# an FMRI *investigation* of intermediate traces

**Emily Finn**
*Advisor*: **Dr. Maria M. Piñango**

May 1, 2009

*Department of Linguistics, Yale University*

2008-2009

**Table of Contents**

*The marvelous thing is that even in studying linguistics, we find that the universe as a whole is patterned, ordered, and to some degree intelligible to us.*

- Kenneth L. Pike

Acknowledgments

First, on the technical side of things: at the Yale Magnetic Resonance Research Center, I would like to thank Jeetu Bhawnani for generous programming assistance, Hedy Sarofin for cheerful data acquisition, Cheryl Lacadie for impeccable data analysis, and Todd Constable for invaluable direction. I look forward to future collaborations with you all.

Well, if this thesis is supposed to represent the culmination of my undergraduate academic experience, these acknowledgments will be inherently incomplete, as my personal and intellectual growth over the past four years has been touched by a great number of incredibly smart, talented and fantastic people. But here, some especially salient ones:

*To Geordie Kaytes:* Thanks for knowing me better than I know myself sometimes.
...Make that most of the time.

*To my subjects, a group that includes nearly all my friends at this point:* Thank you for so willingly submitting to 80-plus minutes of long and complicated sentences inside an MRI machine. And for listening to me obsess about my research (but not in a way that would exclude you from my study).

*To my fellow linguistics majors in the class of 2009:* Thanks for your comments, camaraderie and comic relief (and the occasional hookah session). I look forward to seeing what you all do in years to come.

*To Yovanna Cifuentes and Ilkyu Kim:* Thanks for being the best lab-mates a girl could ask for!

*To Steve Anderson:* Thanks for editing that book with that article by Chomsky...you know, the one from 1973 or whatever? No, in all seriousness, thanks for your insightful critiques, down-to-earth guidance and Wednesday afternoon beers – oh, and for the extension.

*To Maria M. Piñango:* Where do I begin? You were so much more than a senior essay advisor. Thank you for ridiculous analogies (NHITBWTBC – I can't believe I'm actually thanking you for this) and late-night email sessions. Thank you for trusting me with important things, so that I could learn to trust myself. Thank you for never, ever letting me settle. For all your brilliance, kindness and heart, there are no words (and we're linguists, so that's really saying something). We are so not done.

*To my family:* You probably still have very little idea of what linguistics is and why I chose to study it...all I can say is, join the club. But thanks for being supportive anyway. I love you.

**Introduction**

Language is a biological system. As a product of human evolution, it is part of the natural world, subject to scientific inquiry.

The goal of any theory of linguistics, then, should be to model what people do when they use language. We know that the language system must do more than simply combine the words it hears, one by one, in a flat linear sequence. This intuition comes from ideas about constituency structure. It is what makes certain Groucho Marx quotes funny:

(1) Yesterday, I shot an elephant in my pajamas. How he got into my pajamas, I'll never know.

That the same exact sentence can give rise to two very different interpretations – disambiguated only by the second sentence – tells us that there must be some kind of covert structure connecting words and phrases. As we speak or listen, we subconsciously build this structure, which we call syntax.

In this paper, we investigate a phenomenon called intermediate traces. This phenomenon has important implications for syntactic processing as it occurs in real-time. Our experiment hinged largely on the contrast observed in the following sentences:

(2) The captain who the sailor predicted that the weather would frighten turned back towards port.
(3) The captain who the sailor's prediction about the weather had frightened turned back towards port.

At first glance, these sentences look nearly identical. They are matched in terms of length and number of words, and they mean (almost) the same thing. Both contain a long-distance dependency, with *the captain* as the displaced object/experiencer of the verb *frighten.* And yet, they are structurally distinct. This is a contrast we will be returning to throughout this paper. We refer to them as VP extraction and NP extraction, or conditions (a) and (b), respectively.

According to some syntactic theories, the sentence in (2) contains an intermediate trace at the boundary between the two embedded clauses – that is, between *predicted* and *that.* This intermediate trace is absent at the corresponding linear site in (b), between *prediction* and *about.* We will return to what exactly an intermediate trace is, and why certain theories postulate it, in the next section. Our main question was, do intermediate traces play a role in online sentence parsing?

We chose to investigate sentence parsing via functional magnetic resonance imaging (fMRI). FMRI allows for observation of neural activity associated with different tasks. We have reason to believe that syntactic processing happens in the left inferior frontal gyrus (IFG), also known as Broca's area. (Evidence for this comes from observations in brain lesion studies as well as previous fMRI experiments, discussed at length in section five.) So, if there is really a structural difference between (a) and (b) that manifests itself in processing, we might expect it – the intermediate trace – to be visible at that particular moment in the left IFG.

These two sentence types, along with two additional conditions – a violation and a control – formed the basis for our experiment. We used a paced reading paradigm that allowed for comparison of sentence-internal events between conditions. Ultimately, we did find significant differences in left IFG activation for sentences (a) and (b). We attribute these differences to their syntactic composition – specifically, the presence of intermediate structure.

The first section of this paper will examine the original evidence motivating intermediate traces in Government and Binding Theory (GB). The second takes a look at how long-distance dependencies are treated in alternate theories of syntax, namely Simpler Syntax (Culicover & Jackendoff, 2005) and its borrowings from Head-driven Phrase Structure Grammer (Pollard & Sag, 1994). In section three, we reanalyze both formal theories from the perspective of the language processor to see what predictions each makes about *wh*-relatives and filler-gap dependencies in real-time parsing.

The fourth section discusses the psycholinguistic evidence for intermediate structure, focusing on a self-paced reading study conducted by Gibson and Warren in 2004. The fifth section uses evidence from both aphasia and fMRI to present a working hypothesis for what it is that the left IFG does, and why we are looking for intermediate traces there. The sixth section describes our experimental procedure. The seventh presents our (still preliminary) results, and the eighth discusses their implications. Finally, section nine concludes with ideas for future study.

## 1. Intermediate Traces: The Theory

Where – and more importantly, why – did the idea of intermediate traces originate? This section examines the theoretical evidence for intermediate structure.

### 1.1 *Government and Binding Theory*

Government and Binding's transformational syntactic theory holds that extractions across multiple clauses are mediated by an intermediate trace.

In his 1973 "Conditions on transformations," Chomsky proposed the grammatical principle of Subjacency. According to Subjacency, a phrase may only cross one bounding node in each step in its derivation from D-structure to S-structure. In English, NP and IP are bounding nodes. In order to establish a dependency through two clauses, therefore, a phrase must "stop" in the specifier (Spec) position of the intervening CP before moving up to the Spec position of the matrix CP. This movement in steps is known as *successive cyclicity*. If the Spec position of the intervening CP is filled, the phrase cannot stop there and the movement takes place in one long step, crossing two IP nodes and rendering the sentence ungrammatical.

Subjacency and successive cyclicity explain why certain long distance dependencies are grammatical and others are not. For example, contrast (4) and (5):

(4) Who$_i$ did the teacher claim $t_i$ that the idea had inspired $t_i$?
(5) *Who$_i$ did the teacher wonder which idea had inspired $t_i$?

In *(5), SpecCP of the intervening clause is already filled by the NP "which idea," so it blocks movement up the chain by the relative pronoun "who." Examine the structures in (6):

(6)

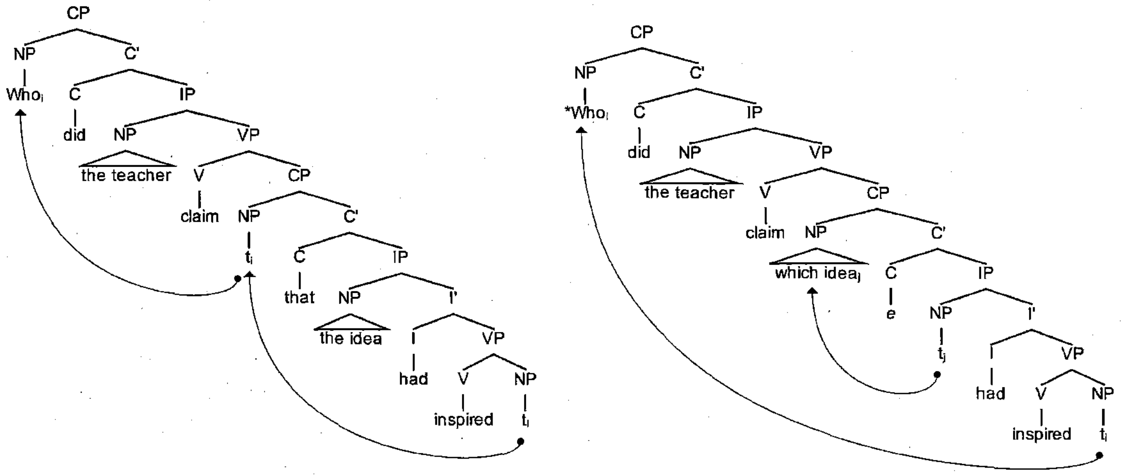


Essentially, intermediate traces block the relativization of more than one *wh*-word. They are also triggered in statement contexts, like the following:

(7)

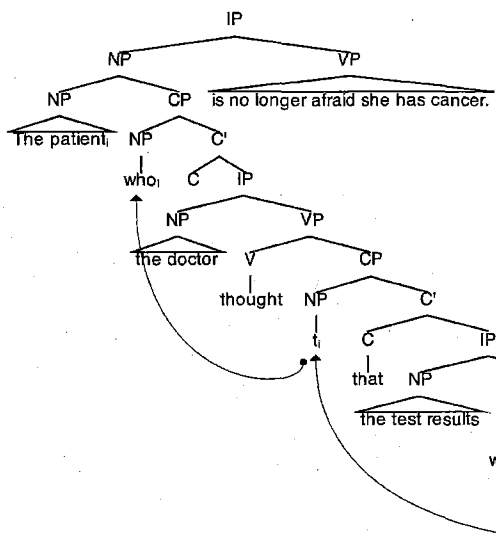


Nearly all mainstream generative syntactic theories follow GB's convention of postulating successive cyclicity and intermediate traces.

1.2 *Linguistic evidence*

Morphology in other languages has been used as linguistic evidence for intermediate traces. In Chamorro, it seems that each trace triggers agreement between the moved element and the immediately higher verb (see Chung, 1982, 1994). So in the case of two-clause *wh*-questions like (4), both the embedded and the matrix verb show agreement with the moved *wh*-word. A theory with intermediate traces predicts this result, while a theory without intermediate traces predicts agreement only for the embedded verb.

Additionally, McCloskey (2000) uses quantifier-stranding in Irish English to support the notion of intermediate traces. Examine the data in (8):

(8) a. What all did he tell him (that) he wanted t?
b. What did he tell him all (that) he wanted t?
c. *What did he tell all him (that) he wanted t?
d. ?What did he tell his friends/Mickey all (that) he wanted t?
e. *What did he tell all his friends/Mickey (that) he wanted t?

*(McCloskey 2000)*

As expected, 8(a) is grammatical, because the quantifier is fully pied-piped along with the *wh*-word to the matrix SpecCP. In 8(b) – also grammatical – McCloskey argues that the quantifier is stranded in the intermediate SpecCP position as the *wh*-word makes its way to the matrix SpecCP, supporting a movement-based theory. A quantifier in a location other than the matrix or intermediate SpecCP, as in 8(c) and 8(e), renders the sentence ungrammatical, suggesting that there is something structurally 'special' about the intermediate position – namely, that it is part of the movement chain. (McCloskey offers that 8(d) is questionable as compared to 8(b) not for any particular syntactic reason, but rather because of prosody.)

Theoretical evidence for intermediate traces as a result of movement is certainly available (see Georgeopoulos 1985 for evidence from Palauan and Torrego 1984 for evidence from Spanish). However, as we will see in the next section, not all theories posit intermediate traces – because not all theories rely on movement to account for long-distance dependencies.

## 2. Alternate Theories: No Movement

### 2.1 *Simpler Syntax Hypothesis and Head-Driven Phrase Structure Grammar*

In their 2005 book, "Simpler Syntax," Culicover and Jackendoff present a unification theory of language processing. Their account cuts down on much of the bulk of transformational syntax and instead reflects an enriched syntax-semantics interface. In contrast to GB, which considers traces as byproducts of syntactic movement, Simpler Syntax (here referred to as SS) places them in direct association with their antecedents, subject to both syntactic and semantic constraints.

