UCLA Working Papers in Linguistics, no. 13, December 2005
Papers in Psycholinguistics 2, Okabe and Nielsen (eds.)

SUBJECT GAPS, A-BAR TRACES, AND PARALLEL PARSING*

CARSON T. SCHÜTZE
cschutze@ucla.edu

This paper discusses two issues related to the parsing of filler-gap dependencies, also known as movement relations. In the first half I suggest how to strengthen or refute the empirical basis for the claim that gaps (left by *wh*-movement) in subject position are processed differently from those posited in post-verbal positions. In the second half I show how, if this claim is true, it can be modeled in a parallel parsing framework, improving upon a proposal in the literature.

1. INTRODUCTION

The parsing of empty categories has long been a focus of research on the human sentence processing mechanism (HSPM). Since empty categories (gaps) do not occur in the input signal, the way they are processed provides good evidence about how the parser uses information of various types to deduce their presence, in particular, grammatical and lexical information that can help to predict where gaps could and could not occur. Existing experimental findings on gap-filling have led to a few different classes of theories of how the parser handles gaps. The basic finding that most such theories strive to account for is that in certain sentential configurations, the parser seems to adopt a preferred analysis containing a gap before the input signal plus the grammar allows it to verify the presence of a gap. Thus, in one way or another, most theories embody a preference for positing gaps as soon as possible.

An important qualification to the description of this phenomenon, however, is that no clear evidence has yet been found to show that this preference for gaps over lexical material extends to subject positions. In particular, the classic filled-gap paradigm (Crain and Fodor 1985; Stowe 1984) fails to expose filled gap effects for local subject extrac-

* This paper incorporates material written in the mid 1990s but never published, including Schütze 1994. I have not tried to systematically update that material in light of subsequent literature. I am making it available because I believe it contains suggestions that are not discussed elsewhere and might still be the source of a fruitful hint or two for those interested in these issues. The original work benefited from extensive discussions with Ted Gibson.

tions, and what little work has been done looking for such effects in long-distance subject extraction has also failed to find a gap preference. If it turns out to be true that no such preference exists, this would have substantial consequences for theories of gap parsing, and hence parsing in general, to the extent that the parsing of gaps is not attributed to an isolated module of the parser. Thus, the main goals of the present work are, first, to seek more empirical evidence about whether subject gap preferences really do not exist; second, to use the parser's treatment of subject gaps to narrow down the range of possible parsing theories; and third, to improve on the implementation of a parsing model that was proposed to capture some of the facts discussed. I should point out that although it is a parallel model of parsing, and I believe parallelism is ultimately the best model of human behaviour, in the earlier sections of the paper I follow the tradition in the literature of describing the parser's actions as if it were acting serially. This is for expository simplicity only; the discussion could be re-stated in parallel terms without affecting the arguments.

2. BACKGROUND ON THE PROCESSING OF SUBJECT GAPS

To date, there have been two sorts of responses to the apparent subject-nonsubject asymmetry with respect to gap preferences, namely, to treat it as real and attempt to derive it in a principled way, or to treat it as artifactual and assume that the parser fundamentally embodies no such asymmetry. In particular, Gibson, Hickok, and Schütze (1994) propose that early gap-filling is driven by the parser's desire to establish θ-relationships as quickly as possible, and since the θ-role of a subject (in languages where it is pre-verbal) is not available at the point when subject position is processed, positing early subject gaps would not help to meet this need. This should be compared with the situation once the (main) verb of the clause has been processed, at which point positing a gap in a θ-position determines the θ-role of whatever argument chain that trace is a part of.

In contrast, Clifton and Frazier (1989) suggest that the lack of local filled-subject-gap effects is an artifact of the rapidity with which an overt subject (when present) or a cue to the lack thereof follows the *wh*-word that would be its potential filler: typically this is the very next word in canonical English experimental items like *I wonder who Ruth will…* vs. *I wonder who will…* Thus, the processor does not have sufficient time to manifest a desire to posit a subject gap before the presence or absence of an overt subject becomes evident from the input signal. Their account of gap filling is therefore independent of position, stating that in the domain of a filler, positing a gap is *always* preferred where possible. Without further elaboration, this theory would presumably not lead one to expect the absence of long-distance filled subject gap ef-

fects, but it appears that no existing experiment was designed specifically to probe for such effects, so the empirical ground is less firm here. Furthermore, one could imagine elaborations of Frazier's (1987) Active Filler Hypothesis that could predict such an absence, if the Active Filler effect were clause-bounded and needed to be "renewed" by postulation of intermediate traces in a Comp position of each new clause, such that the proximity of an embedded subject to its most local Comp would become relevant.

Even the very brief discussion above should make it clear that what the parser does with subject gaps will have a deep-going impact on the account of filler-gap parsing. If we take the further methodological stance that properties of filler-gap parsing should be assumed to follow from those of parsing in general (until we find evidence to the contrary), then subject gaps become central to parsing theory *tout court*. The present work adopts this Occam's Razor assumption and hence seeks to establish what an adequate account of subject gap parsing should look like. The following section lays out why I consider the descriptive question to be an open one and how one might answer it, and then section 4 explores what the theoretical consequences of various possible empirical outcomes would be.

3. WHERE TO LOOK FOR MORE DATA

As already mentioned, the reasons for questioning the nonexistence of filled subject gap effects are 1) local extractions might not be processed quickly enough for their potential subject traces to be posited in advance of encountering overt subject material, which in most (if not all) previous studies has been the very next word after the *wh*-word filler; and 2) long-distance subject extractions have received very little experimental attention, and there are reasons to be skeptical about what findings do exist. Let us tackle these problems in order.

3.1. *Local subject extraction*

There are three basic strategies we can employ to overcome the alleged time-course confound for local subject extractions. The first would be to lengthen the filler expression such that expectations for its trace might be generable before the filler itself has been completely processed. This would involve using a multi-word *wh*-expression in the standard filled-gap paradigm, perhaps containing post-head material which would not be relevant for the parser to decide to predict its gap. Consider example (1).[1]

[1] I follow the convention of the processing literature in indicating the location of the empty element (here, trace) with an underscore.

