8 Understanding Sentence Structure and Meaning

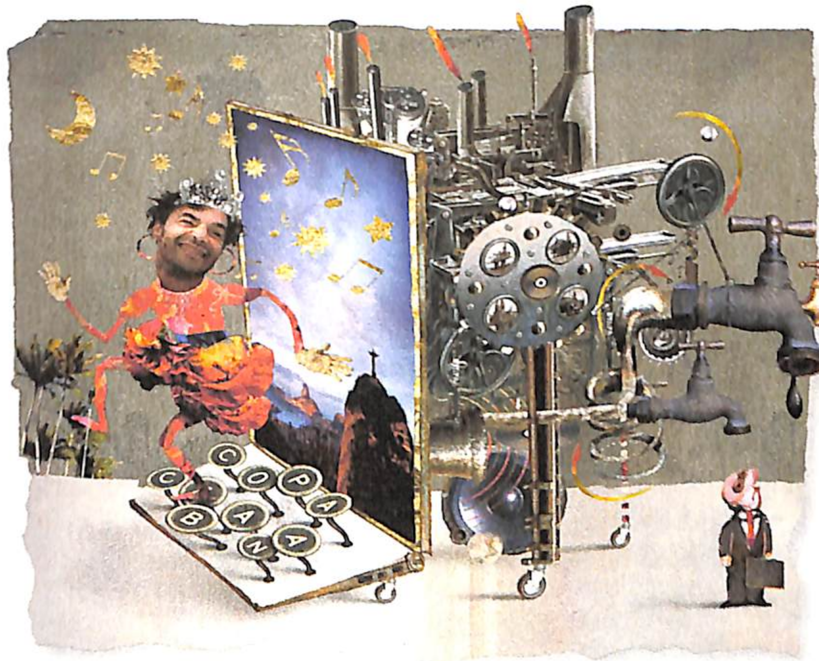

As we saw in Chapter 6, the mind is exquisitely adept at relating syntactic structure to meaning, able to compute the meanings of long, complex sentences, even those containing numerous clauses nested within each other and coding intricate relationships among their elements, of which this particular sentence is an excellent example.

Despite our great parsing prowess, we still stumble over some sentences. We've all had the experience of finding ourselves re-reading certain sentences over and over, not quite sure how to unravel their meanings even though none of the words in them are especially complicated or unfamiliar. Some sentences just feel knotty or clunky. They're the sentences that an English teacher or an editor might single out with the remark "awkward sentence structure" or "this feels clumsy."

Let's see if you have an intuitive sense of the kinds of sentences that strain language comprehension. Of the perfectly grammatical sentences listed on the next page, some slide through the mind with ease, while others seem to bunch up or create hard-to-read word piles. Make a note of the ones you find somewhat taxing or confusing. For the time being, rely solely on your editorial instinct; by the end of this chapter, you should have some scientific understanding of *why* they cause problems for the reader. (And if you ever find yourself working as an editor, you'll be able explain to your authors exactly what's gone awry with many of their bad sentences.)

The boy watched the ball with the yellow stripe float down the river.

Susanne put the toy soldier in the box into the cupboard.

The soup and the vegetables simmered lightly on the stove while the hostess served some wine.

The patient told the doctor that he was having some trouble with about his sore knee.

The baker is going to sue the plumber who botched the installation of his new pipes last week.

The suspect who was seen last night at the crime scene has confessed to the murder.

The mouse the cat chased keeled over.

The boy who stalked the prom queen has been caught.

The teachers taught new math techniques passed the exam with flying colors.

Woody confided to his therapist that he occasionally thought about murdering his mother.

While the pop star sang the national anthem played in the background.

The cruel man beat his puppy with a thick stick.

Samantha explained to her son's dentist that the boy's grandparents kept feeding him candy.

Friedrich cooked the ribs and the corn had already been roasted.

The farmer slaughtered the goose that had intimidated his hens.

The gang leader hit the lawyer with a wart.

As Marilyn smiled the photographers snapped dozens of pictures.

The bartender told the detective that the suspect will try to escape the country yesterday.

The coffee in the red mug was already cold by the time the secretary had time to drink it.

The administrator who the intern who the nurse supervised had accused fudged the medical reports.

At least nine of these sentences would be recognized by most psycholinguists as potential troublemakers, containing elements known to cause problems for readers or hearers. Obviously, being able to recognize what makes sentences tricky to read is exceptionally useful if you're a writer or editor. But if you're a psycholinguist, it also gives you some nifty insights into how human beings go about structuring the words of incoming speech or text into sentences with complex meanings. We can figure out a great deal about how human language comprehension works in real time by observing it at its limits and seeing where it breaks down. Different theories of sentence processing make different predictions about which sentences will be harder to process than others.

In this chapter, you'll get a chance to explore the nature of difficult sentences, whether written or spoken. But this isn't just a story about awkward sentences. As you'll see, in addition to learning how glitches in processing come about, you will ultimately gain an appreciation of why it is that most of the time, comprehension glides along perfectly smoothly.

8.1 Incremental Processing and the Problem of Ambiguity

Processing sentences on the fly

As you saw in Chapter 7, the word recognition system is not especially timid when it comes to making rapid, reasonable guesses about which words are in the midst of being pronounced; it begins to link sounds with possible words and their meanings from the very first moments of utterance, activating a large number of possibilities in the memory store, and gradually winnowing these down to the very best candidates.