*2.1.1 Syntactic constraints.* For *wh*-questions, Culicover and Jackendoff present a sequential mechanism of an operator + variable complex followed by indirect licensing between the operator + variable, the *wh*-word and the syntactic trace position. In the first step, a question operator triggered by the *wh*-word binds the variable in its base position. In turn, both the operator and the variable are linked to the *wh*-word and the syntactic trace – or "target." This creates a two-to-two mapping between CS – the conceptual structure – and syntax. The representation of a simple *wh*-question, under SS, might then look something like this:

(9)

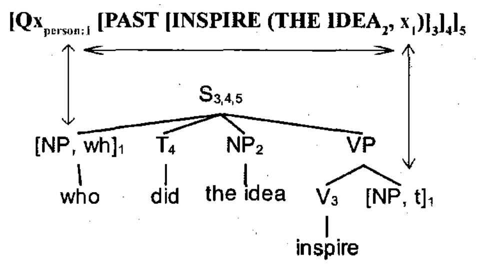


Extrapolating, SS has no need for intermediate traces mediating the movement of *wh*-words, because *wh*-words do not actually originate in lower clauses and thus do not move. Therefore, there is no theoretical reason why *wh*-questions across multiple clauses should be structurally distinct from single-clause questions. So our example from GB of a *wh*-question with an intermediate trace would look like this under an SS reanalysis:

(10)

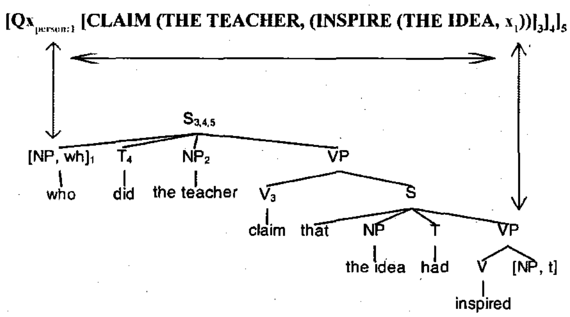


In fact, Culicover and Jackendoff are unequivocal in their stance on the matter: "Since there is no movement, there are no intermediate traces." (Culicover & Jackendoff, 2005, pp. 302)

As we saw above, GB's intermediate traces occur not only in *wh*-questions but also in statements. For these situations, Culicover and Jackendoff adopt *slash categories,* a convention of Head-Driven Phrase Structure Grammar (HPSG) originally proposed by Pollard and Sag (1994). A slash is a feature of all nodes dominating a trace position. In Culicover and Jackendoff's words:

> (11) A slash-category XP/$t_{YP}$ designates an XP which dominates somewhere within it a trace of category YP.

The slash originates at the antecedent, which is any argument that does not immediately receive a theta-role. It is transferred down the tree until theta-role assignment takes place at the gap. In other words, the gap is the lowest slash-category, and when it is reached it satisfies (or "saturates") the slash feature of all the nodes above it. In tree form:

(12)

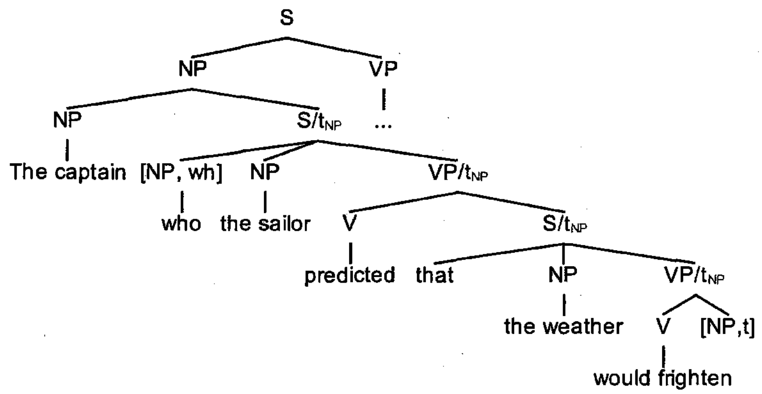


There is no limit to how long the slash chain can be (except, one might imagine, whatever limit is imposed by working memory). According to Culicover and Jackendoff, the "price" for this theory is enrichment of the syntactic category inventory to include slash categories. They justify this enrichment based on ideas about processing: "The passing of the slash-feature down the tree might be taken to represent [the] anticipation of a gap, and to result in increased working memory load." (Culicover & Jackendoff, 2005, pp. 331)

So how, then, does SS/HPSG explain why some extractions are grammatical and others are not? In response to some of the evidence that originally motivated Subjacency, Culicover and Jackendoff (via Pollard and Sag) propose two constraints forbidding the passing of slash-features in certain syntactic configurations, namely between S and NP nodes:

(13) Complex NP Constraint
*NP/t (e.g. *who did you deny [$_{NP}$ the claim [$_S$ that Amy likes t]])
|
S/t

(14) Sentential Subject Constraint
*S/t (e.g. *[$_S$ who did [$_{NP}$ a picture of t] fall on Harry])
|
NP/t

*(Culicover & Jackendoff, 2005, pp. 332)*

*2.1.2 Semantic constraints.*

In addition to these syntactic mechanisms, Culicover and Jackendoff propose semantic restrictions on long distance dependencies. Building on evidence originally presented by Erteschik (1973), they note that certain verbs seem more amenable to extraction than others. One possible generalization has to do with factive versus nonfactive verbs: "Factive verbs like *regret* presuppose the truth of their complements, whereas nonfactive verbs like *think* do not" (Culicover & Jackendoff, 2005, pp. 335). This helps explain why minimal pairs like 15 (a) and (b), which differ only in one lexical item, exhibit different extraction possibilities:

(15) a. Ginny thinks that Don bought a Porsche. → [no entailment about Don buying a

Porsche] → What does Ginny think that Don bought?

b. Ginny regrets that Don bought a Porsche. → Don bought a Porsche. → ??What does Ginny regret that Don bought?

*(Culicover & Jackendoff, 2005, pp. 335-6)*

If the extraction versus non-extraction parameter is a semantic feature of the lexical item itself, the difference between these minimal pairs is accounted for.

Another generalization, from Erteschik (1973), is "bridge verbs" and "non-bridge verbs." This captures the behavior of manner-of-speech verbs like *mumble, yell* or *groan,* which are more resistant to extraction than ordinary speech verbs like *tell* or *say.*

In short: in SS, semantic constraints work in tandem with syntactic operations to account for discontinuous dependencies. No movement means no intermediate traces thereof – at least not in the GB sense.

**3. From Theory to Processing**

GB and SS/HPSG agree that *wh*-relative clauses require some sort of special, or more "complex," mechanism for theta-role assignment. The divergences stem from where and how each theory chooses to insert this complexity into the system.

As we have seen, the theories' representations of long-distance dependencies are quite different on paper. But what predictions do these models make in terms of real-time language processing? Let us look at VP and NP extractions in a side-by-side comparison of GB and SS/HPSG. As our model sentences, we return to the captain and sailor, shown here again:

A. The captain who the sailor predicted that the weather would frighten turned back towards port.
B. The captain who the sailor's prediction about the weather had frightened turned back towards port.

For simplicity in the diagrams, we show the first NP only, because here is where the structural difference between the two conditions lies.

*GB: VP Extraction (a)*

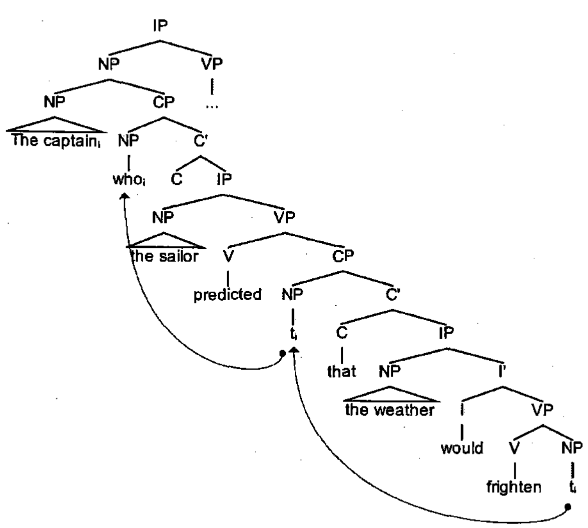


*HPSG: VP Extraction (a)*

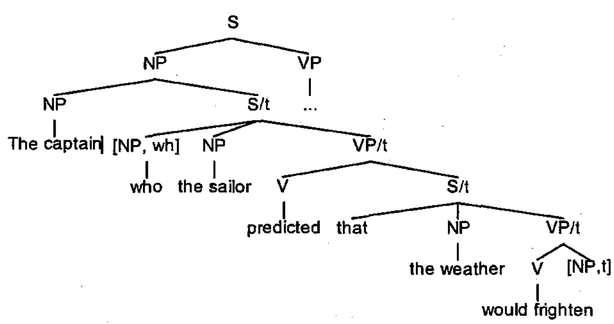

*GB: NP Extraction (b)*

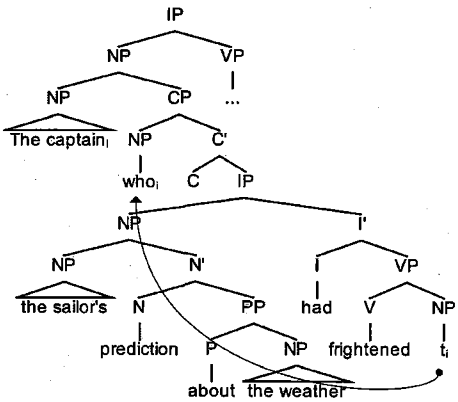


*HPSG: NP Extraction (b)*

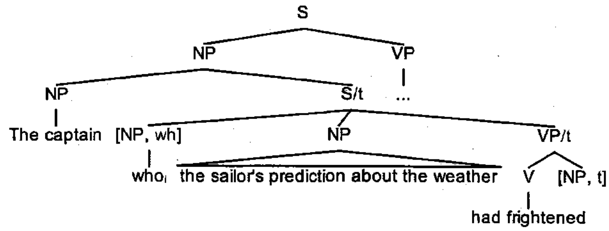


Neither GB nor HPSG is actually formulated in terms of processing, so we must extrapolate from these diagrams. The challenge for both is to account for what the processor does when it encounters a *wh*-word, knowing that this word must fill a gap somewhere downstream.

Of course, we first recognize the similarity between (a) and (b): both have an argument gap after the verb "frighten(ed)," separated from its filler, "the captain," by the same number of words. The Active Filler Strategy, originally proposed by Frazier and Clifton (1989), holds that the processor initiates a search for a gap as soon as it identifies a filler:

(16) When a filler has been identified, rank the option of assigning it to a gap above all other options.

The Active Filler Strategy enjoys much experimental support. Fodor (1995) sums it up nicely: "It seems that if the processor has a filler on its hands, it is inclined to jump at the first gap it finds." The consequence of this strategy is that if the first possible gap does not actually turn out to be the correct gap, the processor falls victim to a garden path and then must "back up" and restructure its analysis of the sentence.

So, upon hearing the word *who,* the Active Filler Strategy should invoked in both (a) and (b). With this in mind, we now explore the mechanism by which the processor embarks on its "gap hunt" in each of these theories in turn.

3.1 *GB: VP extraction (a)*

Right off the bat, we see that GB's movement is from right to left, but the processor works from left to right. We know that in real-time the processor does not benefit from knowledge of the sentence as a whole; it can only make its way along word by word (or morph by morph). And yet, successive cyclicity suggests that the processor should reactivate the filler at the boundary between the two embedded clauses in (a). We might say that the Active Filler Strategy might also predispose the processor to 'jump' at this site as a gap, since it is the object position for the verb "predict." But the semantics of "predict" indicate that this position is extremely unlikely to be the actual gap – after all, sailors don't predict captains. So why should there be reactivation of the filler here? If the processor is at all sensitive to the semantics of the sentence in real-time – as we know logically it must be, since we don't need to wait until the end of sentences to begin understanding them – it should know that "captain" cannot be the object of the verb "predict," and should wait patiently for a more plausible place to posit a gap.

But maybe the processor is tempted by a different garden-path: it could correctly anticipate that the *wh*-word comes from an embedded clause, but incorrectly posit that the word is the *subject* of this clause instead of its object. Examine the possibility in (17):

(17) The captain [who$_i$ the sailor predicted [t$_i$ would drink a lot of whiskey]] …

In this case, we might expect gap-filling after *predicted* just like in (2), since up through the first six words, the sentences look exactly the same. But this reactivation would not be an intermediate trace – it would be the argument gap. Disambiguation comes with the complementizer, *that.* This word closes the door on a subject-relative and essentially tells the processor "no gap here – keep going." The complementizer, then, is what triggers the reading we are aiming for in (a) – namely, an object-relative across two embedded clauses.

What makes this trace 'special' is that it surfaces even in "unambiguously intermediate" circumstances – like in (a) (an animate filler combined with an embedded verb strongly biased toward a CP complement or inanimate NP object). So if there really is reactivation of the filler at this intermediate site, we might say that the Active Filler Strategy is triggered for *purely syntactic* reasons. That is, there must be something encoded in the lexical entry a *wh*-word that tells the processor to reactivate it as every clausal boundary, regardless of the semantics of the sentence.

In the sense that this reactivation represents a feature of syntactic composition, it should be localizable – at this particular instant – to the left IFG. Note that in order to say anything about the presence of intermediate traces we must look for this activation *after* the onset of the complementizer, so as to control for any garden paths of type (17).

3.2 *HPSG: VP extraction (a)*

For HPSG, there is no successive cyclic movement, because there is no movement at all. Instead, the processor passes the slash-feature linearly down the tree from the moment it hears the *wh*-word until the moment it fills the gap. So the processor anticipates the gap at a level that stays constant throughout the intervening material; there is no 'special status' reactivation of the filler at clause boundaries like in GB. But the processor is systematic in the path it chooses for the slash-feature: it does not steer the slash into phrases like [the weather] because it knows – thanks to island constraints – that the gap will not be found within an NP (Stowe, 1986). Instead, the processor waits for the next node capable of containing the gap somewhere beneath it – in this case the VP – and passes the slash to this candidate.

Depending on how 'taxing' this constant anticipation might be, we might predict that the left IFG should be more or less equally active at all points between filler and gap. Of course, this interval includes our point of interest, the onset of the complementizer.