(1) a. Mary wondered which friends that we met last summer Ruth would visit __ this Christmas.
b. Mary wondered which friends that we met last summer __ would visit Ruth this Christmas.
c. Mary wondered whether friends that we met last summer would visit Ruth this Christmas.
d. Mary wondered to which friends that we met last summer Ruth would talk __ this Christmas.

The question of whether a subject gap is preferred over a lexical subject is essentially the question of whether (1a) is harder to parse than (1b), the point being that the need for a NP A-bar trace could be known to the parser after processing *friends*, i.e., long before the subject position is actually constructed. One way to ask this question would be to compare the three-word regions *Ruth would visit* and *would visit Ruth* in (1a) vs. (1b). However, this comparison potentially introduces numerous confounds since these regions are structurally non-parallel. Instead, the standard filled-gap paradigm has compared the no-subject-gap condition (1a) to a highly similar structure where a subject gap would not have been *possible*. When a single-word *wh*-expression is used, this comparison can be made quite cleanly by replacing that *wh*-word with *whether* or *if*, which also unambiguously introduce interrogative clauses but do not involve a filler-gap relation. In an example like (1) this is not feasible, however, because the overt subject *Ruth* would occur much earlier in the non-filler sentence than in the filler-gap sentence, a difference that potentially matters to reading times, independent of possible gaps. The best we seem to be able to do is a control like (1d), which disambiguates against a subject gap by using a *wh*-PP, which could not have a subject trace (except perhaps in an extremely marked word order).[2] (1d) is still one word longer than (1a), but since the seven words preceding the point of measurement are identical, this might be hoped not to make much difference.[3] (At the ex-

[2] It might be argued that a non-NP *wh*-question creates a better control than a *whether/if* question anyway. If we compare reading times at *us* in the sentences *Mary wonders who Ruth will bring us home…* versus *Mary wonders whether Rule will bring us home…*, a longer reading time in the former might simply reflect the extra work required to maintain the filler in memory while processing *us*, even if no gap is entertained at the direct object position. In principle, a PP-filler would not be a candidate for the direct object position, hence should be unable to induce a filled gap effect, but would still need to be maintained in memory until its own trace can be posited. Unfortunately, this elegant logic might be defeated by the fact that, at least for argument PPs, it has been shown that their traces (or something functionally equivalent) can be established as soon as the licensing verb is processed—the parser does not have to wait for the corresponding linear position in the string (Pickering and Barry 1991; Gibson and Hickok 1993).

[3] Whether any difference it would make works for or against the search for a filled gap effect in (1a) depends on whether reading times become longer or shorter at this distance

pense of lexical parallelism, one could massage the control sentence to make it one word shorter.)

The use of a complex *wh*-filler also allows a second comparison, namely (1b) vs. (1c). Here we are able to keep the number of words matched, varying only the *wh*-word to create an unambiguous gapless question in (1c). If, in (1b), a subject gap were *dis*preferred relative to finding an overt subject, then one might expect *would…* to be read more slowly in (1b) than in (1c), where the subject has already been unambiguously identified. While such a finding would not be uninteresting, it would be equivocal with regard to gap preferences, since finding a relative slow-down in this position might simply reflect extra work that the parser must do to hook up the filler to the trace in subject position, work not required in (1c).[4]

We could also try to lengthen the *wh*-filler by adding material *before* its head noun, on the assumption that an initial *wh*-word will still signal the presence of a filler and allow the gap-finding mechanism to "rev up" while the rest of the NP is parsed. This would look like (2).

(2) a. Mary wondered which one of her friends Ruth would visit __ this Christmas.
b. Mary wondered which one of her friends __ would visit Ruth this Christmas.
c. Mary wondered whether one of her friends would visit Ruth this Christmas.
d. Mary wondered to which (one) of her friends Ruth would talk __ this Christmas.

The advantage to this version over (1) is that the subject or the filled subject gap appear immediately after the head of the *wh*-phrase, which is probably the statistically most likely location and hence perhaps where expectations would be maximized.

The second way that one might avoid the potential time-course confound would be to interpose additional material between the *wh*-filler and the subject position, again allowing time for 'trace prediction' to occur. There does seem to be a range of elements that can occur either

from the start of the sentence. One also needs to worry that pied piping carries some stiltedness that might slow readers down; this confound might be ameliorated by using a preposition that strongly disprefers stranding.

[4] However, if local subject extractions actually do not involve movement, then arguably such work would not be required. Thus, finding *no* difference between (1b) and (1c) would suggest that local *wh*-subjects are indeed in situ, and that a subject gap is not dispreferred compared to an overt subject. (Of course, it could also mean that these effects are present but not detectable, or that subject gaps are *preferred*.)

between a non-subject *wh*-expression and the subject or between a subject *wh*-expression and the verb.[5] An example is provided in (3).

(3) a. Mary wonders which friends in all likelihood Ruth will invite __ to the party.
b. Mary wonders which friends in all likelihood __ will invite Ruth to the party.
c. Mary wonders whether in all likelihood Ruth will invite someone to the party.
d. Mary wonders to which friends in all likelihood Ruth will talk __ at the party.
e. Mary wonders whether/when friends in all likelihood will invite Ruth to the party.