But what about sentences? Unlike words, their meanings don't just depend on retrieving meanings that are pre-stored in memory. Remember that the meaning of each sentence has to be constructed anew based on the syntactic relationships among all the words in the sentence. (Otherwise, we wouldn't be able to understand sentences we'd never heard before, having never had the opportunity to memorize their meanings.) This process of structure-building during comprehension is referred to as **parsing**. (Psycholinguists often use "the parser" as a convenient term for the collection of structure-building mechanisms and procedures; the term does not refer to an individual person.) How long does it take for the parser to build these meaningful structures? Does each word in the sentence have to be uttered and fished out of memory first before syntactic grouping and structuring can begin to take place?

parsing The process of assigning syntactic structure to the incoming words of a sentence during language comprehension. The structure-building mechanisms and procedures collectively are often referred to as "the parser."

incrementality The property of synthesizing and building meaning "on the fly" as individual units of speech come in, rather than delaying processing until some amount of linguistic material has accumulated.

shadowing task An experimental task in which subjects are asked to repeat the words of a speaker's sentence almost as quickly as the speaker produces them.

It seems not. The parser happens to be just as eager as the word recognition system in generating guesses about possible meanings, even on the basis of very partial information from the speech stream. Like word processing, understanding sentences is an exercise in **incrementality**—that is, meaning is built on the fly as the speech comes in, rather than being delayed until some amount of linguistic material has accumulated first. One very intuitive way of showing that understanding follows hot on the heels of the uttered speech is by means of a **shadowing task**, in which subjects are asked to repeat the words of a speaker's sentence almost as quickly as the speaker produces them. People who are very good at this can follow at a lag of about a quarter of a second—roughly a lag of one syllable behind the speaker. This suggests that within hundreds of milliseconds of a word, not only has its meaning been recognized, but it has been integrated with the syntactic structure and meaning of the sentence. We know that people are actually analyzing the sentence's meaning rather than just parroting words or even just the sounds of words, because shadowing gets considerably slower when meaningful sentences are replaced with nonsensical sentences (but with recognizable words), and slower yet if the sentence that has to be repeated is in a foreign language or uses made-up words.

WEB ACTIVITY 8.1

Processing language on the fly Try your hand at a shadowing task. You'll first be able to listen to someone shadowing speech, and then have an opportunity to shadow a variety of sensible, senseless, and foreign-language sentences.

But there's a downside to this "hungry" style of language processing. Like an overeager student who blurts out the answer before the teacher has finished asking the question, the parser's guesses, based on only partial information from the sentence, may not all turn out to be the right ones. You might remember that this was also a factor with the eager word recognition system. Along with the correct target word (for example, *candy*), soundalike words, especially

reduced relative clause A grammatical structure in English involving a relative clause in which certain function words have been omitted (for example the reduced relative clause *raced past the barn* derives from the full relative clause *that was raced past the barn*). This structure often leads to ambiguity.

those that begin with the same sounds, also become activated (*candle, Canada, cantaloupe*). Evidence of such spurious activation can be seen in priming tasks or in patterns of eye movements to visual displays. In other words, extreme eagerness or incrementality in word recognition leads to an explosion of potential ambiguity of meanings, at least for a short period of time. The processing system ultimately has to suppress a number of possible interpretations that started out as being perfectly consistent with the uttered speech. This happens at the level of the sentence as well.

"Garden path" sentences

Successful interpretation is often a matter of not getting tripped up by the incorrect meanings that also become activated during comprehension. But sometimes, one of the alternative meanings causes so much disruption during parsing that it becomes really hard to recover the correct meaning of the sentence. Consider the following, probably the single most famous sentence in psycholinguistics, brought to life by Tom Bever (1970):

> The horse raced past the barn fell.

Figure 8.1 Tree structures illustrating two possible syntactic interpretations of the string *The horse raced past the barn*. . . . (A) A main clause interpretation. (B) An interpretation involving a reduced relative clause, with the anticipation of additional content to be uttered as part of the verb phrase (VP).

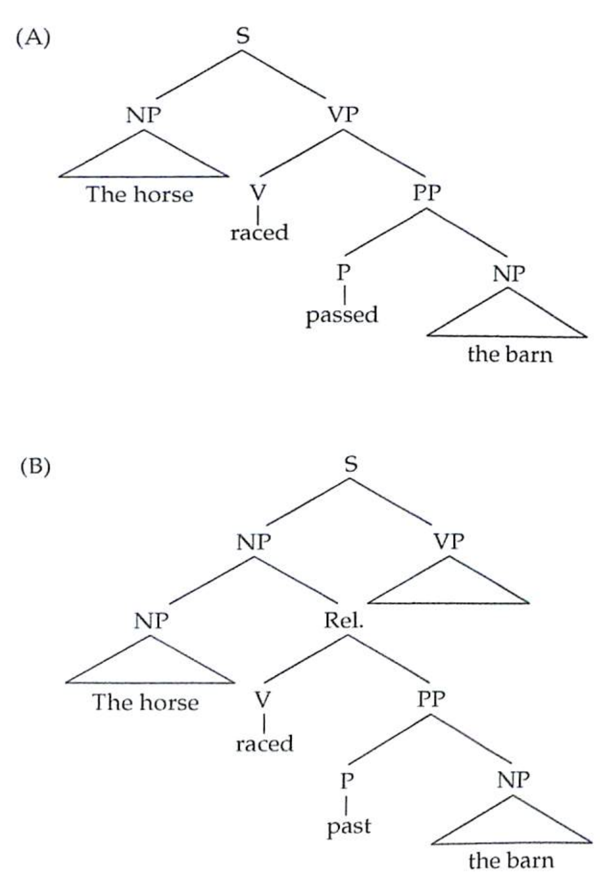

On a first reading, many people think this sentence makes no sense, or is ungrammatical with some words left out. But it's perfectly grammatical. Don't believe me? It means exactly the same thing as this:

> The horse that was raced past the barn (by someone) fell.