3.3 *GB: NP extraction (b)*

For the NP extraction condition, only one IP separates the filler and the gap, so there are no clause boundaries within the matrix NP. This means under GB that there is no successive cyclic movement, and therefore no reason for the processor to reactivate the filler at any special point within the intervening material. Upon hearing [the sailor's prediction…], the processor should recognize the NP island and wait patiently for the gap to occur at the end of that same clause.

Therefore, there should not be an increase in left IFG activity at the onset of *about,* which represents the linear point in (b) corresponding to the clause boundary in (a).

3.4 *HPSG: NP extraction (b)*

For the NP extraction, HPSG's prediction is similar but not identical to GB's. Again, because the processor is sensitive to island constraints, it will not attempt to transfer the slash to the NP starting with [the sailor's prediction...].

The real-time consequence of this is that the processor ends up putting the slash 'on hold' until it emerges on the other side of the NP. When the processor gets to the VP (starting with *had*), it recognizes that the gap might be within this phrase, so it resumes the slash here and maintains it until it finally fills the gap. In terms of the left IFG, then, HPSG does not predict increased activity (reflecting filler reactivation) at least until the processor gets through the NP – after the point of interest, the onset of *about,* has come and gone.

So might GB and HPSG, so different from a theoretical perspective, actually agree somewhat on how these sentences are processed? GB's 'special status' intermediate traces predict reactivation of the filler at the clause boundary in (a), with no corresponding activation in (b). As for HPSG, even though this trace is not 'special' per se, the linear path of the slash-feature causes it to manifest at the complementizer in (a) but not at the corresponding linear position in (b). In this sense, the theories do ultimately arrive at similar predictions in terms of processing.[1] Essentially, something about the structure of (a) mediates between the filler and the gap in a way that the structure of (b) does not. This mediation is a feature of syntactic composition, and therefore we should see it in the left IFG.

3.5 *A note about terminology*

This experiment did not seek to adjudicate between syntactic theories. Rather, we intended to conduct a well-controlled investigation into the processing of long-distance dependencies. We recognized from the outset the possibility that our results might be compatible with more than one theory – even those that are seemingly at odds – or that the results might not be entirely compatible with any existing theories.

In an effort to remain as theory-neutral as possible, at least for the time being, we use the terms "gap" or "trace" and not "movement" to refer to the displacement of arguments from their theta-role positions. When we must use movement-related terms to describe existing accounts for certain data in the literature, we recognize implicitly that the phenomenon would receive its respective alternate explanation in non-movement theories, according to what has been briefly outlined in section two. We use the term "argument trace" to refer to the lowest gap site, at which theta-role assignment takes place.

[1] If we wanted to adjudicate between GB and HPSG, we would have to use condition (b) and introduce another event – maybe between "had" and "pleased" – in order to compare left IFG activation there to left IFG activation at *about*. If the processor really does resume its slash-feature at "had," we should see left IFG activation at this new event but not at *about*. That is a very interesting question, but one for another day and another thesis.

## 4. Intermediate Traces: The Psycholinguistic Evidence

This section will outline the psycholinguistic evidence for intermediate structure. After a glance at language acquisition, we discuss Gibson and Warren's 2004 self-paced reading study, which is a foundation for the present experiment. We also briefly examine a priming study conducted by Shvartsman (2007) that failed to obtain significant results.

### 4.1 *Acquisition*

In a rather endearing line of evidence presented by Thornton (1990, 1995), it seems that young children acquiring English provide us with support for intermediate traces. Sentences like the following are attested in child-speech:

(18) a. Who do you think [$_{CP}$ who [$_{IP}$ the cat chased t ]]?
b. Which mouse do you think [$_{CP}$ who [$_{IP}$ the cat chased t ]]?

There seems to be an optional medial-*wh* that surfaces in child grammar, right in the SpecCP intermediate trace position. Thornton argues that these child speech patterns make ordinarily covert transformations into overt ones, giving us a rare window into what happens between deep structure and surface structure. At face value, this supports a traditional GB view of intermediate traces.

### 4.2 *Self-paced reading: Gibson & Warren (2004)*

The evidence from acquisition is interesting, but not particularly compelling unless it can be coupled with empirical investigations getting at this supposed "covert" mechanism of adult speech. To this end, Gibson and Warren (2004) used self-paced reading times during online sentence comprehension to test the psychological reality of intermediate traces.

Gibson and Warren establish that if intermediate traces exist, they shorten the linear distance between dependent items in the sentence, making integration of those items easier. So the presence of an intermediate trace should facilitate processing at the argument trace. This facilitation should be visible through faster reading times for VP extractions compared to NP extractions.

Gibson and Warren used four types of sentences: VP and NP extractions as experimental items and non-extracted VPs and NPs as controls. For example:

(19) a. *VP extraction*
The manager who the consultant claimed that the new proposal had pleased will hire five workers tomorrow.
b. *NP extraction*
The manager who the consultant's claim about the new proposal had pleased will hire five workers tomorrow.
c. *Nonextracted VP*
The consultant claimed that the new proposal had pleased the manager who will hire five workers tomorrow.
d. *Nonextracted NP*
The consultant's claim about the new proposal had pleased the manager who will hire five workers tomorrow.

As mentioned in the previous section, the challenge in an online experiment such as this is to create traces that are "unambiguously intermediate." That is, stimuli must semantically discourage a situation in which the filler NP is initially interpreted as the subject of the embedded clause. To this end, for their intermediate verbs, Gibson and Warren used bridge verbs that were strongly biased toward a CP complement rather than an NP object. If these verbs could take an NP object, it had to be inanimate, but the filler NP was always animate to further discourage the undesirable interpretation.

The non-extraction controls are necessary because there is a confounding difference between the VP and NP extractions that could influence processing times. In (a), the head of the verb *pleased*'s subject NP is *proposal*, while in (b) the head of *pleased*'s subject is *claim*. There is a shorter linear distance between *pleased* and *proposal* than between *pleased* and *claim*, potentially giving (a) an additional syntactic advantage. The non-extraction conditions control for this difference by requiring the same subject-verb integration at *pleased*, but without the long-distance *wh*-filler integration.

Because of the long-distance dependency, Gibson and Warren predicted that processing for conditions (a) and (b) should be slower than processing for (c) and (d) overall. The true test for intermediate structure, then, is the comparison of (a) and (b), the VP and NP extractions.

The experimental paradigm was a self-paced reading task. The sentences were segmented into seven regions, as shown in (20). Region 5 was designated the "critical region," where the argument trace and the *wh*-filler are integrated:

(20) a. *VP extraction*
[$_1$ The manager who] [$_2$ the consultant claimed] **t** [$_3$ that] [$_4$ the new proposal] [$_5$ had pleased] **t** [$_6$ will hire] [$_7$ five workers tomorrow].

b. *NP extraction*
[$_1$ The manager who] [$_2$ the consultant's claim] [$_3$ about] [$_4$ the new proposal] [$_5$ had pleased] **t** [$_6$ will hire] [$_7$ five workers tomorrow].

Gibson and Warren did, in fact, find that reading times were significantly faster at region 5 for the VP extractions compared to the NP extractions (by an average of 75 milliseconds after corrections for word length and individual variation in reading rates). (As predicted, reading times were faster overall for the non-extraction sentences compared to the extractions, but they were not significantly slower for (c) than for (d).) They interpreted these results as strong psycholinguistic evidence for the presence of intermediate structure.

*4.2.1 Gibson & Warren's results: a discussion*

Of course, not everyone accepts Gibson and Warren's interpretation of their results. Culicover and Jackendoff mention the psycholinguistic evidence for traces in *Simpler Syntax*, and respond: "It is important to recognize that evidence for the activation of a fronted constituent somewhat later in the processing of a sentence is not in itself evidence for a trace. It is simply evidence for the linking of the fronted constituent with something at the point of the activation (i.e. some kind of chain effect.)" (Culicover & Jackendoff, 2005, pp. 303). By way of an alternate explanation, they direct the reader toward Pickering and Barry (1991).

So we, in turn, take a brief look at Pickering and Barry's theory of long-distance dependencies. Using evidence from word order and constituency structure, they present an account of sentence processing without "empty categories," or gaps. Instead, they propose a direct association between the extracted element and its subcategoriser that mediates theta-role assignment. Arguments are thus linked not to their canonical positions within the sentence, but to their subcategorisers. This distinction becomes important in sentences like these:

(21) a. [In which box]$_a$ did you [put]$_a$ the very large and beautifully decorated wedding cake bought from the expensive bakery?
b. [Which box]$_a$ did you put the very large and beautifully decorated wedding cake bought from the expensive bakery [in]$_a$?

For most people, 21(a) is more acceptable than 21(b). But for a trace-based theory, the two are nearly identical in terms of structure: both would have a trace at the very end of sentence, given that canonical word order for *put* is agent-theme-goal (*I put [the cake] [in the box]*, not *I put [in the box] [the cake]*). The only difference between 21(a) and 21(b) is how much gets pied-piped along with *box*. Trace theory, then, has no way of capturing the difference in acceptability. Pickering and Barry's theory, on the other hand, does predict this difference. The shorter linear distance between *in which box* and its subcategoriser *put* in 21(a) makes the dependency easier to process than the one in 21(b), which deteriorates due to the distance between *which box* and its subcategoriser *in*.

It is a nice idea with some intuitive charm. However, contrary to Culicover and Jackendoff's suggestion, it cannot account for Gibson and Warren's data, because it does not predict a difference between VP extractions and NP extractions. The *wh*-filler *who* is separated from its subcategoriser *pleased* by the same number of words in both conditions, so they should look identical under Pickering and Barry's representation. See 22, *i* subscripts:

(22) a. The manager [who]$_i$ the consultant claimed that the new [proposal]$_j$ had [pleased]$_{i,j}$ will hire five workers tomorrow.
b. The manager [who]$_i$ the consultant's [claim]$_j$ about the new proposal had [pleased]$_{i,j}$ will hire five workers tomorrow.

Pickering and Barry's theory does predict a difference between (a) and (b) based on the distance between the verb and its subject NP head (*j* subscripts), with the advantage for (a). But Gibson and Warren did not obtain any difference between their (c) and (d) sentences, which required integrations of the same length as (a) and (b), respectively. Therefore, we can see that something else must be at work to explain Gibson and Warren's results.

In a more minor result, Gibson and Warren also found longer reading times at region 3 for VP extractions than for NP extractions. (This finding did not reach statistical significance at an items analysis, but is still worth mentioning.) They acknowledge two competing explanations – i.e. ones that do not rely on the presence of intermediate traces – but counter these explanations with additional data, ultimately supporting their initial hypothesis.

First, they acknowledge the garden-path possibility discussed in the previous section: subjects could have been positing an object gap for the matrix verb in the VP extraction (Clifton & Frazier 1989, De Vicenzi 1991, Gibson 1998, 2000). Reaching the complementizer informs them that their initial interpretation was wrong and forces a reanalysis, leading to longer reading times. However, as previously mentioned, because of the deliberate semantic constraints, Gibson and Warren declare this possibility "extremely unlikely."

Next, they acknowledge the second garden-path possibility: subjects could have been positing a *subject* gap in the *embedded* clause at region 3 in the VP extraction. Again, the complementizer forces reinterpretation, surfacing as longer reading times. But, perhaps a reactivation of the antecedent at region 3, even one that turns out to be structurally incorrect, still primes subjects to process the integration at region 5 more quickly than no reactivation at all. Such a proposal might predict faster reading times at the matrix verb *hoped* for object-relatives like 23(b):

(23) a. The reporter$_i$ who$_i$ $t_i$ sent the photographer to the editor hoped for a good story.
b. The reporter$_i$ who$_i$ the photographer sent $t_i$ to the editor hoped for a good story.

However, previous reading time experiments (King & Just 1991, Grodner et al. 2000) using relative clauses contradict this proposal. In fact, the opposite effect is observed: these experiments reported faster processing times at the matrix verb for subject-relatives like 23(a). Therefore, we might say that quicker processing time at the gap only seems to occur when the antecedent was reactivated as part of the same syntactic chain.

Despite some lingering questions about their analysis, Gibson and Warren (2004) is the best experimental evidence for intermediate traces we know of to date. Our experiment will build on their work by investigating the presence of intermediate structure via an arguably more direct window into online language processing: fMRI. Section five will discuss how we arrived at our hypothesis and how we went about localizing intermediate traces in the brain.

4.3 *Priming: Shvartsman (2007)*

Before we delve into the neurological literature, worthy of mention is Shvartsman (2007, unpublished), who tackled the question of intermediate structure via a different psychological methodology. Shvartsman reasoned that any reactivation of the antecedent in the intermediate position should be demonstrable through priming.

Shvartsman used questions with embedded clauses. To create unambiguously intermediate traces, he used unagentive extracted NPs with matrix verbs that were ditransitive and took a CP complement. For example:

(24) [Which gift]$_i$ did the man from Arizona tell Bill $t_i$ that Sarah would like $t_i$ for Christmas?

For the experimental paradigm, Shvartsman used a cross-modal lexical decision task, with auditory sentence presentation and visual lexical decisions. The lexical decisions were placed at three different points during the sentence: one experimental position and two control positions. The experimental position corresponded to the site of the intermediate trace. One control position was at the site of the argument trace, while another was at an earlier point in the sentence syntactically unrelated to the chain in question:

(25) Which gift did the man from [CTRL$_1$] Arizona tell Bill [TEST] that Sarah would like [CTRL$_2$] for Christmas?