Obviously, one would want to try to minimize the stiltedness of the examples. Beyond that, a problem that both the paradigms in (1) and (3) share is that they differ from canonical object filled-gap environments because the gap site is not unambiguously predictable. That is, in a typical filled-gap example like *Mary wonders who Ruth will bring (us) home…*, the verb *bring* signals unambiguously that object position must follow immediately. However, in (1) and (3) there seems to be no way to unequivocally warn the parser that subject position is coming up next: in (1), additional adverbials could follow in the embedded clause after *last summer*, and in (3) additional material could follow *likelihood*, e.g., *according to Jane*, albeit thereby creating a stylistically distasteful muddle of a sentence. Of course, a parser that posits gaps as soon as they *might be* grammatically possible should not be put off by this uncertainty, but I am not aware of any existing evidence that active filler effects occur under such conditions, i.e., where not only the gap but the position itself is not assured to occur. A further potential problem when *which*-NP fillers are used is that there is a false gap position immediately following their head noun in examples (1) and (3):[6] that is, upon hearing *Mary wonders which friends*, a gap-seeking parser could posit that *which friends* has moved from subject position, a hypothesis that would be immediately disconfirmed in (1) by *that*. It is conceivable

[5] Not all adverbials are neutral with respect to the gap location. For example, ??*Mary wonders who probably Bill saw* is degraded relative to *Mary wonders who probably saw Bill*. One should gather off-line norms to try to avoid such biases in materials, hoping that off-line acceptability ratings would not themselves be affected by initial gap-positing preferences.

[6] Strictly speaking this problem also exists in standard *who* or *what* questions in the filled-gap paradigm: though stylistically marked, it is possible to modify these question words, e.g., *Mary wonders who of all her friends Ruth will invite….* Indeed, this could be a further factor contributing to the failure of filled subject gap effects to appear in previous experiments.

that after a first possible trace position turns out to be wrong, the parser becomes "gun shy"[7] about predicting further traces.[8,9]

The third approach to local subject extraction would be to try to get more directly at whether the parser is treating the filler as a subject, without trying to induce a filled-gap effect. This is a method that has been applied extensively to German and Dutch topicalization as well as *wh*-movement constructions. The idea is to disambiguate either by the agreement on the auxiliary verb following the filler, or else on the following noun phrase. (4) illustrates the first possibility.

(4) a. Which ideas that we thought of yesterday are the graduate students discussing __?
b. Which ideas that we thought of yesterday is the graduate student discussing __?
c. To which ideas that we thought of yesterday is/are the graduate student(s) subscribing __?
d. Which ideas that we thought of yesterday are they discussing __?
e. To which ideas that we thought of yesterday are they discussing a response __?

At the auxiliary, (4a) is consistent with a subject interpretation of the *wh*-phrase, while (4b) is not. Therefore, a subject preference should yield a slowdown on *is* in (4b) versus *are* in (4a). (4c) is a control for a possible low-level effect of number mismatch between a noun phrase and a linearly adjacent verb—since it is unambiguously a non-subject question, any difficulty on *is* versus *are* cannot be due to a garden path effect. Thus, there would ideally be no difference between *is* and *are* within (4c), or if there is a difference, it must be factored out of the (4a) versus (4b) comparison.

Furthermore, in (4a), if the *wh*-phrase is initially taken as a subject, one should expect a garden path effect when it becomes clear that *students* is actually the subject. It is not obvious when that would become evident in the general case. (Cf. the subject question *Which ideas that we thought of yesterday __ are the graduate students' favorites?* versus the nonsubject question *Which ideas that we thought of yesterday are the graduate students' discussions about __?*) The best reading time comparison might therefore be with a non-NP question, e.g. time spent on the word *they* in (4d), where it disambiguates against a subject ques-

[7] Thanks to David Pesetsky (p.c.) for this metaphor.
[8] Indeed, in (1) there is a second false subject position following *met* as well: twice bitten, thrice shy?
[9] I do not see any way to modify the examples to flag the end of the complex subject.

tion, versus in (4e), where a subject question was never a candidate to begin with.

One could also disambiguate more forcefully by using a case-marked pronoun in post-verbal position, as in (5).

(5) a. Which animal did them harm?
b. Which animal did they harm?
c. Who had them trapped?
d. Who had they trapped?

A potential confound in (5) is that the verbs *have* and *do* could exhibit a frequency-based preference for auxiliary over main-verb uses that might bias the interpretation of the filler before the disambiguating pronoun is processed. (This seems less of a worry with auxiliary versus copular *be* in (4), but it should be checked.) Also, the semantic features of the *wh*-phrase, such as animacy, probably play a strong role in whether it gets taken as a subject or not; therefore, bare *wh*-words as in (5c–d) might be preferable in this paradigm.

3.2. *Long-distance subject extraction*

Let us turn now to the question of preferences surrounding long-distance extraction (potentially) of a subject. One experiment that I am aware of potentially bears on this issue. Frazier and Clifton (1989) were looking for filled *object* gaps in long-distance extraction, which seemed to exist to some degree, so it appears that gap preferences do not disappear across clause boundaries, if their interpretation of the results is correct. However, they note in a footnote that all their examples potentially involved an embedded subject gap as well, yet no filled gap effect was found in this position. I do not believe we can put any stock in this statement, however, because of various problematic aspects of the paper. For one thing, on this very point the reading time data in their table do not seem to match their prose discussion thereof. For another, the comparison they tested is not ideal. Their paradigm was as in (6), where (6b) contains an overt subject that could act as a filled gap at the boldfaced word *she*; the other three sentences lack this property.

(6) a. Who did the housekeeper from Germany urge the guests to consider?
b. Who did the housekeeper say **she** urged the guests to consider?
c. The housekeeper from Germany urged the guests to consider the new chef.
d. The housekeeper said she urged the guests to consider the new chef.

The comparison they describe as nonsignificant is the following.[10] Compare the reading time for *say she* in (6b) with the reading time for *from Germany* in (6a); that difference will include any filled subject gap effect in (6b). Then compare the same two sequences in (6d) versus (6c), where there is no filler, hence no possible filled gap effect, but those sequences are otherwise in the same syntactic positions (introducing the complement clause versus modifying the matrix subject). If the (6d) minus (6c) difference is less than the (6b) minus (6a) difference, that must be because of the combination of the preceding *wh*-filler and the filled subject position. They take their failure to find such a difference to mean that "reading time did not seem to be disrupted by this [(6b)] false gap." But this comparison fails to take account of the possibility that the cost of holding onto the filler across the PP modifier of the matrix subject is not necessarily the same as for holding onto it across a verb and an embedded pronominal subject. (This might arise, for example, because the matrix subject θ-role throughout the region of interest in (6a) whereas it receives one at the first word of the tested region in (6b). Moreover, considering the full stimulus set one can see that the materials were extremely heterogeneous as to the syntax of the material in our region of interest, the position of the filled subject gap within a presented chunk of the sentence, etc.

