuetos, Corley, & Brysbaert, 1995). It might note that more verbs are transitive than intransitive, say, and thus support a transitive analysis of a fairly rare intransitive-preference verb. Unless the verb were extremely rare (with the data about preferences being unreliable), such a heuristic would go against likelihood. Constraint-based models therefore approximate to likelihood, but may diverge from it, and are not directly motivated by it.

Other current models are restricted, in that only some sources of information may be employed during initial processing (e.g., Abney, 1989; Crocker, 1996; Ferreira & Henderson, 1990; Frazier, 1979, 1987; Gorrell, 1995; Mitchell, 1987, 1989; Pritchett, 1992). These models assume that initial parsing decisions are based on structural principles. They are in general serial. They do not approximate to likelihood, because a great deal of information that might be needed to determine likelihood is inaccessible. In many cases, the parsing assumptions support the construction of the structurally simplest analysis. For instance, Frazier's (1979) principle of minimal attachment supports the construction of the analysis with the fewest nodes in a phrase-structure tree. The models differ in some of their predictions, of course, but they overlap considerably.

In such models, the goal of initial processing is not to adopt the most likely analysis. As mentioned above, this is intuitively strange, in that it does not appear to approach optimality (Anderson, 1991). For instance, Frazier (1979) assumes that the parser adopts the minimal attachment analysis because it requires fewer computational steps than alternative analyses. The different analyses compete in a race, and the minimal attachment analysis wins. But since this analysis is not in general the most likely analysis, the system appears to be set up in a way that makes reanalysis more prevalent than necessary (with the average number of steps necessary to reach the correct analysis being higher than if the most likely analysis were adopted immediately). Why is it not designed to slow down a bit and obtain the most likely analysis first time? On the surface, restricted accounts such as Frazier's seem maladaptive in comparison to unrestricted accounts based on likelihood.

But there is experimental evidence against likelihood accounts, though it is highly controversial (e.g., Ferreira & Henderson, 1990; Mitchell, 1987). In the next two sections, we outline experimental results from our laboratory that provide strong evidence against likelihood accounts. It therefore appears that the parser operates in a fundamentally suboptimal way. In the final section, we

address this paradox, by suggesting that adhering to likelihood is not in fact as adaptive as it appears.

10.4 Evidence Against Likelihood from Local Dependencies

We now consider experiments suggesting that the processor does not adopt the most likely analysis, using subordinate clause ambiguities and complement clause ambiguities. In both cases, the ambiguity concerns the question of whether a noun phrase should be treated as an argument of a verb or not. Most structural accounts predict that the parser should prefer to attach a noun phrase as an argument if possible. In our experiments, this analysis is less likely than the alternative analysis, so likelihood models make different predictions. The next section reports a comparable experiment using unbounded dependencies.

As before, we manipulated the plausibility of an analysis that is ultimately shown to be incorrect. In addition, we kept the plausibility of the ultimately correct analysis constant. Hence any plausibility effects must be due to the parser selecting the ultimately incorrect analysis. If this analysis is implausible, then we predict immediate difficulty, but if it is plausible, then we predict difficulty at the point of disambiguation. Either one of these effects would provide a clear index of misanalysis. If misanalysis occurs in our experiments, then the parser must be initially considering an analysis that is not the most likely analysis. Parallel likelihood models may also experience problems, as we discuss below.

One important factor in a likelihood account is subcategorisation frame preferences. Consider (6a, b) (Pickering, Traxler, & Crocker, 1998):

(6) a. The young athlete realised her potential one day might make her a world-class sprinter.
b. The young athlete realised her exercises one day might make her a world-class sprinter.

The verb *realised* can take an NP object, as in *The young athlete realised her potential*, or a complement clause, as in (6a, b). It can also be used in other, irrelevant, ways, such as with a complementiser. Our pre-tests indicated participants more often write sentences using *realised* with a *that*-less complement clause than with an NP object. Another pre-test showed a similar preference when participants completed sentences beginning *The young athlete realised*. When (6a, b) are presented in isolation, these are the only cues available to participants. Hence, the serial likelihood model predicts that the processor should adopt the complement clause analysis initially if the sentence

is presented in isolation. The same prediction is made by Ford et al. (1982), where ambiguity resolution is determined by subcategorisation preferences alone.