Shvartsman used three target conditions: words semantically related to the antecedent (experimental), words unrelated to the antecedent, and non-words (both control). Priming for the related targets was expected at the test location as well as for control 2, but not for control 1.

However, Shvartsman's results did not show significant priming in any location. He offers a number of potential explanations for this. For one, he notes that many subjects reported issues with the volume and clarity of the recordings, which could have interfered with full comprehension of the sentence. Second, the comprehension questions focused on items that came later in the sentence, so subjects may not have paid close enough attention throughout the duration of the sentence. Third, and perhaps most interestingly, it seems that for some speakers, an overt complementizer renders these questions ungrammatical. He notes that perhaps the complementizer forms a CP-island, blocking extraction. So, for instance, while (24) is ungrammatical, a version without the complementizer is acceptable: "Which gift did the man from Arizona tell Bill Sarah would like for Christmas?"

The fact that Shvartsman did not obtain results even for control location 2 – where priming is well documented in the literature – suggests that one or more of the above explanations, or perhaps another factor altogether, confounded the procedure.

## 5. Why fMRI: Our Hypothesis

In using fMRI, one had better have an idea of what one is looking for – or, more precisely, where one is looking for it. We have been alluding to the left IFG, also known as Broca's area, Brodmann's areas 44 and 45 (roughly), or, more specifically, as the pars opercularis and pars triangularis of the inferior frontal gyrus. Why do we think that syntactic processing – specifically intermediate traces – happens here? That is what this section will set out to prove, using evidence from both lesion studies and fMRI.

### 5.1 *Lesion studies*

Since Paul Broca's 1861 characterization of a lesion to the left IFG that impaired what he described as *la faculté du langage articulé,* or speech production, the exact function of this area has been the focus of extensive inquiry. Broca noted that patients with damage to this part of the brain (now known as Broca's aphasics) demonstrated labored and "telegraphic" speech, often lacking function words. Though their speech was largely asyntactic, they seemed to understand the speech of others.

In 1874, Carl Wernicke characterized patients with damage to another perisylvian area of the brain, the left posterior superior temporal sulcus (pSTS). He observed deficiencies on the opposite side of the language system: comprehension. Though the patients' speech might be fluent and perfectly grammatical, it would be inappropriate to the situation at hand and largely devoid of meaning. The left pSTS became known as Wernicke's area and the language deficit, Wernicke's aphasia. Many psycholinguistic studies as well as clinical observations corroborated these classifications, and so they persisted largely unquestioned until the 1970s. Even now, one may sometimes hear the terms "expressive aphasia" and "receptive aphasia," referring to Broca's and Wernicke's aphasia, respectively.

An experiment conducted by Caramazza and Zurif in 1976 spurred a movement away from the comprehension versus production classifications, and toward an understanding of Broca's area better informed by linguistic principles, specifically syntax. Caramazza and Zurif found that Broca's aphasics could not comprehend sentences with object-relatives. When it came to these constructions, they performed at chance on a sentence-picture matching task. But, patients performed above chance on subject-relative sentences. By way of explanation, Caramazza and Zurif proposed that Broca's aphasics were relying on some sort of comprehension heuristic that worked for subject-relatives but failed for object-relatives. It was this heuristic that allowed Broca's patients to "fool" linguists and doctors alike into thinking their language comprehension remained unaffected. Caramazza and Zurif's was one of the first experiments to suggest that language is not divided along a set of modalities (speaking, listening) but a set of systems, at least one of which is algorithmic in nature (namely syntax), and that these algorithms have neural representation.

Since then, a plethora of experiments have documented Broca's aphasics' performance on various sentence types across a number of different languages and task paradigms. In addition to object-relatives, Broca's aphasics also perform at chance on many passive constructions. But they perform above chance on actives, subject-relatives, adjectival passives and truncated passives (see Beretta et al., 1999). With the abundance of data, there have been many attempts to formulate a descriptive generalization accounting for their behavior.

In one notable attempt, Grodzinsky proposed the Trace Deletion Hypothesis (TDH), which generalizes Broca's aphasia in terms of a failure to correctly derive argument structure in sentences with traces (Grodzinsky, 1986, 1995). The TDH establishes that in SVO languages like English, verbs assign <agent> to the argument to their left and <patient> to the argument to their right. For healthy subjects, traces mediate the link between the theta-assignment position and the *wh*-filler in relative clauses. For Broca's aphasics, all traces are deleted, severing this link and disrupting theta-role assignment. The end result (for English object-relatives) is two possible agents, prompting patients to guess and producing the chance comprehension patterns observed.

The TDH has been modified several times since its original version (one significant modification was Mauner et al.'s 1993 Double-Dependency Hypothesis). However, several problems remain. There are structures for which the TDH predicts chance performance but Broca's aphasics perform above chance, and vice versa (i.e., Beretta et al. 1999).

Note that the TDH is a descriptive generalization for comprehension in Broca's aphasics; it is not formulated to account for production. Even if it could be reformulated "in reverse," it would not account for all the data, since observation tells us that Broca's aphasics have production trouble with more than just trace-based structures.

Rather, a line of altogether different descriptive generalizations has been proposed to account for impairments in production. Generally, Broca's aphasics' speech is deficient in functional morphosyntactic elements like determiners, tense, agreement, complementizers and relative pronouns. The Tree-Pruning Hypothesis (TPH) proposed by Friedmann and Grodzinsky (1997) generalizes these deficiencies to an inability to represent certain levels of the syntactic tree. Based on the observation that tense is impaired more often than agreement, and the theoretical assumption that tense is higher on the tree, Friedmann claims that impairment can be localized to nodes above agreement. However, the TPH has been the subject of much criticism, as it, too, fails to account for much of the data (i.e., Burchert et al. 2005).

The TDH and TPH have very little in common other than a basic appeal to syntax to account for Broca's aphasics' behavior. One deals with traces and theta-role assignment, the other with nodes of agreement and tense. Do we have a reason to suspect that Broca's area plays such different roles in speaking and understanding language? Perhaps a descriptive generalization that accounts for both production and comprehension is more empirically desirable.

This decoupling of modalities is one issue, but not necessarily the main one. From our vantage point, the biggest problem with these hypotheses is that they are *offline* accounts of Broca's patients' deficits (and respective compensatory strategies). They are concerned only with predicting the end result: the observed behavior. This means they do not discriminate between knowledge and implementation of knowledge. For example, the TDH cannot distinguish whether Broca's aphasics have lost their representation of traces altogether, or just their ability to process them. This distinction – between language representation and language processing – is crucial.

There are some theories of Broca's aphasia that do appeal to processing to explain the observed deficits. In contrast to the TDH and TPH, in processing-based theories, Broca's patients have not lost something so specific as their knowledge of traces or specific nodes on the tree. In fact, they have not actually lost syntactic representation at all. Instead, the damage to Broca's area impairs resources that are crucial for constructing syntactic representation *in real-time*. This is true for both encoding (production) and decoding (comprehension).

These processing theories can be modeled in a number of ways. Avrutin (2005, "weak syntax") describes an economy hierarchy for language in which syntactic mechanisms are normally less "costly" than semantic mechanisms and thus are relied on more heavily for transferring information. In Broca's aphasia, this hierarchy shifts so that syntax becomes just as costly – if not more so – than semantics, so the two compete. Processing can also be modeled in terms of a chronological hierarchy (Piñango, 2000, "slow syntax"), which holds that syntactic processing normally occurs before – and provides crucial information for – semantic interpretation. The impairment in Broca's aphasia causes a slow-down in syntax, so that syntactic and semantic processing end up occurring at the same time. Regardless of how it is packaged, the main idea is this: the two systems (syntax and semantics) end up working in parallel rather than in series. If their outputs differ, subjects are forced to guess, deriving the characteristic chance performance.

Different outputs seem to arise from structures with non-canonical word order (i.e. object-relatives but not subject-relatives). So, in the absence of a universally agreed-upon model, we can use the evidence from aphasia to (somewhat safely) establish this much: Broca's patients experience a breakdown in argument structure for structures that deviate from canonicity.

But where does this breakdown occur – in the syntax (automatic lexical reactivation at the gap site) or at the interface with semantics (derivation of argument structure)? Can the two be dissociated? In other words, is it possible to fill a gap and then fail to assign theta-roles, or fail to fill a gap but still derive the correct theta-role assignment? To investigate this question, we turn to evidence from online sentence processing, chiefly from Swinney and Zurif (1995).

According to Swinney and Zurif, the short answer is yes, the two processes can be dissociated. Broca's patients do not show priming at the gap in cross-modal lexical decision tasks, regardless of whether they ultimately interpret the sentence correctly. This plainly suggests that regardless of how they compensate, Broca's aphasics do not engage in correct gap-filling, and therefore that the gap-filling mechanism must rely on resources housed in Broca's area.

Wernicke's patients, on the other hand, *do* show priming at gaps even in sentences that they ultimately fail to understand. This suggests that Wernicke's area must *not* be crucial for the syntactic operation of gap-filling, and that the Wernicke's deficit must occur somewhere downstream, most likely in mapping to argument structure.

This dissociation is further supported by a series of studies conducted by Shapiro et al., in which the experimenters teased apart argument structure as a semantic process distinct from its syntactic counterpart (subcategorization) by selecting verbs for which the two differ in complexity. Complexity is defined as the number of possible configurations. In a 1989 experiment with neurologically normal subjects, Shapiro et al. showed that semantic complexity caused processing interference, demonstrable through slower reaction times in a cross-modal lexical decision task. Repeating the same study with aphasic patients, Shapiro et al. (1993) showed that Broca's aphasics patterned with normals, while Wernicke's aphasics did not experience this interference. So, it seemed that Wernicke's aphasics had lost sensitivity to semantic complexity, while Broca's aphasics had retained this sensitivity.

To answer our previous question, in light of all of this evidence, it appears that for Broca's patients the breakdown occurs in syntactic processing. Indeed, syntax is what unites the generalizations behind the TDH, TPH, weak syntax, SSH and so on. This evidence has theoretical implications, suggesting that gap-filling and argument structure, are, in fact, not one and the same.

From aphasia, we take that the left IFG is crucial for certain kinds of syntactic work – that is, the implementation of algorithms that operate at the interface between phonology and meaning. Exactly which of these theoretical operations exist from a processing standpoint, and which can be localized to this particular brain region, is an empirical question.

This is what makes Broca's area a worthy target for an investigation of long-distance dependencies. Specifically, intermediate traces, insofar as they are a feature of syntactic composition, will help further clarify the role of the left IFG.

5.2 *Neuroimaging: fMRI*

Lesion studies are extremely valuable, but because they exploit pathology, by definition they cannot provide us with a complete picture of normal brain-language relations. (That is, lesion studies can exclude the involvement of Broca's area and observe the resulting aberrant behavior, but only imaging studies can positively associate normal behavior with its neural substrate – not only in one focal area, but across the whole brain.) FMRI in particular has been used to try to pinpoint language processing in the brain.

FMRI relies on a measure of hemodynamic response through BOLD, or blood-oxygen level dependent. Oxygenated blood causes differences in magnetic signal that can be detected using an MR scanner.

FMRI studies of sentences with long-distance dependencies have converged in many ways with lesion studies – that is, these sentences seem to elicit more left IFG activation than those without. But the imaging data has introduced new subtleties into the mix, and many questions remain unanswered. This section will highlight some of the past decade's major fMRI investigations of trace-based and other relevant structures, complete with their overlaps and divergences, and its subsection will present a working hypothesis to serve as the basis of the present study.

Cooke et al. (2001) conducted one of the first fMRI studies of sentences containing traces. They used four conditions in a 2 x 2 paradigm, varying subject- and object-relatives as well as long and short distances between antecedent and gap (seven and three words, respectively). Their findings can be summarized as follows: 1) increased activation in left IFG for long object-relatives only, 2) right pSTS activation for long distances (both subject- and object-relative), and 3) left pSTS activation in all conditions.

Cooke et al. interpreted these results in the following ways: 1) left IFG is recruited for sentences with simultaneously non-canonical word order and high short-term memory demands (but not for each individually), 2) right pSTS is implicated for relatively long dependencies and

3) left pSTS is involved in some sort of common semantic processing, independent of syntactic structure.

Most interesting for our purposes at the moment is (1). Note that the increased activation for left IFG in long object-relative sentences compared to the other three sentence types does not mean the area was only active during long object-relatives. Rather, it most likely means that the area was recruited somewhat equally for all three other sentence types, but activation was particularly strong for long object-relatives. One issue was that Cooke et al. did not include a trace-less control condition, so by its very nature the study was not able to localize traces themselves.

A 2003 Ben-Shachar et al. experiment was better suited to do so. They compared Hebrew sentences containing object-relatives (traces) to those with CP complements (no traces), i.e. "I helped the girl that Rina met *t* in the garden" versus "I told Rina that the girl slept in the garden." They reported increased activation in the left IFG and bilateral pSTS for sentences with traces compared to those without.

Next, Ben-Shachar et al. (2004) explored two additional syntactic contrasts with traces: topicalization and embedded subject and object *wh*-questions. They used two additional conditions without traces – dative shift and yes/no embedded questions – as controls. The sentences with traces activated a set of brain regions consistent with the previous study, including left IFG and bilateral pSTS. Crucially, they reported no difference in activation for subject- and object-relatives.