Still don't see it? It has exactly the same structure and a very similar meaning to this:

> The horse driven past the barn fell.

Assuming you've eventually been able to parse the sentence correctly, the puzzle becomes this: What is it that makes that first sentence so dastardly difficult to understand, much more so than the other two, even though all of them are permitted by the grammar of English? The answer is that only the first has the potential for a temporary ambiguity; until the very last word of the sentence, it's most natural to understand the sentence as being structured so that the word *horse* is the subject of the verb *raced*—that is, the horse is *doing* the racing, rather than being raced. This interpretation would have worked out just fine if the sentence had continued this way:

> The horse raced past the barn and fell.

Which is why people often tend to have the feeling that the sentence is missing a word or two. But in order to get the right reading for our original sentence, you have to ignore this highly tempting meaning for the first six words of the sentence, and instead interpret them as being structured such that the phrase *raced past the barn* is a separate clause adding more information about the horse—this is called a **reduced relative clause** because it's just like the (non-reduced) relative clause *that was raced past the barn* with some of the function words taken out. In this more complicated structure, notice that *the horse* acts semantically as the direct object of *raced*—that is, it is being raced, rather than doing

garden path sentences Sentences that are difficult to understand because they contain a temporary ambiguity. The tendency is for hearers or readers to initially interpret the ambiguous structure incorrectly, and then experience confusion when that initial interpretation turns out to be grammatically incompatible with later material in the sentence.

the racing—but *also* acts as the subject of the main clause verb *fell*. **Figure 8.1** shows one way in which to graphically capture the two different interpretations. (And, if grammatical terms like *direct object* and *relative clause* aren't part of your daily vocabulary, **Box 8.1** on the next page offers a quick refresher.)

Here are a few more sentences with reduced relative clauses, and with the same potential for ambiguity as our horse-racing example—though a number of them seem to cause somewhat milder reading hiccups, for reasons we'll see later:

> The swimmers drowned in the lake were not found until the following spring.
>
> The general presented copies of the report was aware of the problems.
>
> The boat floated downstream will never get past the rapids safely.

Difficult sentences like these are poetically called **garden path sentences**, since their effect is to lead readers down a garden path, misleading them into interpreting the sentence one way, but then veering off in another direction entirely. In English, temporary syntactic ambiguities aren't limited to cases in which a word sequence can be read either as a reduced relative clause or as belonging to a simple main clause. A generous variety of opportunities exists for garden path sentences. For example:

> After the police stopped the car...

Here, *the car* could be read as either the direct object of the verb *stopped* (as in *After the police stopped the car they noticed its license plate was missing*), or as the subject of the following clause (*After the police stopped the car sped off down the road*). Most likely, the first of these possible continuations feels more graceful to you, while the second is a bit bumpy, presumably because you've given in to the desire to interpret a noun phrase immediately following a verb as a direct object of that verb. This creates problems when you then encounter the second verb *sped off*, which nonsensically appears to have no subject, so you need to go back and reanalyze the sentence so that *the car* is interpreted as its subject.

Now try supplying a plausible continuation for this sentence:

> The married man promised his mistress that he would soon abandon...

If you suggested that the next words should be *his wife*, then you've taken the heavily traveled path of interpreting *that he would soon abandon* as introducing a sentence complement of the verb *promised*—that is, what the married man promised was that he would soon abandon someone, perhaps his wife. But another structure is possible, as evident in this sample continuation:

> The married man promised his mistress that he would soon abandon a diamond ring.

Since it's decidedly bizarre to abandon a diamond ring, the reader is nudged into a different interpretation and, if successful, will eventually settle on a reading in which *that he would soon abandon* is a relative clause, providing additional information about his mistress (similar to *The married man gave the woman that he would soon abandon a diamond ring*.)

Once you start looking, you'll be able to diagnose many awkward-sounding sentences as garden path sentences. For instance, each of the sentences below contains a temporary ambiguity. Try to identify the fragment of the sentence that could be interpreted in two ways, and then think about another possible

BOX 8.1
Key grammatical terms and concepts in English

Subject

The noun phrase (NP) that appears to the left of a verb phrase (VP) and combines with it to form a sentence. The subject is often described as "what the sentence is about." When the sentence is in the active voice, the subject is typically (but not always) the cause or the instigator of the event described by the verb:

Copernicus made an important discovery.

The third man on the left is acting suspiciously.

She died on August 12, 1542.

Direct object

A noun phrase that appears inside of the verb phrase, to the right of the verb. Not all verbs take a direct object; whether a verb allows/requires one or not must be specified as part of the lexical knowledge of that verb. When the sentence is in the active voice, the direct object is usually the entity that is acted upon, or comes into being, as a result of the actions of the subject:

Copernicus made **an important discovery**.

The president fired **his chief of staff**.

The police stopped **the car**.

Cyrano wrote **an eloquent love letter**.

Indirect object

Occasionally, two noun phrases occur inside a verb phrase; in this case, only one is the direct object; the other NP is the indirect object. Rather than expressing the acted-upon entity, the indirect object usually expresses the recipient of the acted-upon thing. An indirect object either appears immediately after the verb or is introduced by a preposition:

Cyrano wrote **Roxanne** an eloquent love letter.

Cyrano wrote an eloquent love letter for **Roxanne**.

Copernicus explained his discovery to **many skeptical theologians**.

The hooded man passed **the bank teller** a note.