Given the problematic, inconclusive nature of this experiment, it seems worthwhile to make further attempts to ascertain the parser's preferences in cross-clausal extractions. In addition to refining the kinds of materials used before, I suggest an additional technique that seems not to have been tried, which would use a comparison like that in (7).

(7) a. Who do you think Mary saw?
b. Who do you think that Mary saw?
c. When do you think that Mary saw Bill?

At least for speakers who exhibit standard *that*-trace effects, the presence of *that* in (7b) should unambiguously cue the absence of a subject gap.[11] Thus, if subject gaps are actively predicted, the time to read *Mary* should be longer in (7a) than in (7b): in (7a) it signals a filled gap, requiring the abandonment of the subject gap analysis, while in (7b) that abandonment already happened (by hypothesis) at *that*. (Again, if the extra word is of concern, (7a) could be padded with an

[10] According to their table, the difference in corrected reading time differences is 54 ms in the direction that a filled-subject gap effect would predict, which seems rather high to be nonsignificant.

[11] Using verbs like *think* that disallow matrix NP objects has the benefit of avoiding possible interference from that additional ambiguity, as would arise in *Who do you believe…*

adverbial like *probably* following *think*, though my intuition once again is that this has the side-effect of making a subject-gap reading more strongly preferred.) Conversely, in (7c) although a *wh*-phrase is waiting for a gap position, it is not a candidate for the embedded subject position, so *that* is not disambiguating against that option, whereas by hypothesis it is doing so in (7b), so the reading time on *that* should be longer in (7b) than in (7c).

4. IMPLICATIONS OF POTENTIAL RESULTS

Having now laid out the ways in which the presence or absence of a preference to posit gaps in subject positions could be more thoroughly investigated, let us consider the ramifications of the possible outcomes to these experiments. I catalog the major possibilities in (8).

(8) *Possible findings from subject gap experiments:*
   i. Both local and embedded extractions do trigger filled subject gap effects.
   ii. Long-distance extractions trigger filled subject gap effects, local extractions do not.
   iii. Local extractions trigger filled subject gap effects, long-distance extractions do not.
   iv. Neither local nor embedded extractions trigger filled subject gap effects.
   v. Local and/or long-distance extractions show a dispreference for subject gaps.

Let us consider how each of these scenarios might arise. Option (8i) would suggest that Frazier's alleged time-course confound was real, distorting previous studies, and that gap-positing in general exhibits no subject-nonsubject asymmetry. Option (8iv) suggests that this proposal is not the correct explanation for previous failures to find filled subject gaps, but that whatever drives post-verbal filled gap effects is not operative vis-à-vis subjects. Option (8iii) would suggest that an active filler strategy does apply in subject position but disappears across a clause boundary; one would then want to re-examine the purported evidence for cross-clausal filled object gap effects. Option (8ii) might suggest that local *wh*-subjects are in subject position, not moved to Spec-CP, hence not fillers in the relevant sense, but subject position does trigger active filling in general. Option (8v) (which might be hard to establish, but it would mean that encountering the absence of a subject is harder to deal with than encountering an overt subject) would suggest that putting a trace in subject position does not make the parser any happier than not positing the trace, and doing so furthermore is more work than dealing with an overt NP. This might seem unlikely prima facie, but as has been noted in the syntax literature (Hegarty 1990),

long-distance subject extraction often feels less felicitous than long-distance object extraction, which seems to be exactly the reverse of findings for local relativization, for instance.[12]

Thus, at a more general level, we can hope to establish whether subject gaps are or are not "actively filled" in the way that post-verbal gaps are. This clearly makes a big difference to the kind of account of active filler effects that can be entertained, since it immediately rules out one of two classes of theory: those that embody subject-object asymmetries or those that do not. In the category of asymmetrical active-filler accounts, one proposal has already been mentioned, namely θ-role resolution. There are other candidates here as well. For instance, focus-structure could be relevant: all other things being equal, subject position wants to contain old information, so subject questions could be dispreferred because the parser wants to find some old material before structuring new material. Not incompatible with this story would be a frequency-driven parser conditioned on the fact that subject questions occur less frequently than nonsubject questions.[13] The active filler effect in either case would then constitute trying to put a trace in a favorable position as quickly as possible, while not putting it in an unfavorable position unless forced to do so. A third kind of theory that could yield a subject-nonsubject split would be one in which predictions are driven by lexical frames of certain heads, including verbs. That is, post-verbal positions would be filled as quickly as possible if the lexical frame of the verb preferentially contains some element in that position, i.e., by positing a gap, while other positions would only be filled when necessary. (Indeed, most studies have found that filled gap effects disappear for objects of optionally transitive verbs that prefer their intransitive frame.) Since a subject can never be predicted by a lexical entry in this way, it would not be actively filled.

There are also several types of theory that would predict no asymmetrical behavior of subject gaps. The Active Filler proposal does this by directly stipulating the desire for early gap-positing whenever possible, but presumably this would follow from some general theory of memory load under which holding a *wh*-filler always incurs memory cost and positing its trace always relieves that cost, though one wants to know exactly how this cost metric would work in full generality. Another possible account would be one such as Pritchett's (1992), which

[12] It is possible that what makes many nonlocal subject extractions sound funny is a prosodic problem: with monosyllabic verbs, the resulting sentences necessarily contain main stresses on adjacent syllables, as in *Who do you think sáw Mary?* Introducing an unstressed syllable perhaps sounds better: compare *Who do you belíeve admíres Mary?*

[13] Although I am not aware of precise figures for this comparison, it strikes me as very likely to be true.

claims that satisfaction of (among other principles) Case Theory on-line is a driving force for the parser. Positing the tail of a filler's *wh*-chain results in the chain getting Case, in subject position as well as post-verbally.