So the Cooke et al. and Ben-Shachar et al. results are largely convergent on the left IFG as an important area for sentences with traces. At first glance, one discrepancy surfaces between the two: antecedent-gap distances in Ben-Shachar et al. (2003) were of similar length to the "short" condition in Cooke et al. (about three words), which would predict similar activation patterns. But Cooke et al. obtained left IFG activation only for their "long" dependency condition (seven words). Remember, though, that Cooke et al. did not use a traceless control. The two findings can perhaps be reconciled under an incremental view: the left IFG processes dependencies of any length, but it is recruited more heavily for increasingly long distances between antecedent and gap.

So, sentences with traces seem to activate the left IFG. But let us take a step backwards. This does not necessarily entail that the traces themselves are the source of the activation. Bornkessel et al. (2005) point out that in the literature thus far, "syntactic factors," or traces, covary with complexity of deriving argument structure. As they define it, this complexity increases with non-canonical word order.

For their experiment, Bornkessel et al. took advantage of some unique grammatical features of German to manipulate word order without invoking traces. They employed a three-way contrast: word order (subject-object or object-subject), verb class (active versus object-experiencer) and morphology (unambiguous versus ambiguous). Interestingly, they reported that left pSTS and left IFG were activated by different variables: 1) the left IFG was sensitive to non-canonical word order (object-subject for active verbs) regardless of morphology; it was also independently sensitive to ambiguous morphology regardless of word order and verb class, and 2) left pSTS was sensitive to the interaction between all three conditions: word order, verb class and ambiguous morphology. In light of these results, the authors claim that the view advanced by Cooke, Ben-Shachar and other previous studies – that the left IFG is specific to syntactic transformations, or traces – is too narrow.

Continuing chronologically, we come to Fiebach et al. (2005), who used stimuli similar to Cooke et al.'s: short and long subject- and object-relatives. German follows an SOV word order in embedded clauses, so subject-relatives preserve the subject-object order while object-relatives reverse it. Therefore, given Bornkessel, we might predict increased left IFG activation for object-relatives. But Fiebach et al. found no difference in activation anywhere in the brain between subject- and object-relatives; only the length contrast elicited increased left IFG activation. Fiebach et al. interpreted their results to mean that the left IFG is involved in working memory as opposed to traces in particular. (But note that Fiebach et al. suffer from the same problem as Cooke et al.: a lack of traceless control sentences.)

Next, Santi & Grodzinsky (2007) used trace sentences along with another syntactic feature involving co-dependent, non-adjacent elements: reflexive binding. They argued that reflexive binding, like traces, requires working memory to retrieve an antecedent at a particular site. However, binding structures did not selectively recruit the left IFG, while object-relatives did.[2] (Santi and Grodzinsky did not test subject-relatives.) The left IFG was recruited incrementally as the linear distance between antecedent and gap increased. In contrast to Bornkessel's analysis, Santi & Grodzinsky used this data to claim that the left IFG is specific to the processing of trace structures.

One final study worth mentioning is Roder et al. (2002), who varied both canonical/non-canonical word order and words/pseudowords in German. They found that the left IFG was most active for real words in non-canonical order.

At this point, it is quite clear that imaging has arrived at the same tension present in lesion studies: is the left IFG necessary for traces in particular, or is it involved in word order (for argument structure) more generally? The data presented thus far is summarized in (26).

Here, it is valuable to recall Swinney and Zurif's findings for gap-filling in aphasia as well as the Shapiro et al. experiments dissociating argument structure from subcategorization. Both suggest that traces (syntactic form) can be decoupled from their role mediating argument structure. Both suggest that Broca's area does the traces, or syntactic work.[3]

Given that we can dissociate gap-filling from argument structure, and we know gap-filling to a be a syntactic process, perhaps we can reconcile all the data involving the left IFG in terms of *syntactic integration.*

---

2 Lesion studies converge with the Santi & Grodzinsky results: Broca's patients are able to understand binding dependencies (Blumstein et al., 1983).

3 It is important to remember that just because Broca's patients arrive at the correct interpretation for subject-relatives does not mean that they are processing them in the same way as normals. (In fact, we know they are not, thanks to the data from priming.) Therefore, the left IFG might still process traces in subject-relatives and other structures that Broca's aphasics perform above chance on – it is simply that Broca's aphasics' compensatory strategy allows them to comprehend these sentences via alternate means. Rather, it is the Wernicke's patients who crucially inform our analysis here: they show automatic reactivation at the gap site (a syntactic process), but ultimately fail to derive the correct argument structure. Because their Broca's area is intact, this is consistent with the idea that the left IFG is involved in syntactic processes.

(26) *The Left IFG: A Summary of the Literature*

| | **Left IFG** | **Control items (no left IFG)** | **Left pSTS** | **Right pSTS** |
|---|---|---|---|---|
| Cooke (2001) | Long object-relatives | (none) | All relatives | Long object- and subject-relatives |
| Ben-Shachar (2003) | Object-relatives | CP complements | Object-relatives | Object-relatives |
| Ben-Shachar (2004) | Topicalization, embedded subject- and object-*wh*-questions | Dative shift, embedded yes/no questions | Topicalization, embedded *wh*-questions | Topicalization, embedded *wh*-questions |
| Bornkessel (2005) | Non-canonical word order (OS for active verbs), ambiguous morphology | Canonical word order (SO for active verbs) | [Non-canonical word order + morphology] – interaction | (none) |
| Fiebach (2005) | Embedded subject and object *wh*-questions (long > short) | (none) | Embedded object *wh*-questions (long > short) | (none) |
| Santi & Grodzinsky (2007) | Object relatives (long > short) | Reflexive binding | (none) | (none) |
| Roder et al. (2002) | Non-canonical word order (real words > pseudo-words) | Canonical word order, pseudo-words | Real words ( > pseudo-words) | (none) |

5.3 *The Syntactic Integration Hypothesis*

So we want to say that the left IFG does "syntactic integration." What does that mean?

The syntactic integration hypothesis (SIH) states that the left IFG plays home to the syntactic processor, whose job it is to build enriched syntactic structure in the anticipation of theta-role assignment. Simple SVO sentences do not rely on enriched structure. But sentences with long-distance dependencies, or with their arguments displaced in some other way, do rely on this structure, and therefore activate the left IFG.

What do these sentences look like in real-time from the processor's perspective? Here, we take several examples, culled from both the fMRI and aphasia literature, to explore the predictions of the SIH. For (1)-(5), all of the (a) sentences produced more left IFG activation than the (b) sentences:

<u>Data from fMRI</u>

1. a. [The messy boy]$_i$ who Janet the very popular hairdresser grabbed $t_i$ was extremely hairy.
   *long object-relative*
   b. [The strange man]$_i$ in black who $t_i$ adored Sue was rather sinister in appearance.
   *short subject-relative, Cooke et al. (2001)*

2. a. I helped the girl$_i$ that Mary saw $t_i$ in the park.
   *object-relative*
   b. I told Mary that the girl ran in the park.
   *CP complement, Ben-Shachar et al. (2004) – Hebrew*

3. a-i. The waiter asked which tourist$_i$ $t_i$ ordered an alcoholic drink in the morning.
*embedded subject question*
a-ii. The waiter asked which drink the fat tourist ordered $t_i$ in the morning.
*embedded object question*
b. The waiter asked if the tourist ordered an alcoholic drink in the morning.
*embedded yes/no question, Ben-Shachar et al. (2004) – Hebrew*

4. a. Yesterday, someone said that [the boy]$_{DAT}$ [the teachers]$_{NOM}$ help.
*object-subject, active verb*
b. Yesterday, someone said that [the boy]$_{NOM}$ [the teachers]$_{DAT}$ helps.
*subject-object, active verb, Bornkessel et al. (2005) – German*

5. a-i. Thomas asks himself, who$_{NOM}$ on Tuesday afternoon after the accident the$_{ACC}$ doctor has called.
*long subject-relative*
a-ii. Thomas asks himself, who$_{ACC}$ on Tuesday afternoon after the accident the$_{NOM}$ doctor has called.
*long object-relative*
b-i. Thomas asks himself, who$_{NOM}$ the$_{ACC}$ doctor on Tuesday afternoon after the accident has called.
*short subject-relative*
b-ii. Thomas asks himself, who$_{ACC}$ the$_{NOM}$ doctor on Tuesday afternoon after the accident has called.
*short object-relative, Fiebach et al. (2005) – German*

For (6) and (7), Broca's patients' observed behavior is given:

Data from Broca's Aphasia

6. The tiger that chased the lion is big.
- *Did the tiger chase the lion/Did the lion chase the tiger?* ABOVE chance
- *Is the tiger big/Is the lion big?* AT chance
*(Hickok et al. 1993)*

7. Right-branching:
i. Point to [the man]$_i$ who$_i$ $t_i$ is calling the boy. ABOVE chance
ii. Point to [the boy]$_i$ who$_i$ the man is calling $t_i$. AT chance
*(Beretta et al. 1996) – Spanish*

Let us first establish this much: the processor expects, due to some feature of conceptual structure beyond the scope of this paper, that the first thing it hears will be an agent of some sort – or an "actor" – while the second thing should be a patient – or an "undergoer." For lack of a better term to capture this conceptual regularity, we continue to refer to it here as canonicity.

We also establish that when the processor encounters an argument that does not immediately get a theta-role, it must hold this argument in mind until it does get a theta-role. Linking an argument with its theta-role is what we mean by "integration." We assume that something inherent to *wh*-words prompts the processor to begin building its "enriched structure" in anticipation of theta-role assignment.

So for example, in a sentence like 1(a), up until the *wh*-word, the processor is expecting [the messy boy] to be the agent by default. But upon hearing *who*, the processor must accept the possibility that, depending on how the relative clause pans out, the boy could be acting as a local subject or object. No matter what, though, [the messy boy] will still end up being the subject of the matrix sentence.

If [the messy boy] ends up being the subject of the embedded clause, the dependency is resolved fairly quickly. But even in this case, if the clause is center-embedded, the processor still must wait to integrate the matrix subject with the matrix predicate (in this case, [was extremely hairy]). In this sense, center-embedded subject-relatives are just as complex as center-embedded object-relatives with respect to integration. But, right-branching relatives should be simpler overall than center-embedded relatives because they do not require this final long-distance integration.

This prediction is borne out by evidence from positron emission tomography, suggesting that the left IFG is more active for center-embedded relatives than for right-branching relatives (Stromswold et al. 1996). It is also borne out in aphasia: for center-embedded subject-relatives, Broca's patients can handle the local integration within the first NP and relative clause, because canonical word order is preserved. (For object-relatives, they are at chance even locally because, hearing two arguments before a verb, their compensatory strategy cannot apply and they are forced to guess.) But for both subject- and object-relatives, the integration of matrix subject with matrix predicate proves too much, and they perform at chance on questions probing this integration. Hence the observed behavior in (6) (Hickok et al. 1993).

For right-branching structures, on the other hand, subject-relatives *should* be less complex with respect to integration than object-relatives. This prediction is supported in that Broca's aphasics comprehend right-branching subject-relatives but not object-relatives, as seen in (7) (Beretta et al. 1996).

So, with the idea that *wh*-words trigger enriched structure in both subject- and object-relative contexts, we have accounted for the results in (2) and (3) (Ben-Shachar et al. 2003 and 2004). Remember that if one only compares long-distance dependencies to *longer*-distance dependencies, the longer ones will surface with greater left IFG activation (Cooke, Fiebach, Santi & Grodzinksy). This does not mean that the shorter dependencies did not recruit the left IFG at all. Length is not the core variable triggering left IFG activation – it is the presence of the dependency itself. With this in mind, we can (somewhat safely) say that the data in (1) has also been accounted for.

Languages like German, with overt case, are a different story. If case morphology makes theta-roles clear right away, the processor only needs to build enriched structure when the sentence does not conform to the expected subject-object word order. This accounts for the Bornkessel et al. (4) and Roder et al. results. (Note that Bornkessel et al. found increased left IFG activation for all sentences with ambiguous morphology, regardless of word order or verb class. This is also predicted by the SIH, because theta-roles were *not* made clear right away.)

But Fiebach et al. also varied subject-object with object-subject word order and did not find a difference in left IFG activation. How can we reconcile these results with Bornkessel et al.? Here it is important to remember that activation is defined in *relative* terms. Perhaps the source of the discrepancy was Fiebach et al.'s additional variable: distance between arguments (kept constant in Bornkessel et al.). Right-branching relatives like in (5) are costly only until argument hierarchization is complete. In 5(b), because the two arguments appear one right after another, the hierarchization is resolved more quickly than in a sentence like 5(a), where the processor must wait through intervening material between the two arguments. Hence the increased activation for 5(a). It would be interesting to repeat the Bornkessel et al. subject-object active condition with filler material between the arguments and the verb, to see if this would trigger left IFG activation and replicate Fiebach et al.

To sum up:

The SIH accounts for the Cooke et al. data in (1) because the dependency must be maintained for longer in the (a) sentence. (Remember that Cooke et al. did not include sentences without dependencies, so all of their left IFG activations are relative to shorter dependency conditions.)

It accounts for the Ben-Shachar et al. (2003, 2004) data in (2) and (3) because the (a) sentences contain dependencies to be integrated whereas the (b) sentences do not.

It accounts for the Bornkessel et al. (2005) data in (4) because in (a) the processor does not expect to hear the object first (marked in real-time by the dative case morphology).

It accounts for the Fiebach et al. (2005) data in (5) because the processor must wait longer in (a) before it gets the complete picture of argument hierarchization within the sentence.