Sentence complement

A clause, or a sentence unit, that appears inside the verb phrase, to the right of the verb. It is often (but not always) introduced by the *complementizer* word *that*, and may occur on its own, or together with a noun phrase:

His friends warned Copernicus **that the authorities were planning his arrest.**

The president claimed **he had created many new jobs.**

Main versus subordinate clauses

A sentence may have multiple sentence units contained within it, but usually only one is the main clause: this is the sentence that is at the top of the tree (as in Figure 8.1), while the other subordinate clauses are embedded inside other phrases. Subordinate clauses might appear within the verb phrase as sentence complements, attached to nouns as relative clauses, or introduced by adverbial words such as *although*, *despite*, *after*, etc. Here, the subordinate clauses appear in bold:

The horse raced **while the cow stared**.

Though her audience loved her, Marilyn was riddled with doubt.

The man **who was acting suspiciously** turned out to be an escaped convict.

She died **because she had no money for the operation**.

Conjoined clauses

Two or more constituents of the same type can be conjoined by *and* or *but*, including clauses. When a main clause is conjoined with another main clause, they carry equal weight, and both are considered to be main clauses:

The police stopped the car and the driver jumped out.

Copernicus made a discovery but he hesitated to reveal it.

Relative clause

A sentence unit that is embedded within a noun phrase, usually (but not always) introduced by a relative pronoun such as *who* or *that*:

The boy **who stalked the prom queen** has been caught.

The discovery **that Copernicus made** was controversial.

Notice that, confusingly, *that* can also be a complementizer introducing a sentential complement inside a verb phrase (see the entry "Sentence complement").

BOX 8.1 *(continued)*

Active versus passive voice

This is one of the most commonly misunderstood grammatical concepts. The active voice is simply the default structure for expressing an event or a situation:

Many theologians denounced Copernicus.

The principal drove the prom queen home.

The passive voice is an alternative way of expressing the same event:

Copernicus was denounced by many theologians.

The prom queen was driven home by the principal.

A passive sentence can only be related to an active voice sentence that contains a direct object. The passive sentence rearranges the order of the noun phrases so that the direct object from the active version becomes the subject of the passive version, and the subject from the active sentence (usually the instigator of the action) becomes embedded in a prepositional phrase introduced by the word *by*. In the process, the auxiliary verb *to be* (or occasionally *to get*) is inserted between the subject and the verb, and the verb appears in the past participle form (for example, *drove* becomes *was driven*; past tense verbs that end in *-ed* appear in exactly the same form as past participles—more on that later in the chapter). Thus:

The orangutan is being fed too often by the staff.

Copernicus got arrested for his controversial ideas.

The prom queen is chosen by the students every year.

A noteworthy feature of passive sentences is that the *by* phrase may be dropped, thereby leaving the instigator of the action completely implicit, unspoken, or unknown:

The orangutan is being fed too often.

Mistakes were made.

Copernicus was denounced.

Many people mistake other structures for passive sentences because they contain superficial similarities such as the presence of *to be*. But if a sentence does not contain all of the elements described above, it is not passive. Try identifying the passive sentences among the following examples:

(a) The maid is stealing money.

(b) Those ideas were considered revolutionary.

(c) The prisoners were being tortured on a daily basis.

(d) Now you are just being picky.

(e) Copernicus got famous for his controversial ideas.

(f) Every car is searched at the border.

For many relative clauses, it's possible to drop the relative pronoun (such as *that* or *who*), leading to a reduced relative clause. When this happens to a relative clause that's in the passive voice, the auxiliary *to be* (or *to get*) also gets dropped:

The astronomer (who got) arrested for his ideas was controversial.

The prom queen (who was) driven home by the principal was drunk.

Answer to the passives quiz: The passive sentences are (b), (c), and (f).

way to continue the sentence that would be consistent with that second reading of the ambiguous string of words:

The investigation revealed the error resulted from gross negligence.

Sam ate the hot dog and the vegetables went untouched.

The gangster shot the cop with the gun.

This company only hires smart women and men who may be dumb but who have clout.

Visiting relatives can be annoying if they overstay their welcome.

The government plans to raise taxes were defeated.

The army houses soldiers and their families.

Once you've identified the potential ambiguities, you can check your interpretations against **Figure 8.2**, where some of these examples are graphically mapped out.