Both within and across these two groups of theories, each one makes a different set of predictions for other types of filler-gap configurations that have not received experimental attention. (They also vary in the degree to which the principle needed to get the filler-gap facts is independently motivated from other parsing phenomena, an important consideration as well.) For instance, the theories split on what they predict about adjunct fillers: Case- and θ-theories predict no pressure to unload these quickly, while memory-based theories do predict such pressure, and lexically-driven theories might make differing predictions depending on whether particular verbs preferably occur with particular kinds of adjuncts. Argumental PPs make a different cut among the theories: they should be active fillers under a θ-theory or a memory theory but not under a Case theory, and again a lexical theory might make varying predictions. They should thus differ from adjunct PPs only under a θ-Theory. A-movement gaps could in principle distinguish Case- from θ-driven theories via their treatment of gaps as complements to passive verbs, for instance, but since these are much less ambiguous than A-bar movements anyway, we might not expect to find the same kind of phenomena at work at all.

All of these theories predict that questions and relative clauses should behave the same way with respect to their fillers and gaps, while a focus-driven theory might predict substantial differences here, as suggested in work by Kaan (1997) and Meng (1997). The theories also differ on their predictions concerning gap-first dependencies such as Heavy NP Shift (HNPS). These involve an open θ-role, but no element that lacks Case (though it might be relevant that there is a Case waiting to find the head of its assignee chain). One should also consider the possibility that the pressures of role-less fillers and filler-less roles are different, at least when it comes to long-distance dependencies. Whether HNPS should incur memory load under a Frazier-type theory depends on what units are relevant to such a theory—overt material vs. semantic relations vs. incomplete chains etc. Lexical theories might predict two loci of processing difficulty with HNPS structures: at the point where it seems a preferred argument has been omitted, and again when reversion to the preferred argument frame is forced, for preferably transitive verbs, but just one slow-down for preferably intransitive verbs, whereas θ-theory would predict no difficulty if the post-gap material is itself an argument of the verb as opposed to an adjunct (e.g., *John gave __ to his mother the scarf she had seen in the store window* vs. *John talked __ on Tuesday to the people he had met in China*).

In the remainder of this paper, I assume that further experimentation along the lines proposed above confirms the finding that there are no filled gap effects in subject position, and pursue a model that is intended to capture that result.

5. THE PARSING MODEL: MOTIVATION

Gibson, Hickok, and Schütze (1994) present a proposal for incorporating the processing of *WH* filler-gap dependencies into Gibson's (1991) parallel parsing theory. (I will not attempt to rigorously present those approaches, rather, I will informally describe just those concepts necessary for following the discussion, as they arise.) Here I propose some refinements to the implementation of Gibson et al.'s ideas. These refinements are motivated by three respects in which their original account seems unsatisfactory. The first is that it required an appeal to the concept "unambiguous head of a chain" as the primary object that incurred processing costs, which is not needed elsewhere within Gibson's parsing theory and is not motivated from the underlying grammar. While the notion 'head of chain' *is* an essential grammatical construct, Gibson et al. had to appeal to the status of a hypothesized NP-node as either unambiguously the head of a chain or potentially the tail of a chain. The latter possibility depended on the presence of a c-commanding potential antecedent (e.g., a *wh*-phrase) elsewhere in the sentence structure under construction, i.e. this was a non-local property. This notion will be eliminated under the new account. The second problem in Gibson et al.'s account was the way in which new traces were posited in the sentence structure. This was done by inserting them freely in the input stream (so that they would then be parsed as if they were audible words), which is conceptually somewhat odd, and goes against the intuition that there should be a difference between traces that are expected because a filler has been encountered versus those that must be posited in advance of any filler. (This intuition is presumably part of the motivation for Frazier's Active Filler Hypothesis.) My account will involve a different mechanism for trace-positing that embodies this intuition. The third problem was that, as stated, the new cost metrics made clearly wrong predictions with regard to rightward movement, e.g. HNPS. Since a trace of HNPS in direct object position could be freely posited and was not (it turned out) an unambiguous head of a chain, this account predicted that a structure containing such a trace should be favored over one in which a lexical (i.e., overt) NP object is hypothesized, because that *would* be the unambiguous head of a chain, hence would incur costs. This yields the intuitively implausible prediction that sentences involving HNPS should be easier to process at the word following the verb than non-shifted sentences. The new account that I will propose makes more plausible predictions about such cases.

There is one more respect in which Gibson et al.'s theory seems less than completely intuitive, namely the stipulation that cost assessment for partial structures (parses of an initial substring of the sentence) takes place only once for each input word that is processed and attached, and that this counting crucially precedes the formation of new chains (which could have the just-processed word as a member). However, this stipulation will be maintained in the current proposal, since it has the correct consequences within my framework of assumptions. One would hope eventually to eliminate this requirement, or else show that it follows from some more basic assumption. Perhaps some ideas from footnote 10 of Gibson et al. (1994) can be developed in this regard, but I leave that task for future work.

6. PROPOSED PARSING PRINCIPLES

There are several assumptions crucially required to make this account viable. In this section, I attempt to give them each some intuitive motivation. First, a general property of the parsing algorithm should be made explicit. I assume that in all cases where the grammar allows for an ambiguity between a trace and a (hypothesized) lexical element in a given position, such ambiguity is represented by creating separate structures in parallel, neither of which is itself ambiguous on this point. That is, one structure unambiguously contains a trace at the given location, the other unambiguously contains a hypothesized lexical phrase; there is no further allowance for the possibility of a trace in the latter structure. This eliminates the need for reference to an unambiguous head of a chain. What is needed now is to specify how these two structures are evaluated with respect to the cost metrics of Gibson's system. The general approach is to assess costs to structures that (temporarily) violate the θ-Criterion; the lowest cost (least penalized) structure is predicted to be preferred in situations involving ambiguity.