So according to the SIH, the left IFG is not about traces or no traces, subject-relatives or object-relatives. It is, more generally, about building structure in preparation for argument structure integration. The longer the processor has to wait to assign theta-roles, and the harder it has to work to determine those theta-roles, the more active the left IFG will be.

### *5.3.1 The bigger picture: Back to intermediate structure*

We acknowledge that a number of factors can complicate the interpretation of fMRI data, including less than ideal temporal resolution and low signal-to-noise ratio. We also acknowledge that the case for Broca's area is not closed (see Santi & Grodzinsky 2008 for a review), and indeed that the region may very well be multi-functional. Still, the relative convergence of evidence from fMRI and aphasia suggests that the building of syntactic structure is a distinct neurological process, localizable to the left IFG.

With the SIH in mind, despite some minor lingering discrepancies (which may be due to methodological and sensitivity differences), we can begin to piece together a story about sentence processing from this data. The left IFG syntactic processor builds enriched structure to facilitate ultimate theta-role assignment. But deriving argument structure itself does not necessarily take place in the left IFG. It seems that the information decoded in the left IFG is then implemented in the left pSTS, where the semantic process of theta-role assignment can proceed.

Here is where an investigation of intermediate traces comes into play. As we saw from the processor's perspective in section three, theories predict a difference between VP and NP extractions. Turning briefly back to psycholinguistic evidence, we recall that Gibson and Warren described intermediate traces as a feature of online sentence processing that facilitates eventual integration. So while according to the SIH we might expect the left IFG to be more active in sentences with intermediate traces compared to those without, this extra activity corresponds to extra structure-building, which ultimately makes argument structure easier to derive.

Hence, our hypothesis: increased left IFG activation for VP extractions over NP extractions, supporting the idea that this area is concerned specifically with building syntactic composition. This is what we set out to test.

### 5.4 *A note about data analysis*

Note that for an investigation of intermediate traces, it is necessary to examine intra-sentential events when analyzing the imaging data. In contrast to previous literature, it is not enough to simply subtract whole sentence conditions: both VP and NP extractions contain argument traces, so left IFG activation might be similar when observed across the time course of the whole sentence. Instead, we must obtain an activation map for the specific time at which the intermediate trace occurs, to see if there is a difference in activation – at the left IFG in particular – for sentences with intermediate structure compared to those without. This is a novel analysis technique that has not been used in the literature presented here.

**6. Experiment**

This section will describe our experimental procedure.

6.1 *Subjects*

Fifteen undergraduate native speakers of English participated in the study (eight women, seven men). All except one were right-handed. The left-handed subject did not show significant differences in language lateralization from the rest, so his data was included in the analysis. All had normal or corrected-to-normal vision, with no history of neurological disease or head injury. We obtained written informed consent from all subjects.

6.2 *Materials*

Experimental items were written in groups of four sentence conditions (quadruples).

The first type of sentence, condition (a), contained an intermediate trace. It essentially replicated the structure of Gibson and Warren's VP extraction condition. Following Gibson and Warren's design, sentences were created using bridge verbs as matrix verbs that tended toward CP complements rather than NP objects. To the extent that the verbs could take NP objects, there was a strong requirement for an inanimate object, and matrix subjects were always animate objects.

We took one step further to preempt subjects from interpreting the matrix subject as the subject of the embedded clause as in (17): we added adverbials to our sentences after the matrix verb and before the complementizer or preposition (*that* or *about*). Adverbials were either temporal or locative in nature ("this morning," "during the meeting," "in New York," etc). These were used to fill the position immediately following the matrix verb, discouraging any unwanted gap-filling effects until the onset of the complementizer, which would then steer the processor toward the correct interpretation and trigger the intermediate trace.

(27) is an example of an (a) sentence:

(27) The captain, who the sailor predicted yesterday that the weather would frighten, turned back towards port.

Our second condition, (b), provided the main contrast to (a). Following Gibson and Warren's NP extraction condition, condition (b) was created by nominalizing the embedded CP from condition (a). This lead to a structure that, though identical to (a) in length and linear position of arguments, differed importantly from (a) in that there is no intermediate trace:

(28) The captain, who the sailor's prediction yesterday about the weather had frightened, turned back towards port.

The third condition, (c), was an ungrammatical sentence created by illegally filling the gap at the site of the argument trace:

(29) *The captain, who the sailor's prediction yesterday about the weather had frightened *the crew*, turned back towards port.

The fourth and final condition, (d), served as a control for features of the other sentence types. (A) and (b) differ in that (a) uses the complementizer "that" while (b) uses the preposition "about," an unavoidable lexical consequence of the contrasting structures. By inserting a verb directly after the subject, turning the embedded CP into the direct complement of the verb and the entire utterance into a compound sentence, condition (d) eliminates all traces from its structure. It uses the complementizer "that," and therefore controls for this feature of the (a) sentences. It also uses the same embedded object as (c), but in a way that is grammatical:

(30) The captain believed the sailor's prediction yesterday that the weather would frighten the crew and turned back towards port.

Putting it all together, (31) is an example of a full quadruple:

(31) a. The captain, who the sailor predicted yesterday that the weather would frighten, turned back towards port.
b. The captain, who the sailor's prediction yesterday about the weather had frightened, turned back towards port.
c. The captain, who the sailor's prediction yesterday about the weather had frightened the crew, turned back towards port.
d. The captain believed the sailor's prediction yesterday that the weather would frighten the crew and turned back towards port.

Note that the same adverbial as for condition (a) was also used in (b), (c) and (d), to keep all sentences length-matched and as close to minimal pairs as possible. Commas setting off the embedded CP were present to suggest prosodic contours and contribute to ease of comprehension, in all conditions except (d).

Comprehension questions followed all sentences of type (a), (b) and (d). No questions followed sentences of type (c) – due to the ungrammaticality, it is difficult to ask questions that have a clear yes/no answer. Questions probed different combinations of the matrix subject, embedded subject, matrix verb and embedded verb. A sample set of possible questions for sentence 31(a) above:

(32) i. Did the sailor predict that the weather would frighten the captain? (Y)
ii. Did the captain predict that the weather would frighten the sailor? (N)
iii. Did the captain turn back towards port? (Y)
iv. Did the sailor turn back towards port? (N)

Types of comprehension questions were varied, so as to discourage subjects from developing a strategy that would allow them to answer the questions without constructing a full mental representation of the events in the sentence.

*6.2.1 Plausibility norming survey*

The semantic relationships in the sentences were intended to be equally plausible across conditions (except, of course, for the violation in (c)). To ensure that this was indeed the case, all sentences were normed using a sentence judgments questionnaire prior to the experiment. In this way, we were able to further curtail any confounding variables affecting differences between conditions, especially between condition (a) and condition (b). Note that, though close to minimal pairs, the structure of sentences 31(a) and 31(b) create meanings distinct from one another. See (32):

(32) In (a): The sailor predicted that the weather would frighten the captain.
In (b): The sailor's prediction about the weather had frightened the captain.

In other words, in (a), it is the weather that frightens the captain, while in (b), it is the sailor's prediction that frightens the captain. If there is a plausibility difference between condition (a) and condition (b) within the same quadruple, one sentence may be more taxing to the system for a reason entirely separate from the presence or absence of an intermediate structure, confounding the results. Specifically, we were concerned about (a) being more semantically plausible on top of its syntactic "advantage" over (b).

Gibson and Warren normed their 20 sentences using a 7-point scale. We omitted sentences created by Gibson and Warren for which the VP extraction condition received higher average plausibility ratings than the NP extraction condition. We retained sentences for which ratings were the same or higher for the NP extraction condition. This gave us 16 pairs of sentences, for which we created items for condition (c) and (d) to give us 16 quadruples.

We then created 44 original quadruples for a total of 60 quadruples, 240 individual sentences. We assembled two questionnaires of 120 sentences each (two sentences from each quadruple presented in a random order) paired with a 5-point acceptability scale (1 = least acceptable, 5 = perfectly acceptable). We administered each questionnaire to 13 subjects. All were undergraduates and native speakers of English, the same population that participated in the fMRI study.

No significant plausibility difference was found between conditions (a) and (b), which received mean plausibility ratings of 2.66 and 2.67, respectively.

It was expected that plausibility ratings for condition (c) would be significantly lower than ratings for the other conditions, and this was indeed the result. Condition (c) received a mean plausibility rating of 1.73 ($p < .001$).

Plausibility ratings for condition (d) were significantly higher than for conditions (a) or (b), with a mean of 3.80 ($p < .001$). But because (d) was a control condition that would eventually be subtracted from both (a) and (b) (see the next subsection), its higher rating does not pose a problem for our analysis.

6.3 *Contrasts*

The contrasts between sentence conditions allow us to examine a variety of processes with respect to traces and integration of long-distance dependencies.

As mentioned in the previous section, a proper analysis of intermediate traces must look "inside" the sentence to dissociate the effect of single-clause argument traces from double-embedded intermediate traces. To this end, we defined two events within each sentence. Event 1 corresponded to the site of the intermediate trace in (a) and its linear equivalent in (b), (c) and (d). Event 2 corresponded to the site of the argument trace in (a) and (b), the illegal gap filler in (c) and the grammatical object position in (d).

Event 1 was defined as the onset of "that" or "about." We were aiming to maximize the window of potential activation while ensuring that the system would have recognized the need for a trace. Event 2 was defined as the offset of the lowest embedded verb so that we could capture the full contrast expected to surface immediately following it. To illustrate:

(33) a. The captain, who the sailor predicted yesterday [1]that the weather would frighten[2], turned back towards port.
b. The captain, who the sailor's prediction yesterday [1]about the weather had frightened[2], turned back towards port.
c. The captain, who the sailor's prediction yesterday [1]about the weather had frightened[2] the crew, turned back towards port.
d. The captain believed the sailor's prediction yesterday [1]that the weather would frighten[2] the crew and turned back towards port.

To compare brain activation at events across conditions, we employed subtractions. The following subsection lists the relevant subtractions and elaborates on our reasoning and hypotheses for each.

*6.3.1 Event-based subtractions*

(1) i. A1 – B1
ii. A2 – B2

A comparison of (a) and (b) should reveal processing differences between sentences with intermediate structure and those without. If intermediate traces exist, we predicted left IFG activation at the site of the intermediate trace in (a) – event 1 – but no corresponding activation between in (b).

Both (a) and (b) contain argument traces in the lowest object position. Based on previous literature (i.e. Santi & Grodzinsky 2007), we predicted left IFG activation in the argument trace position for both conditions (a) and (b), and therefore very little difference between the two sentences at event 2.

(2) i. A1 – D1
ii. B1 – D1

This pair of contrasts, like that in 1(i), serves to isolate intermediate traces from other structures that induce similar memory loads without intermediate traces. Comparing both A1 and B1 to a length-matched control in D1 allows us to distinguish what activation is due to the (structural) effect of a trace, and what is due to the (simple memory) effect of an argument waiting for theta-role assignment.

If intermediate traces exist, we should see left IFG activation for A1 over D1, but not for B1 over D1.

(3) i. A2 – D2
ii. B2 – D2

For each of these contrasts individually, we predict activation corresponding to the argument trace in the (a) and (b) conditions, but not in (d). We expect this activation to occur somewhere in the left IFG, again replicating Santi & Grodzinsky (2007).

Taken together, then, we do not predict major activation differences between the results of 3(i) and 3(ii). The contrast between A2 and D2 should look the same as the contrast between B2 and D2 – that is, left IFG activation – because both (a) and (b) contain an argument trace at this point whereas (d) does not.

(4) i. B1 – C1
ii. B2 – C2

We expect no difference in activation between B1 and C1, because the two conditions are identical at this point. (Assuming this prediction is borne out, this helps control for B2 – C2.)

Condition (c) contains an illegal filler at the point of the argument trace in (b). B2 – C2, then, should show significant processing differences reflecting the difference between normal gap-filling and the violation. This should isolate the argument trace in (b) as well as shed light on where in the brain the ungrammaticality in (c) is processed, and which regions are involved in attempts at reinterpretation.

*6.3.2 Sentence-based subtractions*

(5) A – B: This will help us further isolate intermediate traces from long-distance dependencies more generally.

(6) A – D: This will isolate activation due to both intermediate and argument traces.

(7) B – D: This will isolate argument traces (clarifying the potential left IFG activation predicted for contrast 6).

(8) B – C: This will isolate argument traces (in B) as well as shed light on where the ungrammaticality in C is processed.

6.4 *Procedure*

Sentences were presented visually in a paced word-by-word format, each word lasting 500 milliseconds. We arrived at this presentation timing after extensive piloting.

We were constrained to paced reading by individual differences in reading speed: we could not present the whole sentence at once, as it would then be impossible to determine the exact time when a subject reached the individual events within the sentence. But due to the length and complexity of the sentences, we were concerned with determining the ideal pacing and grouping for the words. To this end, we conducted dry-run pilot studies using stimuli and comprehension questions presented on a computer monitor using Tempo software. We evaluated performance based on the number and distribution of errors on comprehension questions. Note that although it was preferable to minimize errors, a certain number of errors were acceptable as long as they occurred equally across sentence conditions (a), (b) and (d). Following the task, subjects were asked how they felt about the difficulty and the timing of the presentation, as well as if they had developed any strategies for answering the questions.

Comprehension questions paired the matrix subject or embedded subject with the matrix verb or embedded verb, to create a pseudo-randomized mix of pairings and yes/no responses.