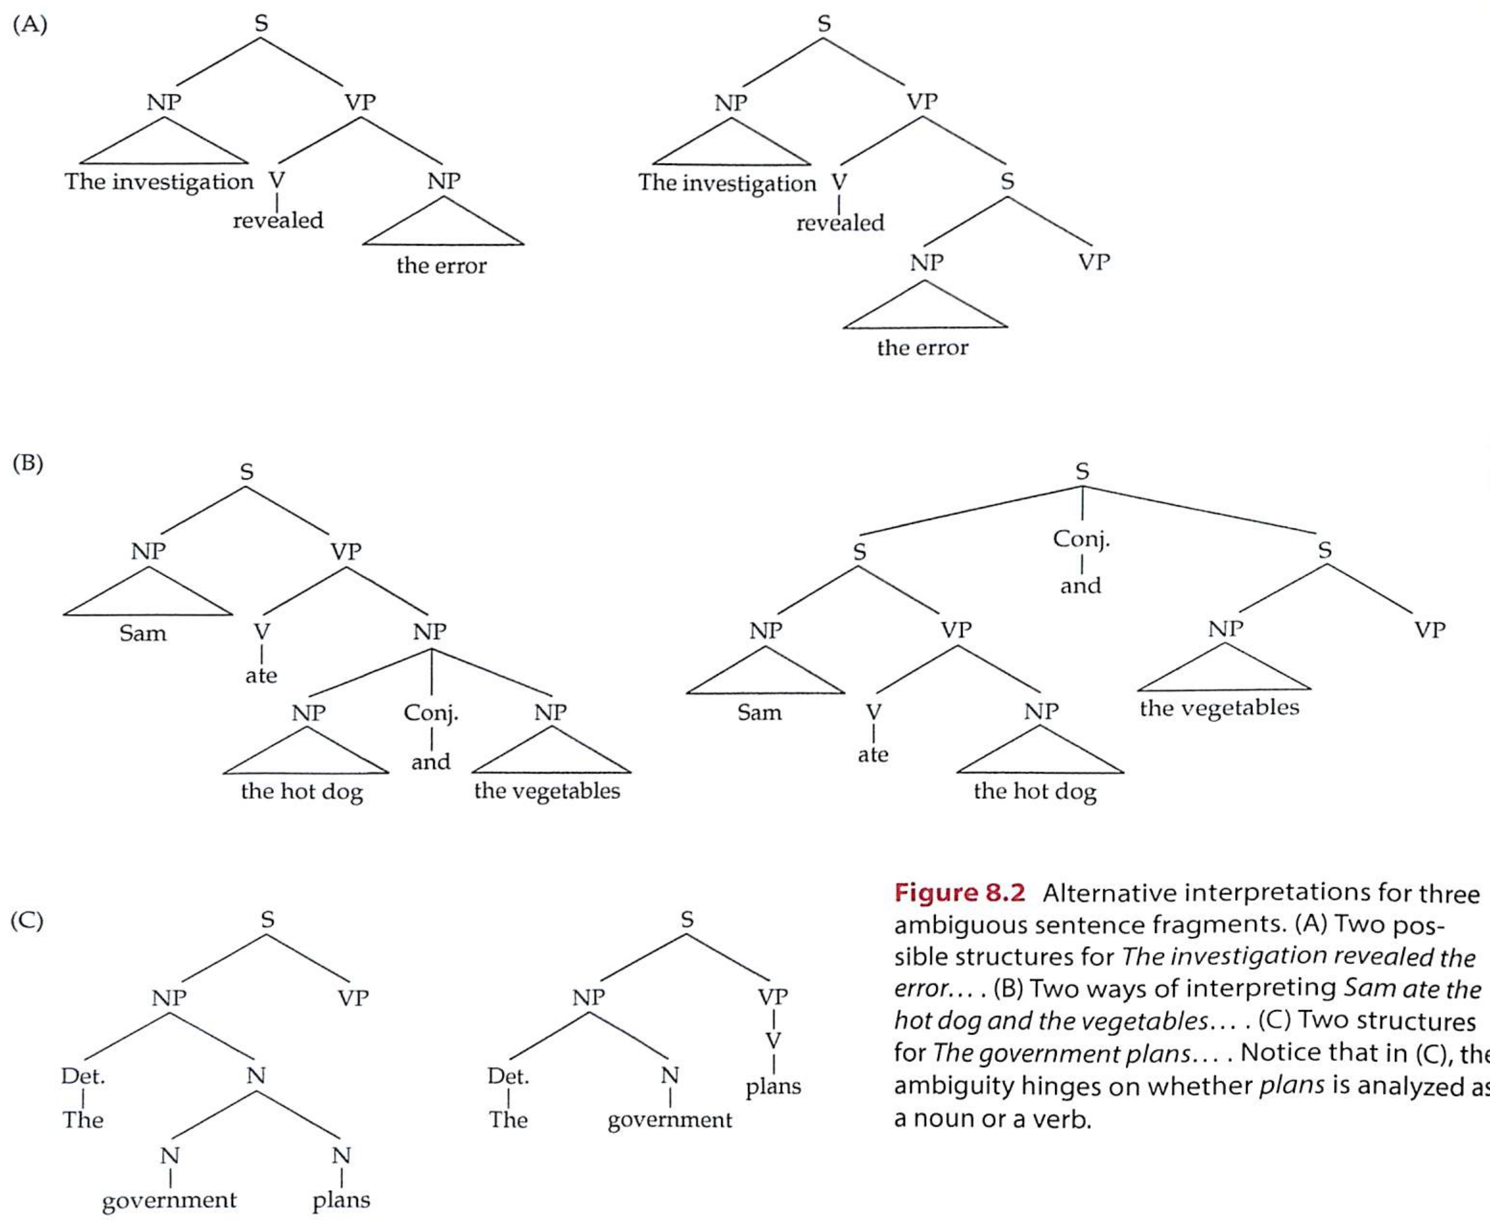

Figure 8.2 Alternative interpretations for three ambiguous sentence fragments. (A) Two possible structures for *The investigation revealed the error. . . .* (B) Two ways of interpreting *Sam ate the hot dog and the vegetables. . . .* (C) Two structures for *The government plans. . . .* Notice that in (C), the ambiguity hinges on whether *plans* is analyzed as a noun or a verb.

Measuring processing difficulty

Many garden path sentences are easy enough to identify, because the reader becomes aware of some processing tangle or even of an initial misreading. But psycholinguists often want to study subtler ambiguity effects that bypass conscious awareness. Once again, careful timekeeping is a useful tool for detecting mild but real processing disruptions.

A garden path effect is detected by measuring how long people take to read the disambiguating region of the sentence—that is, the first point at which any alternative but incorrect readings stop being consistent with the unfolding sentence. In the following examples, disambiguating regions are in bold:

> The general presented copies of the report **was aware** of the problems.
>
> The investigation revealed the error **resulted from** gross negligence.