I assume that hypothesized nodes are treated no differently in this environment (i.e. when a *wh*-phrase c-commands) than anywhere else, that is, a hypothesized but empty node in a position to which a θ-role is assigned incurs a lexical requirement (LR) violation. (That is, it does not provide enough substance for the head to discharge its θ-role.) The new proposal is that a newly-posited trace is treated exactly the same way:

- A trace that is not co-indexed with any overt elements does not satisfy a lexical requirement of a verb, i.e. a θ-assigned position filled by such a trace incurs a LR violation.

Intuitively, the motivation for this is that a lexical requirement means that a head's θ-role must be associated with some overt element in or-

der to be considered "assigned," i.e., in order for the parser to know what that role is assigned to. A trace that is co-indexed with an overt *wh*-phrase can satisfy this requirement, since we can identify the relevant lexical material, but a trace on its own, not part of any chain, cannot. On the flip side, a trace in a position to which a θ-role is assigned does not incur a thematic reception (TR) violation, regardless of whether it has been linked into a chain, since the precise identity of its θ-role is known, the role-assigner having already been determined. (A TR violation is incurred by a chain that needs a θ-role but is not assigned one in the current structure.)

The two remaining assumptions concern the difference, alluded to above, between gaps that are "expected" by virtue of the parser's having encountered a *wh*-filler, versus gaps that are unexpected. Let us consider the former case first. Intuitively, we know after encountering a *wh*-phrase in an A-bar position that we must eventually posit a trace to receive the role of that phrase. I suggest that this intuition is directly implemented in the parser, in the following way:

- When a *wh*- or similar phrase is attached in an A-bar position to which it must have moved in a given partial structure, a trace must immediately be added to that same structure. The trace may either be attached to an existing tree, or a new stack node may be created for it. The latter option must always be pursued, while the former might be subject to restrictions imposed by the grammar.[14]

The idea here is to encode immediately on-line the fact that no valid sentence structure can contain a *wh*-phrase in a Comp without having a *wh*-trace associated with it. This may be seen as a redundancy, in that the grammar will eventually rule out structures that do not meet this requirement, but empirically it seems that the human parser takes a more proactive approach to ensuring the fulfillment of this requirement. If the only remaining parses at the end of a sentence have unattached traces in their stack representations, the parser knows immediately that no valid structure has been found, without attempting to verify that every *wh*-phrase has a valid associate, which could be a costly process. (In general, the stack contains lexical items and subtrees that have not yet been integrated into the matrix sentence.)

[14] I do not wish to commit as to the nature of those restrictions. It may be that as long as any attachment sites of the appropriate syntactic category are available, a trace will be initially attached in all such sites, regardless of eventual grammatical restrictions such as island constraints.

How are structures that are created in this way to be evaluated with respect to thematic requirements? I claim that nothing special needs to be said about this. A *wh*-phrase (at least of category NP) is a (single-membered) chain that requires a θ-role and hence incurs a TR violation if it is not in a θ-position and not associated (by chain formation/co-indexing) with a trace in such a position. A *wh*-trace is also, itself, a (single-membered) chain that requires a θ-role, and if it is in a stack node by itself it certainly is not assigned a θ-role, hence it too incurs a TR violation. Even though the *wh*-filler prompted the introduction of the trace into the structure, they cannot possibly be part of a single chain until they are at least within the same tree, and then they must still be co-indexed. I assume that they do not start out co-indexed, perhaps because co-indexation across different stack nodes is meaningless.

Now we must consider the other environment for trace-positing, a gap encountered in the *absence* of a previously-encountered filler, as in HNPS. I assume that traces of A-bar movement (but probably not of A-movement) can be freely posited at any stage of parsing a sentence, i.e. after each new input item is processed.[15] However, by the same intuition appealed to immediately above, the parser should embody the intuition that positing a gap "out of the blue" entails that a filler for that gap must eventually be processed. Paralleling the implementation above, I propose to capture this "expectation" for a filler by adding a hypothesized node of the appropriate category to the current representation. On the assumption that string-vacuous rightward movement does not occur,[16] this will always involve creating a new stack node for the hypothesized constituent.[17]

- *Wh*-traces are freely posited (perhaps in all hypothesized positions). In a given possible parse, if a new *wh*-trace is added, a hypothesized empty constituent of the same category must also be immediately added to that stack.

How are costs to be counted in this instance? Again, I claim no special stipulation is required. The trace is treated just like the other traces mentioned above, i.e. it cannot satisfy a lexical requirement, but it does not incur a TR violation either. In contrast, the hypothesized category that occupies its own stack node clearly cannot be getting a θ-role from anywhere, hence it does incur a TR violation.[18]

---

[15] Again, I do not concern myself with what the constraints on initial A-bar trace sites might be.

[16] This assumption merely simplifies the exposition; nothing crucially rests on it.

[17] In principle it could be attached to an existing partial structure elsewhere on the stack, but in practice I cannot envision a case where that would be possible.

[18] This is inconsistent with Gibson's (1991) original proposal that hypothesized arguments do not incur TR violations. One way to address this would be to posit some new

Taking the two preceding sets of assumptions together, we see that they both represent an implementation of the idea that the parser cannot process one end of a A-bar chain without immediately incorporating the other end somewhere in its working partial representation. (The status of intermediate *wh*-traces remains unclear.) I would expect the same to be true of chains of NP-movement.

The final assumption, which I will not attempt to motivate further, is essentially carried over from Gibson et al. I spell it out here for the sake of explicitness.

- When a *wh*-trace is first attached in a larger structure, that structure is evaluated (with respect to the properties of Thematic Reception and Lexical Requirement) before any chain formation (co-indexation) involving that trace is performed, and the structure is not re-evaluated after chain formation before the next input item is processed.

7. APPLICATION TO CRUCIAL CONSTRUCTIONS

Gibson et al. (1994) discuss four construction types whose processing they wish to account for; I refer the reader to that paper for the relevant data. In this section I will show how my account handles these four constructions, plus a fifth, namely HNPS.