Further pre-tests showed that the plausibility of the object analyses of (6a) and (6b) differed, since it is plausible to realise potential, but implausible to realise exercises, but that the plausibility of the complement clause analyses did not differ. However, eye-tracking data showed that readers found (6b) more difficult than (6a) before reaching the disambiguating word *would*, and that they found (6a) more difficult than (6b) after disambiguation. Hence subjects displayed the crossover pattern demonstrated in Pickering and Traxler (1998a). The only explanation of these results is that readers considered the object analysis of (6a) and (6b), even though this analysis was not the most likely analysis. Hence, serial likelihood is incorrect.

Pickering et al. (1998) present two other experiments that reinforce this conclusion. In one, (6a, b) were placed in short discourse contexts consisting of a title, an initial sentence, the target sentence, and a final sentence. The initial sentence mentioned the two nouns that were manipulated between conditions (here, *potential* and *exercises*). The results were similar, which indicated that the results did not reflect any special strategies that might be employed when a series of unrelated sentences are encountered in isolation. In the other experiment, subordinate clause ambiguities were employed, using the same design as (6a, b):

(7) a. While the pilot was flying the plane that had arrived stood over by the fence.
b. While the pilot was flying the horse that had arrived stood over by the fence.

Pre-tests showed that *flying* is preferentially intransitive, and also that participants preferentially complete *The pilot was flying* intransitively. According to serial likelihood, the following noun phrase should not be attached as its object. However, eye-tracking data suggested that it was attached: readers had more difficulty with (7b) before disambiguation, but more difficulty with (7a) after disambiguation.

Our results therefore present serious problems for serial likelihood accounts. Do parallel likelihood accounts fare any better? Consider (6) again:

(6) a. The young athlete realised her potential one day might make her a world-class sprinter.
b. The young athlete realised her exercises one day might make her a world-class sprinter.

After *realised*, the parser could consider both the object and complement analyses, but would foreground the more likely complement analysis. It might then integrate only the foregrounded analysis with background knowledge. If so, no plausibility effect is predicted at any point, because plausibility is only manipulated on the initially backgrounded and eventually abandoned analysis.

Alternatively, it might integrate both analyses with background knowledge. This is only possible if readers can think about two things at once (contra J.A. Fodor, 1983). Readers might experience processing difficulty when the parser reordered its preferences. This might happen in (6a) during the ambiguous region. After *The young athlete realised*, readers know that the probability of the complement analysis is higher than the probability of the object analysis. But after *The young athlete realised her potential*, the probability of an object analysis might be higher than the probability of a complement analysis, thus leading to reordering. But readers would not reorder preferences in (6b), because the object analysis becomes even more unlikely after *her exercises* (as *The young athlete realised her exercises* is very implausible on the object analysis).

Thus, the parallel likelihood account has two possible predictions about processing during the ambiguous region: either (6a) will be harder to process than (6b), or (6a) and (6b) will not differ. In fact, we found that (6b) causes difficulty in comparison to (6a). Hence parallel likelihood is very difficult to reconcile with our findings.

10.5 Processing Unbounded Dependencies: Evidence for Immediacy and Against Likelihood

The processing of unbounded dependencies has traditionally been treated rather separately from other work on syntactic ambiguity resolution (e.g., J.D. Fodor, 1978). In contrast, we suggest that they can be integrated into a unified theory (e.g., Pickering, 1994). One purpose of this section is to show how this can be done. Additionally, we provide evidence that unbounded dependencies are formed immediately, at the verb, even when this is not the most likely analysis. We have addressed two issues in the processing of unbounded dependencies. First, Traxler and Pickering (1996b) found that the processor formed and interpreted unbounded dependencies at the verb. Consider (1b) again, repeated here:

(1) b. That's the garage with which the heartless killer shot the hapless man yesterday afternoon.