LANGUAGE AT LARGE 8.1

Crash blossoms run amok in newspaper headlines

"VIOLINIST LINKED TO JAL CRASH BLOSSOMS." That's the headline of a 2009 article that ran in the newspaper *Japan Today*, prompting an American editor on an online discussion forum about copyediting to query, "What's a crash blossom?" In fact, the article was about the blooming musical career of a violinist whose father had died in a 1985 Japan Airlines crash. The headline is a garden path sentence rivaling the famous "horse raced" example in its inscrutability, and it involves a similar ambiguity of structure. It contains a reduced relative clause, and would be somewhat more understandable had it been written like this:

> The violinist who has been linked to the JAL crash blossoms.

There's also a second ambiguity contained in the word sequence *JAL crash blossoms*. Under the intended meaning, *crash* is supposed to be a noun, and *blossoms* is supposed to be a verb. But it's tempting to read the phrase such that *crash* is acting as a modifier (much like an adjective), and *blossoms* is a noun. Here, a lexical ambiguity involving the syntactic category of the word *blossoms* contributes to a completely nonsensical way of structuring the sentence overall. (Another example of this is with a sentence like *He saw her duck*, where *duck* could be either a noun or verb, resulting in very different readings.)

Function words like *who*, *that*, and *was* can serve the practical function of disambiguating a sentence, but in many cases such words are, perversely, optional. When they get dropped, this can lead to potential syntactic ambiguity. And in newspaper headlines, where there is a premium on terse expression, the disambiguating words are often swept away.

Hilarious results frequently ensue. Spotting garden path sentences in headlines (or crash blossoms, as they've come to be called at the suggestion of journalist John McIntyre) is a favorite sport of language geeks, and you can find many examples discussed on the popular linguistics blog Language Log (just do a search for "crash blossoms" on the blog's website). Here are a few examples of especially humorous crash blossoms that have been honored in the blogosphere:

> POLICE INVESTIGATE DEATH BY BALZAC (Balzac turns out to be a town, not the famous nineteenth-century French writer)
>
> CHINESE COOKING FAT HEADS FOR HOLLAND
>
> MANSELL GUILTY OF MISSING BUSINESSMAN'S MURDER
>
> SMOKING MORE DANGEROUS FOR WOMEN THAN MEN
>
> SISTERS REUNITED AFTER EIGHTEEN YEARS IN CHECKOUT COUNTER
>
> GREECE FEARS BATTER MARKETS AGAIN
>
> BRITISH LEFT WAFFLES ON FALKLAND ISLANDS
>
> MINERS REFUSE TO WORK AFTER DEATH
>
> JUVENILE COURT TO TRY SHOOTING DEFENDANT

You might try keeping a log of the crash blossoms you find in headlines over the period of, say, a week. Some might be funny, like the ones above in which a second meaning is readily accessible. Others might be nothing more than uninterpretable word mash-ups. Then try rewriting the headlines so that you preserve as much of their meaning as possible while removing the hazardous ambiguities.

For purposes of comparison, it's often possible to create virtually identical sentences that don't contain hazardous ambiguities, simply by introducing disambiguating function words, such as those in italics below:

> The general *who was* presented copies of the report **was aware** of the problems.
>
> The investigation revealed *that* the error **resulted from** gross negligence.

If the bold regions in the first pair of sentences (that is, the sentences containing the potential ambiguity) take significantly longer to read than the same bold

METHOD 8.1

Using reading times to detect misanalysis

How do you know for sure that a sentence is hard to read because readers are deceived into taking the wrong fork in the interpretive road? As a psycholinguist, you have several options. Most of the work on ambiguity resolution has relied on measures of reading time. A cheap and low-tech way of gathering these is by means of a **self-paced reading task**. Subjects read through sentences on a computer one word or phrase at a time, pressing a button to advance through the sentence, while a program records the amount of time they spend reading each segment. One downside to this method is that it's a bit unnatural, and breaks up the flow of reading by requiring subjects to keep pushing a button. Researches have tried to make the task a touch more natural by embedding it in a **moving window paradigm**, simulating somewhat the experience of bringing each word into focus by moving the eyes over the sentence. In this experimental situation, subjects first see the entire sentence with each character covered over by dashes:

```
--- ------- --------- ------ -- ---
------ --- ----- -- --- -------.
```

Each button press reveals a portion of the sentence and covers up the previous one, so a sequence of button presses would reveal chunks like this:

```
1. The general --------- ------ -- ---
   ------ --- ----- -- --- -------.

2. --- ------- presented copies -- ---
   ------ --- ----- -- --- -------.

3. --- ------- --------- ------ of the
   report --- ----- -- --- -------.

4. --- ------- --------- ------ -- ---
   ------ was aware -- --- -------.

5. --- ------- --------- ------ -- ---
   ------ --- ----- of the problem.
```

Even so, there may be a tendency for subjects to apply a steady rhythm to their button presses and artificially fit their reading into that time frame. When they hit an especially difficult phrase, they may have already pressed the button to advance to the next phrase, even if they haven't fully processed the earlier phrase. This can lead to *spillover effects*, where the reading time that should be recorded for a particular phrase ends up being recorded for the following one instead, as subjects struggle to catch up in processing. This is a bad outcome, since so much of the theorizing about garden path sentences hinges on firm knowledge about exactly which regions of the sentence cause hiccups in processing. In general, there's a concern that by breaking up the flow of reading, and in a way that doesn't allow subjects to backtrack, you're not getting the most accurate picture of how processing works in typical reading situations.