7.1. *Preference for subject relatives over object relatives in Dutch*

Consider the word-by-word parsing of the Dutch literal translation of a sentence like *The boy who the woman saw…*, which in Dutch is globally ambiguous between a subject and an object reading for the relative pronoun. We wish to model the observation that speakers are more likely to adopt the subject reading.

At the point of encountering and attaching a relative pronoun in Spec-CP, a *wh*-trace must be added to the structure. There are two ways this can be done: the trace can be placed in Spec-IP (since the possibility of an empty C head entails the prediction of IP), but nothing forces

---

kind of element which is not merely hypothesized but required, although its contents are not yet known. This is ad hoc and unsatisfying. Instead I propose the following: hypothesized NPs *do* need θ-roles, but when they are in θ-positions, they in fact *get* their θ-role immediately, so no violation is incurred. A hypothesized A-bar-moved NP, however, is alone in a stack node and not in a θ-position, hence if it needs a role it clearly cannot get one and thus incurs a TR violation, which will turn out to be the desired result. This modification predicts that if there are ever situations where an NP is predicted in an A-position that is not a θ-position, it should incur a TR cost. This might happen for a subject in a VSO language.

it to be there, so the other option is to add it to the stack. We thus have two possible partial structures. (In this and following examples, "h" represents a hypothesized node, "t" is a trace, "WH" and "NP" are overt phrases, and separate lines represent different stack nodes.) In (9a), the first stack node contains a tree with the *wh*-phrase in Spec-CP and a separate stack node consisting of an NP-trace. In (9b) there is just one stack node, consisting of a tree with the same CP-layer but this time the trace has been attached into Spec-IP. Following the stipulation above, there is initially no coindexing relationship established between the potential chain members in (9b).

(9) a. $[_{CP}$ WH $[_{IP}$ $[_{I}$ h]…
t_{NP}

b. $[_{CP}$ WH $[_{IP}$ t_{NP} $[_{I}$ h] …

These parses both have the same cost, namely 2 violations. They each involve TR violations for the *wh*-phrase and for the trace. Therefore no preference is predicted at this point. After this evaluation is performed, co-indexation can apply in the (9b) structure, but no re-evaluation takes place. Then a lexical NP is encountered. It can be attached as a subject in the (9a') structure, but in the (9b') structure it can only be put onto the stack: it will turn out to be an internal argument, but we have not yet constructed a VP into which it can be attached.

(9) a.' $[_{CP}$ WH $[_{IP}$ NP $[_{I}$ h]…
t_{NP}

b.' $[_{CP}$ WH_i $[_{IP}$ t_i $[_{I}$ h] …
NP

Now there is a cost difference: the (9a') structure has 3 TR violations: the *wh*-phrase, the trace, and the subject NP. The (9b') structure, now with a complete *wh*-chain, has only 2: one for the *wh*-chain,[19] one for the overt NP. Thus, the subject-relative structure is slightly preferred, as desired. Note that this account differs from that of Gibson et al. in predicting that the preference emerges as soon as the ambiguous lexical NP is processed, rather than at the end of the sentence. In this respect, it agrees with the predictions of the Active Filler Hypothesis. It remains to be seen whether empirical tests can distinguish these possibilities.

[19] I assume the definition of TR as given in (6) of Gibson et al., stated in terms of chains.

7.2. *Absence of filled subject gap effects in English*

Here we derive the fact (assuming it is verified) discussed at length in sections 2 and 3. The point at which a filled-subject gap effect *could* occur in an English sentence is essentially that represented by the pair of structures in (9a–b) above (except that there might be an auxiliary verb in Comp). That is, until we discover the presence or absence of a lexical subject (at which point there is no more ambiguity), we are entertaining either a trace in subject position or an empty subject position and a trace on the stack, and there is no difference in cost between these structures, as desired. Of course, once the next word has been input, chain formation will lower the cost of the subject-trace structure, but by then it is too late: nothing can intervene between the "inverted" auxiliary and the lexical subject, if there is one.[20]

7.3. *Filled object gap effects in English*

This is the canonical filled postverbal gap effect. At the point of attaching the verb, the following two structures are possible:

(10) a. Wh… V $[_{NP}$ h]
t_{NP}

b. Wh… V t_{NP}

That is, seeing the *wh*-phrase in Comp forced a trace to be added to the stack. When a transitive verb is input, it hypothesizes an NP complement which can match that trace, yielding the (10b) structure. However, this match is not forced; if the hypothesized NP node is left empty, the trace remains on the stack, as in (10a). Costs are as follows: In (10a), there are 2 TR violations (*wh*-phrase and trace) and 1 LR violation (the hypothesized overt NP), for a total of 3. In (10b), there is 1 TR violation (*wh*-phrase) and 1 LR violation (the verb's internal argument), for a total of 2. Recall that the trace in a θ-position does not incur a TR violation. Thus, the (10b) structure with an object gap is slightly preferred over the (10a) structure with no object gap, as desired.

[20] An interesting prediction might ensue for cases where the next word is categorially ambiguous between the first word of an NP versus some other category that can follow a subject gap. At that point, the subject-gap reading should become preferred. A possible example would be *Mary asked which man fires frighten* __ vs. *Mary asked which man* __ *fires lazy workers*. Obviously, the relative frequencies of the senses of the ambiguous word need to be factored out.