Readers experienced difficulty with this sentence as soon as they encountered *shot*. However, the purported gap location is after *man*, as the canonical form

of the relative clause is something like *The heartless killer shot the hapless man with the garage yesterday afternoon*. Hence, readers do not wait until the purported gap location before forming the unbounded dependency, in accord with Pickering and Barry (1991) and Pickering (1993).

Some of our other work attempts to explicate the mechanisms of initial choice at a point of ambiguity and makes the same assumptions. Sentences containing unbounded dependencies are often locally ambiguous (J.D. Fodor, 1978). For instance, the fragment *We like the book that the author wrote* can continue as in (8a) or (8b):

(8) a. We like the book that the author wrote during the winter.
b. We like the book that the author wrote unceasingly about.

In (8a), *the book* forms the unbounded dependency with *wrote*; in (8b), it forms the unbounded dependency with *about*. Hence, the sentences are locally ambiguous after *wrote*. Fodor discussed various possible (serial) strategies that the parser might use, balancing the advantages of incremental processing against the disadvantages of excessive reanalysis. Current evidence indicates that the parser sometimes forms unbounded dependencies immediately. This rules out many cautious models in which the parser waits for fairly certain evidence before forming the unbounded dependency.

There is little doubt that readers form the unbounded dependency immediately if the verb (e.g., *wrote*) preferentially takes an argument of the same category as the filler. Traxler and Pickering (1996b) demonstrated that this occurred during normal reading, using (2a, b), repeated here:

(2) a. We like the book that the author wrote unceasingly and with great dedication about while waiting for a contract.
b. We like the city that the author wrote unceasingly and with great dedication about while waiting for a contract.

As discussed earlier, readers experienced more difficulty with (2b) than with (2a) from the first fixation on *wrote* onwards. Because the plausibility only differed on the misanalysis, readers must have formed the unbounded dependency at *wrote* and treated the filler as the object of *wrote*. Many other experiments suggest the same conclusion, using plausibility (e.g., Boland et al., 1995; Garnsey et al., 1989; Stowe, Tanenhaus, & Carlson, 1991), the "filled-gap" technique (e.g., Boland et al., 1995; Crain & Fodor, 1985; Stowe, 1986), or cross-modal priming (Nicol & Swinney, 1989). We also found a suggestion of the crossover effect, as (2a) may have been harder than (2b) after the disambiguating word *about*.

But the parser does not appear to form an unbounded dependency, even with transitive-preference verbs, if the dependency is rendered ungrammatical by island-constraint information (Ross, 1967). In the same experiment, readers also encountered sentences like:

(9) a. We like the book that the author who wrote unceasingly and with great dedication saw while waiting for a contract.
b. We like the city that the author who wrote unceasingly and with great dedication saw while waiting for a contract.

In (9), there is a potential unbounded dependency between the filler and *wrote*. However, it is ungrammatical, as the contrast between (10a) and (10b) shows:

(10) a. The author who wrote that book died.
b. *That book, the author who wrote died.

We found no evidence that readers formed the ungrammatical potential unbounded dependency in (9), in contrast to (2). This suggests that parsing is closely bound by the constraints of linguistic knowledge (see also Stowe, 1986; but cf. Clifton & Frazier, 1989).

What happens when the verb does not preferentially take an argument of the same category as the filler? According to serial likelihood, the processor does not form the unbounded dependency at the verb; according to parallel likelihood, this analysis may be constructed but is backgrounded. Pickering and Traxler (1998b) conducted two self-paced reading experiments to investigate this, using sentences like (11):

(11) a. That's the plane that the pilot landed (carefully) behind in the fog at the airport.
b. That's the truck that the pilot landed (carefully) behind in the fog at the airport.
c. Those are the lines that the actor spoke (briefly) about to the acting coach after the rehearsal.
d. Those are the props that the actor spoke (briefly) about to the acting coach after the rehearsal.

The adverb in brackets was included in one experiment and excluded in the other; the results from both experiments were similar. Sentences (11a, b) contained the verb *landed*, which preferentially takes an NP argument; they were essentially similar to (9a, b). Sentences (11c, d) contained the verb *spoke*, which preferentially does not take an NP argument. We concentrate on (11c, d) here.