You can encourage a smoother flow of reading and also get a more detailed snapshot of reading with a slightly higher-tech version of the reading time measure. Subjects read sentences presented on the screen in their entirety while their eye movements are being recorded. This allows you, the researcher, to measure the amount of time readers

self-paced reading task A behavioral task intended to measure processing difficulty at various points in a sentence. Subjects read through sentences on a computer screen, one word or phrase at a time, pressing a button to advance through the sentence. A program records the amount of time each subject spends reading each segment.

moving window paradigm A version of the self-paced reading task in which dashes initially replace each alphabetic character in a sentence, and participants press a button to successively "uncover" each portion of the sentence. This method of presentation simulates a fairly natural reading rhythm.

regions in the second pair (the unambiguous versions), then a garden path effect is considered to have occurred, and is taken as evidence that the reader ran into some processing trouble from having initially been misled, at least in part, by the wrong interpretation of the ambiguous portion of the sentence. You can get a more detailed feel for the reading time task in **Method 8.1**.

METHOD 8.1 *(continued)*

spend on each region of interest. It also allows you to get a picture of how long subjects take reading each portion on their first pass, how often they backtrack and where, and how long they spend rereading portions of the sentence.

By and large, there's a lot of overlap in the data gathered by these two methods—certainly enough that the self-paced reading task remains a viable, easy way to detect most robust garden path effects. But the finer detail of eye tracking can provide clues that might help distinguish between competing theories. For instance, as you'll see in Section 8.2, different theories make different predictions about the nature of the parser's initial analysis, as opposed to later revisions to that analysis. Some researchers have argued that first-pass reading times align neatly with the parser's first guesses, while backtracking, or regressive eye movements, are informative about later revision processes. It's doubtful that eye movement behavior lines up this cleanly with the inner workings of the sentence processing system, but the distinction between earlier versus later reading time can often add an extra dimension to the reading time profile.

Regardless of how they're collected, interpreting reading times is tricky, because the measure is somewhat of a blunt instrument. All kinds of factors can cause increases in reading times, not just misanalyzing the syntactic structure of the sentence. Reading times can be elevated because uncommon words are used, because the sentence depicts an unexpected event, because a low-frequency structure is used (regardless of ambiguity), or because the sentence is somehow unnatural in a particular context. Ideally, in order to be able to attribute long reading times to misanalysis, you want to—wherever possible—compare sentences that contain an ambiguity to *unambiguous* versions of the same sentences, with the same meaning, containing the same content words. This is a standard design feature of experiments that rely on reading times. Usually, each subject sees only one version of each sentence, so the second version cannot be predicted, and reading times across subjects are compared.

8.2 Models of Ambiguity Resolution

What's the big deal about garden path effects?

Some of the most passionate discussions in psycholinguistics (often referred to as the "parsing wars") have revolved around questions of ambiguity resolution and how to correctly explain garden path effects. Part of me would like to say that this passion was fueled by a burning desire to equip writers and editors with a definitive list of do's and don'ts of sentence construction—no doubt the skills of an entire generation of professional writers might improve as a result. But the reality is that researchers were mobilized into spending long hours (and indeed, many years) in studying ambiguity resolution because this area of psycholinguistics became the battleground for several "Big Ideas" in the field.

In order to work up to the Big Ideas, we have to start with a simple question: Why is it that *one* reading of an ambiguous string of words often seems to be so much more attractive than another? For instance, in the classic horse-racing example, why are we so tempted to read the sentence as a simple main clause, and why does it seem to be so hard to access the reduced relative clause reading? Notice that this kind of confusion is strikingly different from what happens in word recognition. Remember that with words, multiple candidates are usually activated in parallel, with competitors dropping off in activation over time until a winner remains standing. We don't ever seem to have the experience of hearing the syllable *can-* and committing *so* strongly to the word *candle* that we just can't seem to recover and recognize the word as *cantaloupe* once it's been fully uttered. But with full-blown garden path sentences, there's a strong

garden path theory A theory of parsing that claims that an initial "first-pass" structure is built during comprehension using a restricted amount of grammatical information and guided by certain parsing principles or tendencies, such as the tendency to build the simplest structure possible. Evaluations of plausible meanings or consideration of the context only come into play at a later stage of parsing.

early commitment to one particular structure, a strategy that sometimes leads to epic failure when this commitment turns out to be wrong. When faced with a syntactic ambiguity, why don't we just remain non-committal until there's enough information to help us decide on the right structure? This would avoid situations in which the interpretation eventually falls apart. What is it about the preferred structure that makes it so magnetic?

The garden path theory

Explanations for the existence of severe garden path effects have followed several very different approaches. One of the earliest accounts of these troublesome sentences is known as the **garden path theory** of parsing, proposed by Lyn Frazier and her colleagues (see Frazier & Fodor, 1978, for an earlier version; and Frazier & Clifton, 1996, for a reformulation). These researchers noticed that for a number of different types of garden path sentences, people seemed to be drawn toward a structure that was simpler than what turned out to be the correct structure. For example, if you look back in Figure 8.1 at the rendition of the sentence fragment *The horse raced past the barn*, it's evident that the more alluring structure—with its single clause and basic word order—is not as complex as the reduced relative clause, which involves a passive structure and two clauses. Frazier and her colleagues argued that a number of other garden path effects follow the same pattern (for example, look back at the sentence fragments diagrammed in Figure 8.2A). They proposed that when faced with an ambiguity between a simpler and a more complex structure, people have a strong preference for the simpler structure, and this causes an interpretation crash when the sentence reveals itself to be consistent only with the more complex structure.

Elaborating on this observation, Frazier and colleagues suggested that, rather than activating multiple meanings at once, as is the case in word recognition, the parser computes only one structure and its associated meaning. To justify why there are such different mechanisms for recognizing words and processing sentences, they argued that the two processes are really very different in nature. After all, one of them involves pulling pre-stored items out of memory while the other involves actual *computing*; new structures have to be built out of the stream of words coming in as input, not just matched up with items in memory. As you saw in Chapter 6, words can't just be assembled in any old way for a given language, but have to follow certain rules of syntax that specify the possible ways to group words into constituents—we might think of these syntactic rules as being like a set of very stringent building codes. So the parser has to consult these building codes and propose a way to group words into a sensible structure that meets their legal requirements. It's easy to imagine that all this computation might be fairly expensive in terms of processing resources. In order to be able to interpret sentences incrementally—as we've seen that people manage to do—the parser needs to build sentences very quickly. So to achieve this blazing speed, the parsing system builds whatever legal structure is easiest to build, and runs with it. If the whole thing runs aground at some later point, then a reanalysis of the sentence is initiated.