The first pilot presented each sentence in three large chunks, for example:

(34) The captain who the sailor predicted/
that the weather would frighten/
turned back towards port.

The whole sentence lasted five seconds, with the timing of the segments 1600 msec/2200 msec/1200 msec, respectively. Although these chunks seem like a natural way of dividing the sentence, subjects performed significantly worse on comprehension questions that tested the embedded subject or verb. Indeed, many went so far as to say that they had developed a strategy to answer the question by only paying attention to segments 1 and 3. Of course, this was an extremely undesirable outcome: if subjects fail to read and pay attention to the whole sentence, they do not create a complete mental syntactic representation, rendering the experiment largely pointless.

Our next version attempted to invoke a more even attention pattern across the entire sentence by using smaller chunks. For example:

(35) The captain/
who/
the sailor/
predicted/
that/
the weather/
would frighten/
turned back/
towards port.

Each chunk lasted 500msec, for a total of four to five seconds (depending on the length of the matrix predicate). Error rates were similar to the initial version, and subjects felt that the presentation was choppy and unnatural due to the variation in length of the chunks.

Finally, we created a purely word-by-word presentation with each word lasting 500 msec. This mimicked the presentation from Cooke et al. (2001), who found no difference between 500 and 750 msec per word. Although the presentation did not feel entirely natural, it was "predictably unnatural." Performance on the behavioral task indicated that error rates were lower and evenly distributed across conditions, the desired result.

The final version of the experiment was written using E-Prime software.

*6.4.1 Pre-scan*

Prior to the experiment, we obtained informed consent from each subject. Subjects were exposed to practice items in two phases. First, they read several practice sentences out loud, in order to familiarize themselves with the structure and learn to mentally impose the helpful prosodic contours necessarily absent in the word-by-word, visual presentation. Next, we presented subjects with a half-run (12 sentences) on a laptop running E-Prime, so that the display was identical to what subjects would see inside the scanner. Comprehension questions followed both paper and computer practice items, and subjects were asked to answer the questions out loud so we could check their performance. All practice items were unique, mimicking the structure of the experimental items but not replicating any of them.

*6.4.2 Scan*

Sentences were presented in 10 runs of 24 sentences each. Each run contained six sentences from each of the four conditions. Sentences from the same quadruple never appeared together in the same run. Sentences within runs appeared in an order that was randomized across conditions. Comprehension questions were written so that each run contained an equal number of yes/no correct responses.

Each subject saw the 10 runs in a different, random order. This was done to control for any learning effects at the beginning or fatigue at the end that might weaken performance.

Comprehension questions appeared following sentences of condition (a), (b) and (d), and were displayed for 4000 msec. Subjects were instructed to respond as quickly as possible, but to prioritize accuracy over speed. Condition (c) sentences were followed by a "+" fixation mark, which subjects were instructed to focus on until the next sentence began.

We were able to track subjects' responses in real-time, and offer feedback and encouragement in between runs.

*6.4.3 fMRI data acquisition*

Imaging was performed using a 1.5 T Siemens Magnetom Sonata syngo MR A30 system. Subjects were placed supine in the MR scanner, and their heads were placed within the standard quadrature head coil. Care was taken to ensure that the subjects were looking straight up; foam was placed on either side of the head to minimize motions. Also to minimize motion, a piece of tape was placed across the forehead of each subject and attached to the head holder within the head coil. Once inside the magnet, subjects could view a back-projection screen placed at their feet, through a mirror assembly in the head coil. A fiber optic button box held in the right hand was used to record the subject's responses, and MR compatible headphones were used to reduce scanner noise and allow for communication between the subject and experimenters between runs. An LCD panel was used to project output from a Dell desktop computer running E-Prime onto the back-projection screen.

Four structural scans were obtained, three before and one after the functional task. The parameters for these scans were as follows:

1. Voxel size 2.2 x 1.1 x 10.0mm, FoV read 280mm, slice thickness 10mm, TR 20ms, TE 5ms, flip angle 40 deg.
2. Voxel size 1.3 x 0.9 x 5.0mm, FoV read 240mm, slice thickness 5mm, TR 500ms, TE 7.7ms, flip angle 90 deg.
3. Voxel size 1.0 x 0.8 x 5.0mm, FoV read 200mm, slice thickness 5mm, TR 485 ms, TE 11ms, flip angle 90 deg.
4. Voxel size 1.8 x 1.3 x 1.3mm, FoV read 340mm, slice thickness 1.3mm, TR 24ms, TE 4.66ms, flip angle 45 deg.

The parameters for the functional (BOLD) scan were as follows: voxel size 3.4 x 3.4 x 5.0mm, FoV read 220mm, 28 slices, slice thickness 5mm, TR 2000ms, TE 30ms, flip angle 80 deg, BOLD threshold 4.00.

*6.4.4 fMRI data analysis*

Data were motion corrected using SPM99 and aligned to the reference anatomical scans using the Yale BioImage Suite software package. The data were normalized and smoothed to account for variations in the location of activation across participants. Intra-sentential event-based analyses were conducted using regressions that produced activation maps based on the complementizer onset and verb offset definitions for events 1 and 2, respectively. Whole-sentence analyses used regressors encompassing the onset of the first word through the offset of the last word for each sentence.

Data were corrected using Monte Carlo clusters. For those subtractions that met the correction for $p < .05$, $p < .01$ was also run.

## 7. Results

*7.1 Behavioral data*

The overall accuracy rate on comprehension questions was 81.5 percent. Out of 900 sentences per condition, subjects made 155 errors on A sentences, 167 errors on B sentences and 177 errors on D sentences. There were no significant differences in number of errors between conditions ($p = 0.46$ for the A/B comparison, 0.18 for the A/D comparison and 0.55 for the B/D comparison). Due to the complexity of the sentences as well as the duration of the study (90 fatigue-inducing minutes), the somewhat elevated error rate is not entirely surprising. Crucially, the lack of significant differences in error rates between conditions suggests that any activation effects are due to the syntactic composition of the sentences and not simply the difficulty of comprehension.

*7.2 Imaging data*

The images presented here represent the most illustrative slices from whole-brain maps of 20 slices each. Note that the slices are shown in the radiological convention, so image-right is actually patient-left, and vice versa. Red indicates areas that were more active in the subtractor while blue indicates areas more active in the subtractee. (The subtractee is subtracted from the subtractor.) Unless otherwise noted, all subtractions met the Monte Carlo corrected cluster for $p < .01$.

Event-based subtractions are presented on the next page.

*7.2.1 Event-based subtractions:*

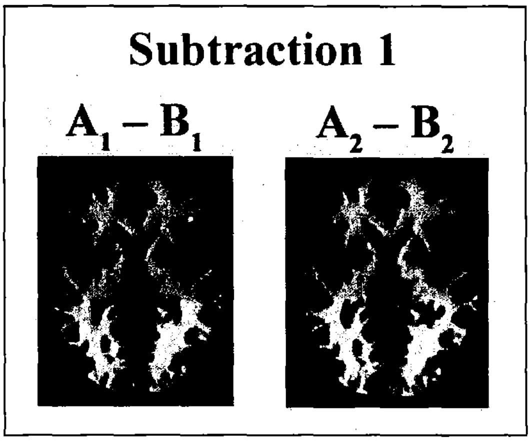


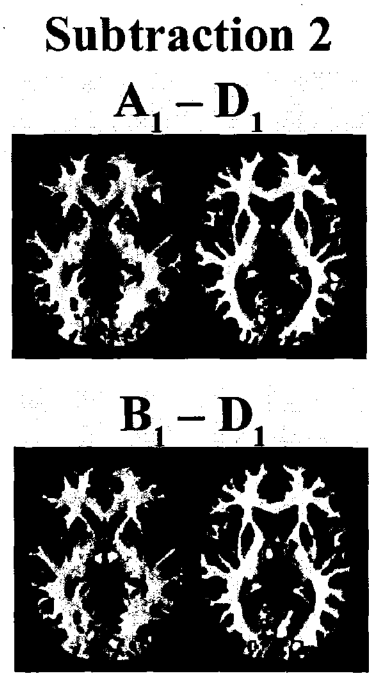


*Subtraction 1.* Neither the (A1 – B1) nor the (A2 – B2) subtraction met the Monte Carlo cluster for $p < .05$, so the figures here represent uncorrected clusters of 10 at $p < .05$. Here, we see a small amount of left IFG activation at event 1 for A over B. There is no significant difference in the left IFG for event 2.

*Subtraction 2.* An analysis of event 1 shows increased activation in both left IFG and left pSTS for A over D, but not for B over D. In both A and B, we see temporo-occipital activation not present in D.

(3)

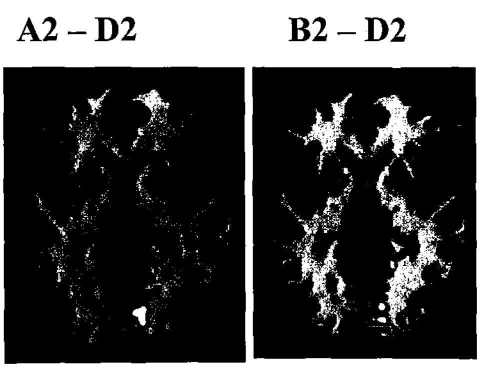


At event 2, there is essentially no difference in left IFG activation between A and B as compared to D. We continue to see temporo-occipital activation in A and B over D.

(4) **B1 – C1** **B2 – C2**

Here we see no activation difference at event 1 between B and C (as expected, because the sentences are identical at this point). At event 2 there are significant differences in multiple brain areas, including left IFG, left pSTS, temporo-occipital regions and right-hemisphere homologues.

*7.2.2 Whole-sentence subtractions:*

(5) **A – B**

This subtraction, like (1), did not meet the Monte Carlo corrected cluster for $p < .05$, so shown here are results for a cluster of 10 at $p < .05$. Though clearly there is very little difference between the two conditions when viewed across the whole sentence, there is a hint of increased activation for A near both the left IFG and left pSTS. For the left IFG, the precise location of the activation seems to be deeper in the brain here than in subtractions (1) and (2).

(6) **A – D**

Here we find increased activation in the left IFG as well as its right-hemisphere homologue, along with prefrontal cortex, for A over D.

(7) **B – D**

This subtraction largely mimics subtraction (6): increased activation in the left IFG as well as prefrontal cortex for B over D.

(8) **B – C**

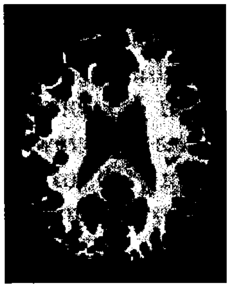

This figure represents a Monte Carlo correction at $p < .05$. Here we see increased activation in the left IFG for B over C, and a large network of right hemisphere regions for C over B.

**8. Discussion**

This section will discuss the results presented above in the context of our initial hypothesis. As outlined in our predictions, we will focus on the effects of both intermediate and argument traces, as well as more general effects observed.

8.1 *Intermediate traces*

Subtraction (2) provides the most compelling support for our hypothesis, suggesting that intermediate structure is, in fact, a feature of real-time syntactic composition, and can be localized to the left IFG.[4] Also in this subtraction, we found left pSTS activation for A over B at the intermediate position.

Similarly, the outcome of subtraction (1) shows increased activation in the left IFG for A over B, and subtraction (5) suggests increased activation in both the left IFG and left pSTS for A over B. (Note that the deeper activation in subtraction (5) is not entirely surprising, as deeper brain areas – including the subjacent white matter, operculum and insula – have been implicated for syntactic processing along with the traditionally cortical Broca's area in lesion studies, electroencephalographic and imaging; see Aboitiz et al., in Grodzinsky & Amunts 2006, for a summary.) Though these outcomes did not reach significance, their trends are bolstered by the outcome of subtraction (2).

That both the left IFG and left pSTS showed heightened activation in A supports the general intuition that A requires more structure-building in real-time. We did not isolate the left IFG even at event 1, suggesting that the left pSTS might play a more immediate role in syntactic structure-building than previously thought.

8.2 *Argument traces*

Whole-sentence subtractions (6), (7) and (8) replicate previous literature (i.e. Ben-Shachar et al. 2003, 2004, Santi & Grodzinsky, 2007), showing increased left IFG activation for sentences with long-distance dependencies compared to those without.

Subtraction (8) is particularly interesting in that it represents the comparison of normal gap-filling to a violation. It seems that the left IFG is activated only when gap-filling proceeds as

[4] There is still the possibility that the left IFG activation for A represents argument gap-filling (as opposed to intermediate gap-filling) due to a subject-relative garden-path. However, this is highly unlikely not only due to the semantics of the intermediate verbs, but also the cumulative effect of exposure to many sentences of similar types, none of which were subject-relatives. After the first run or two, subjects should have grown to expect the object-relative structure.

expected: the ungrammaticality in C did not trigger the left IFG, instead recruiting several regions in the right hemisphere, which could reflect the processor's attempts to interpret the impossible sentence.

Interestingly, the argument trace effect did not appear in any event-based subtractions. In subtraction (1), this was expected, as we predicted equal left IFG activation in A and B for event 2. However, it is somewhat surprising that no effect surfaced in subtraction (3). Perhaps the time course for the gap-filling activation cannot be temporally resolved to this particular moment as we originally expected, but is only visible across a more extended interval.