7.4. *Reduced relatives in English*

These include sentences like *The horse raced past the barn fell* and *The children taught by the Berlitz method learned faster*. It is unclear how difficult these really ought to be once lexical biases are factored out, [21] but let us see what sort of result we can get. I crucially assume, following a suggestion by Alec Marantz (p.c.), that reduced relatives involve A-movement from the internal argument position to the subject position within the relative clause. Then we have two choices: the moved element could be an operator, in which case it subsequently undergoes A-bar movement to a Spec-CP position, or it could be PRO, in which case it moves no further. Take the operator option first. At the verb, the reduced relative structure looks as follows:

(11) the children OP t_{WH} taught t_{NP}

In keeping with the story so far, the OP and its trace have not yet been co-indexed. As it stands, this seems to predict too many violations in this structure. In addition to the matrix subject TR (*the children*), we have a TR for the operator and a TR for the *wh*-trace, since neither is in a θ-position, plus potentially a LR for the patient role of *taught*, since it is not tied to any lexical material at this point. This is clearly an undesirable result. Not wishing to abandon any existing assumptions, I propose to fix this account by adding new claims about A-traces as distinct from A-bar traces.[22] Specifically, I propose that an A-trace is immediately co-indexed with its antecedent before evaluation of the newly-posited structure takes place, unlike a *wh*-trace. That leaves this structure with 2 or 3 violations: TR for *the children* and TR for OP; since the two traces are now co-indexed by stipulation, the higher trace gets a θ-role and does not incur TR, but the LR of *taught* might still count, since the patient role is associated with two traces but with no "lexical" element. (As noted above, the null operator itself will count as lexical for this purpose, but it is not yet indexed.)

Consider now the alternative structure with PRO:

(12) the children PRO_i taught t_i

---

[21] One way to get around these lexical biases in principle might be to give readers/listeners sentences with novel content words that have no lexical properties (so-called Jabberwocky sentences), e.g.

(i) The wug gorped with/in the zoog parged.

Of course, this introduces a host of new issues as well.

[22] I believe one can motivate these claims on the basis of two differences between A- and A-bar movement: A-movement is more strictly local, and the head of an A-chain, unlike a Wh in Comp, does not immediately signal the need for a trace further down in the structure. More specifics must await me thinking them through.

Here we see that the load is only one violation: *the children* lacks a role, but the PRO-chain receives its role via the co-indexed A-trace in object position, and this chain must be assumed to satisfy the lexical requirement of *taught*. Thus, the PRO analysis predicts a slight preference for the main clause reading of reduced relatives, whereas the operator analysis predicts a strong garden path. Both analyses require a contrast between traces of A- versus A-bar movement to get the story to work out. This is not inconsistent with the underlying grammar, since such traces (and/or the movements they encode) have substantially different properties. This account obviously makes strong predictions about processing differences between these kinds of traces, which ought to be tested.[23]

7.5. *Heavy NP Shift*

I return finally to one of the motivations for this proposal, namely an account for the presumed slight difficulty of processing a sentence with HNPS, e.g. *John saw __ in the park a man who looked vaguely familiar*, as compared to an unshifted counterpart. Since there are no *wh*-fillers in this sentence, we must appeal to the proposed free gap-positing mechanism. Thus, at the point of processing the verb, the two alternative structures are as follows:

(13) a. John saw [$_{NP}$ h]

b. John saw t_{NP}
[$_{NP}$ h]

Structure (13a) has a cost of 1 LR for the object role. Structure (13b) has that same violation, since traces of A-bar movement that are not part of a chain do not satisfy a lexical requirement. In addition, it has a TR cost associated with the hypothesized NP on the stack. As a result, there is a slight preference for (13a), so we predict a slow-down at the word *in* in the HNPS sentence, a sort of "unexpected gap" or "unfilled argument" effect, but not a strong garden path.

8. CONCLUDING REMARKS

A careful examination of a variety of filler-gap constructions holds the promise of substantially narrowing the range of general parsing theories

[23] One prediction that might fall out and seems plausible is that the parser would *never* consider structures containing A-chains that violate locality constraints: a trace of NP-movement must immediately be indexed with a valid antecedent, or else the structure is discarded. Thus, one could look for early signs of rejection in examples like **John seems that Mary likes __* or **John seems that it is likely __ to leave*.

that should be pursued. In this paper I have suggested that the parser's behaviour with respect to (potential) gaps in subject positions has been under-studied and could be a particularly rich source of data.

REFERENCES

CRAIN, STEPHEN, and JANET DEAN FODOR. 1985. How can grammars help parsers? In David R. Dowty, Laurie Karttunen, and Arnold M. Zwicky (eds.), *Natural Language Parsing: Psychological, Computational, and Theoretical Perspectives*, pp. 94–128. Cambridge: Cambridge University Press.

FRAZIER, LYN. 1987. Syntactic processing: Evidence from Dutch. *Natural Language and Linguistic Theory* 5, 519–559.

FRAZIER, LYN, and CHARLES CLIFTON JR. 1989. Successive cyclicity in the grammar and the parser. *Language and Cognitive Processes* 4, 93–126.

GIBSON, EDWARD. 1991. *A Computational Theory of Human Linguistic Processing: Memory Limitations and Processing Breakdown.* Ph.D. dissertation, Carnegie Mellon University.

GIBSON, EDWARD, and GREGORY HICKOK. 1993. Sentence processing with empty categories. *Language and Cognitive Processes* 8, 147–161.

GIBSON, EDWARD, GREGORY HICKOK, and CARSON T. SCHÜTZE. 1994. Processing empty categories: A parallel approach. *Journal of Psycholinguistic Research* 23, 381–405.

HEGARTY, MICHAEL. 1990. On adjunct extraction from complements. In Lisa L. S. Cheng and Hamida Demirdash (eds.), *Papers on Wh-Movement, MIT Working Papers in Linguistics* 13, 101–124.

KAAN, EDITH. 1997. *Processing Subject-Object Ambiguities in Dutch.* Ph.D. dissertation, Rijksuniversiteit Groningen.

MENG, MICHAEL. 1997. *Die Verarbeitung von W-Fragen im Deutschen: Präferenzen und Reanalyseeffekte.* Ph.D. dissertation, Universität Jena.

PICKERING, MARTIN, and GUY BARRY. 1991. Sentence processing without empty categories. *Language and Cognitive Processes* 6, 229–259.

PRITCHETT, BRADLEY L. 1992. *Grammatical Competence and Parsing Performance*. Chicago: University of Chicago Press.

SCHÜTZE, CARSON T. 1994. A-bar traces in a parallel parsing framework. Ms., MIT.

STOWE, LAURIE A. 1984. *Models of Gap-Location in the Human Language Processor.* Ph.D. dissertation, University of Rochester. Distributed by Indiana University Linguistics Club.