8.3 *Other effects*

We found that networks of both prefrontal and temporo-occipital regions are recruited in sentences with long-distance dependencies (A and B) compared to those with no dependencies (D). This is visible in subtractions (2), (3), (4), (6) and (7). The temporo-occipital activation replicates Cooke et al. (2001). Along with the prefrontal areas, this activation most likely reflects domain-general processes of "active maintenance" working memory and attentional demands (i.e. Cohen et al. 1997).

As mentioned above, subtractions (4) and (8) reveal a network of widespread activation in the right hemisphere in ungrammatical sentences. This most likely reflects the detection of the violation and subsequent attempts at reinterpretation.

## 9. Conclusion

In the only neurological investigation of intermediate traces we are aware of to date, our findings lend exciting support to the role of the left IFG in building syntactic composition.

This study used the visual (reading) modality, but we have good reason to believe that the effects, especially in the left IFG, can be generalized to auditory parsing as well (i.e., Constable et al. 2004).

How can we model the building of this enriched composition? Let us return briefly to our discussion of theoretical implications for processing. If we accept the SIH as a working model for the function of the left IFG, the results presented here suggest that *wh*-dependencies across two or more clauses trigger some sort of marker at the clausal boundary that works to facilitate eventual integration/theta-role assignment. The presence of this marker is consistent with more than one theoretical model of syntactic structure: it could be a "special status" trace in the vicinity of the complementizer (GB), or it could persist more or less all the way through the two clauses (SS/HPSG). As we saw in section 3, both theories predict our results but for slightly different reasons:

(36) Theoretical Predictions for Processing

| **At event 1:** | Left IFG activation in A? | Left IFG activation in B? |
|---|---|---|
| GB | Yes *(special status trace)* | No *(no clause boundary)* |
| SS/HPSG | Yes *(slash-feature path)* | No *(inside the NP island)* |

As outlined in the SIH model, we believe that our results can be used to reconcile the movement-based (i.e. Ben-Shachar et al. 2003, 2004, Santi & Grodzinsky 2007), semantic role-based (i.e. Bornkessel et al. 2005), and working memory-based (i.e. Fiebach et al. 2005) hypotheses for the function of the left IFG. It is most likely a combination of all three: the left

IFG builds syntactic structure in anticipation of theta-role assignment. This can be seen in sentences with long-distance dependencies arising from "movement," but it can also be seen in sentences where derivation of argument structure is made more complex in some other way. Now, we have shown that it can even be seen in real-time intermediate syntactic positions that trigger enriched composition to maintain dependencies.

How might this need for extra structure be encoded for the processor? We know that in all statement contexts, the appearance of a *wh*-word not in-situ entails at least one embedded clause. But perhaps if a second level of embedding is introduced before the *wh*-filler gets its theta-role, the processor knows it must go inside this clause in search of the theta-role. This decision as to a search path engages both the left IFG and left pSTS. This is, of course, a preliminary

Another main contribution of this investigation was methodological: we pioneered a novel intra-sentential event-based regression analysis, which can be used in future studies to achieve an ever more more precise view of the trajectory of sentence comprehension. Future work might examine additional intra-sentential events to further clarify the time course of left IFG and left pSTS activation in sentences with differing structures. (This will be increasingly useful in conjunction with technological developments improving the temporal resolution of fMRI.)

Worth noting is that in tandem with the fMRI study, we are in the process of investigating intermediate traces via priming, in a paradigm similar to Shvartsman (2007) as discussed in section four. In this study, A-type sentences serve as experimental items, while we use traceless syntactic structures as controls. Probes (related, unrelated and non-word) are placed at the intermediate trace and argument trace as well as control locations within the sentence. If the left IFG activity seen in the present study corresponds to subconscious reactivation (or simply continued activation) of the antecedent, we expect priming for related words at the intermediate trace site, along with the argument trace site (replicating previous literature). At this time, we are still in the testing phase and no results are available yet.

If we do show priming in the intermediate location, this would provide psycholinguistic evidence for intermediate structure in real-time that would nicely complement the imaging data presented here. If we are to continue in this line of research, another potential cross-modal study might introduce B-type sentences to compare priming at event 1 in A versus B, further teasing apart structural implications for processing. (According to our current theoretical models, we would not expect priming at event 1 in B. Although, as noted in footnote 1, pp 15, SS/HPSG might predict priming inside the embedded VP, which might make for an interesting experiment.)

Combining neurolinguistic and psycholinguistic evidence gives us the opportunity to link the physical substrates of brain-language relations with the result of their metabolic activity: the functioning of the mind. It is at this intersection that the field finds its application to medicine, computer science, and the like.

At this point, it is clear that determining the exact process of syntactic composition as it gives rise to argument structure will require both novel theoretical developments and additional experimental evidence. It is our hope that the ideas, methods, and data presented here will inform and inspire future inquiry into the fascinating, uniquely human capacity of language.

References


Aboitiz, F., Garcia, R., Brunetti, E. & Bosman, C. 2006. "The origin of Broca's area and its connections from an ancestral working memory network." In Grodzinsky, Y. & Amunts, L. (eds.) *Broca's Region.* New York: Oxford University Press.

Avrutin, S. 2006. "Weak syntax." In Grodzinsky, Y. & Amunts, L. (eds.) *Broca's Region.* New York: Oxford University Press.

Ben-Shachar, M., Hendler, T., Kahn, I., Ben-Bashat, D. & Grodzinsky, Y. 2003. The neural reality of syntactic transformations: evidence from functional magnetic resonance imaging. *Psychological Science 14:5,* 433-440.

Ben-Shachar, M., Palti, D. & Grodzinsky, Y. 2004. Neural correlates of syntactic movement: converging evidence from two fMRI experiments. *NeuroImage 21,* 1320-1336.

Beretta, A., Piñango, M., Patterson, J. & Harford, C.1996. Recruiting comparative crosslinguistic evidence to address competing accounts of agrammatic aphasia. *Brain and Language 67:3,* 149-168.

Blumstein, S.E., Goodglass, H., Statlender, S. & Biber, C. 1983. Comprehension strategies determining reference in aphasia: A study of reflexivization. *Brain and Language 18,* 115-127.

Bornkessel, I., Zysset, S., Friederici, A., von Cramon, D. & Schlesewsky, M. 2005. Who did what to whom? The neural basis of argument hierarchization during language comprehension. *NeuroImage 26,* 221-233.

Burchert, F., Swoboda-Moll, M. & De Bleser, R. 2005. Tense and agreement dissociations in German agrammatic speakers: underspecification versus hierarchy. *Brain and Language 94:2,* 188-199.

Caramazza, A. & Zurif, E. 1976. Dissociation of algorithmic and heuristic processes in language comprehension: evidence from aphasia. *Brain and Language 3:* 572-582.

Carnie, A. 2007. *Syntax: A Generative Introduction.* Malden, Mass.: Blackwell Publishing.

Chomsky, N. 1973. "Conditions on transformations." In Anderson, S. & Kiparsky, P. (eds.) *A Festschrift for Morris Halle.* USA: Holt, Rinehart and Winston, Inc.

Chung, S. 1982. Unbounded dependencies in Chamorro grammar. *Linguistic Inquiry 13:* 39-77.

Chung, S. 1994. *Wh*-agreement and "referentiality" in Chamorro. *Linguistic Inquiry 25:* 1-44.

Clifton, C. Jr. & Frazier, L. 1989. Comprehending sentences with long-distance dependencies. In Carlson, G. & Tanenhaus, M. (eds.) *Linguistic Structure in Language Processing.* Dordrecht: Kluwer.

Cohen, D., Perlstein, W., Braver, T., Nystrom, L., Noll, D., Jonides, J. & Smith, E. 1997. Temporal dynamics of brain activation during a working memory task. *Nature 386:10,* 604-608.

Constable, R., Pugh, K. Berroya, E., Mencl, W., Westerveld, M., Ni, W. and Shankweiler, D. 2004. Sentence complexity and input modality effects in sentence comprehension: an fMRI study. *NeuroImage 22:1,* 11-21.

Cooke, A., Zurif, E., DeVita, C., Alsop, D., Koenig, P., Detre, J., Gee, J., Piñango, M., Balogh, J. & Grossman, M. 2001. Neural basis for sentence comprehension: grammatical and short-term memory components. *Human Brain Mapping 15,* 80-94

Culicover, P. & Jackendoff, R. 2005. *Simpler Syntax.* New York: Oxford University Press.

De Vincenzi, M. 1991. Syntactic parsing strategies in Italian. Dordrecht: Kluwer.

Erteschik, N. 1973. On the nature of island constraints. Doctoral dissertation, MIT.
Fiebach, C., Schlewesky, M., Lohmann, G., von Cramon, D. & Friederici, A. 2005. Revisiting the role of Broca's area in sentence processing: syntactic integration versus syntactic working memory. *Human Brain Mapping 24,* 79-91.
Fodor, J. 1995. "Comprehending sentence structure." In Osherson, D., Gleitman, L. & Liberman, M. (eds.) *An Invitation to Cognitive Science: Language.* Cambridge, Mass.: The MIT Press.
Frazier, L. & Clifton, C. 1989. Successive cyclicity in the grammar and the parser. *Language and Cognitive Processes 4:* 93-126.
Friedmann, N. & Grodzinsky, Y. 1997. Tense and agreement in agrammatic production: pruning the syntactic tree. *Brain and Language 56,* 397-425.
Georgeopoulous, C. 1985. Variables in Palauan syntax. *Natural Language & Linguistic Theory 3,* 59-94.
Gibson, E. 1998. Linguistic complexity: locality of syntactic dependencies. *Cognition 68,* 1-76.
Gibson, E. 2000. The dependency locality theory: a distance-based theory of linguistic complexity. In Miyashita, Y., Marantz, A. & O'Neil, W. (eds.) *Image, Language, Brain.* Cambridge, Mass.: MIT Press.
Gibson, E. & Warren, T. 2004. Reading-time evidence for intermediate linguistic structure in long-distance dependencies. *Syntax 7:1,* 55-78.
Ginzburg, J. & Sag, I. 2000. *Interrogative Investigations: The Form, Meaning and Use of English Interrogatives.* Stanford, Calif.: CSLI Publications.
Grodner, D., Watson, D. & Gibson, E. 2000. Locality effects on sentence processing. Paper presented at the 13th CUNY Sentence Processing Conference, University of California, San Diego.
Grodzinsky, Y. 1986. Language deficits and the theory of syntax. *Brain and Language 27,* 135-159.
Grodzinsky, Y. 1995. A restrictive theory of agrammatic comprehension. *Brain and Language 50,* 27–51.
Hickok, G., Zurif E. & Canseco-Gonzales, E. 1993. Structural description of agrammatic comprehension. *Brain and Language 45,* 371–395.
King, J. & Just, M. 1991. Individual differences in syntactic processing: the role of working memory. *Journal of Memory and Language 30:* 580-602.
McCloskey, J. 2000. Quantifier float and *wh*-movement in an Irish English. *Linguistic Inquiry 31:* 57-84.
Pickering, M. & Barry, G. 1991. Sentence processing without empty categories. *Language and cognitive processes 6:* 229-59.
Piñango, M.M. 2000. "Syntactic displacement in Broca's agrammatic aphasia." In Bastiaanse, R. & Grodzinsky, Y (eds.) *Grammatical Disorders in Aphasia: A Neurolinguistic Perspective.* London, England: Whurr Publishers Ltd.
Pollard, C. & Sag, I. 1994. *Head-Driven Phrase Structure Grammar.* Chicago: University of Chicago Press.
Roder, B., Stock, O., Neville, H., Bien, S. & Rosler, F. 2002. Brain activation modulated by the comprehension of normal and pseudo-word sentences of different processing demands: a functional magnetic resonance imaging study. *NeuroImage 15,* 1003-1014.
Santi, A. & Grodzinsky, Y. 2007. Working memory and syntax interact in Broca's area. *NeuroImage 37,* 8-17.

Santi, A. & Grodzinsky, Y. 2008. The battle for Broca's region. *Trends in Cognitive Sciences 12:12,* 474-480.

Shapiro, L., Zurif, E. & Grimshaw, J. 1989. Verb processing during sentence comprehension: contextual impenetrability. *Journal of Psycholinguistic Research 18:2,* 223-243.

Shapiro, L., Gordon, B., Hack, N. & Killackey, J. 1993. Verb-argument structure processing in complex sentences in Broca's and Wernicke's aphasia. *Brain and Language 45:3,* 423-47.

Shvartsman, M. 2007. Adjudicating between linguistic theory types based on evidence from priming: the case of successive cyclicity. Undergraduate thesis, Yale University Dept. of Linguistics (unpublished).

Stowe, L. 1986. Evidence for on-line gap location. *Language and Cognitive Processes 1*, 227-245.

Stromswold, K., Caplan, D., N. Alpert N. & Rauch, S. 1996. Localization of syntactic comprehension by positron emission tomography, *Brain and Language 52*, 452–473.

Swinney, D. & Zurif, E. 1995. Syntactic processing in aphasia. *Brain and Language 50:2,* 225-239.

Thornton, R. 1990. Adventures in long-distance moving: the acquisition of complex *wh*-questions. Doctoral dissertation, University of Connecticut, Storrs.

Thornton, R. 1995. Referentiality and *wh*-movement in child English: juvenile D-linkuency. *Language Acquisition 4:1,* 139-175.

Torrego, E. 1984. On inversion in Spanish and some of its effects. *Linguistic Inquiry 15,* 103-129.