It may have crossed your mind, though, that there are other factors that could potentially affect how an ambiguous phrase is interpreted. For example, some interpretations might be more plausible than others, or fit in better with the preceding context. Wouldn't this information be taken into consideration while assembling the structure of a sentence? According to the garden path theory, no, not during the parser's first attempt at building structure. This may seem odd, since information about plausibility or contextual fit could be a great

help in averting a parsing meltdown. Why would the parser ignore such useful information? This is where the Big Ideas come in.

On the face of it, a parsing system that disregards helpful information looks like a badly designed system, because it would result in a number of otherwise avoidable errors—it just seems needlessly dumb. But it's worth pointing out some intellectual historical context. At the time that psycholinguists began trying to explain how ambiguity resolution works, there was quite a bit of emphasis in cognitive psychology on the fact that sometimes, human information processing *does* appear to be a bit dumb. Evidently, we humans are saddled with certain inherent processing limitations. Intelligent, thoughtful analysis consumes a lot in the way of processing resources, and it takes a lot of time. In the 1970s and 1980s, it was becoming increasingly apparent to psychologists that because of our limitations, we can easily deplete the processing resources that are needed for solving some difficult cognitive problems. As a work-around, we often rely on cognitive processes that are fast and easy, but that fail to take into consideration all of the relevant information.

This notion had come to be a highly influential one. The work of well-known psychologists such as Amos Tversky and Daniel Kahneman (1974) had shown that in many situations, people rely on very quick but error-prone **heuristics**—that is, shallow but very fast information-processing shortcuts that often lead people to leap to "illogical" or incorrect conclusions based on very superficial cues. (As an example, consider this problem: The surface area of water lilies growing on a pond doubles every day. If it takes 24 days for the pond to be completely covered, in how many days is the pond half covered? If you answer 12, then you've fallen prey to a common heuristic.) Cognitive psychologists began talking about an important distinction between these fast, automatic, but often buggy processes, as opposed to slow, deliberate, but ultimately more accurate thought machinery. The scientific thinking evolved that an enormous amount of human cognition depends on the faster but dumber cognitive processes, and that we have a limited mental budget to spend on the more "intelligent" ones, which require considerably more effort and processing time. Such distinctions continue to be important ones in psychology to this day—you can find an accessible overview of this body of research in Daniel Kahneman's (2011) book *Thinking, Fast and Slow*.

heuristics Shallow but very fast information-processing shortcuts that often lead to incorrect conclusions based on superficial cues.

The garden path theory is very consistent with this line of thinking, and represents one of several attempts to apply the ideas about fast versus slow cognition to problems in language processing. The idea is that the parsing system is able to conserve processing resources by relying on a set of quick and simple structure-building heuristics in its first guess at sentence structure. In this first pass, a great deal of information that *could* be pertinent to resolving the ambiguity is ignored, because this would take too much time and mental energy to integrate—in fact, if the parser did have to consider all potentially relevant information in the first-draft structure, doing so might strain memory capacity beyond the breaking point. All things considered, a fast-and-cheap solution may be the best strategy, even if it does result in frequent errors—if an error is detected, a deeper analysis of the sentence is triggered.

If you've worked your way through the Digging Deeper discussion in Chapter 7, this theoretical approach should seem very familiar to you. There, I discussed the tension between *modular* and *interactive* models of word recognition. In a modular system, "higher-level" information doesn't have a direct effect on "lower-level" processes. Instead, it comes into play at a later stage, once the lower-level processes have been completed. For example, in a modular model of word recognition, information about the context or lexical knowledge can't affect the activation levels of individual sounds; it can only help decide which sounds are most appropriate, based on how activated they are on the basis of the bottom-up input. In an interactive model, on the other hand, information from the higher levels can flow from the top down to affect the activation of sounds.

When it comes to parsing, the garden path model is an instantiation of a modular system. Structures are initially built by a lean, fast parser on the basis of a very limited amount of purely syntactic information, and entirely without the benefit of semantic knowledge or knowledge about the context of the sentence. Once a partial structure is built, it gets sent off to the higher-level semantic interpretation component, which then evaluates its sensibleness. But no information about the *meaning* of the sentence fragment to that point can influence the parser's first stab at structure-building. For some researchers, the division between syntactic and semantic knowledge could be seen as a division between fast, automatic mental processes and slow, thoughtful ones, with semantic information lagging behind syntactic processes (and sometimes cleaning up some messes in interpretation). This distinction was discussed in quite a lot of detail by the philosopher Jerry Fodor in his well-known book *The Modularity of Mind* (1983).

In Chapter 7, I suggested a metaphor for thinking about the distinction between modular and interactive systems: A highly modular system is structured like a company or factory in which lower-level employees do their jobs with very limited knowledge and responsibilities, and then simply pass on their work to higher-level workers who later make all the important decisions. On the other hand, an interactive system is more like a company in which all levels of workers have a shared stake in making decisions, with information flowing freely from upper to lower levels. As in business, the competing models offer different advantages: the first option may be fast and cheap, but lacks a certain flexibility (which you may have experienced, if you've ever asked a customer service representative to deal with a problem outside of his area of expertise). The second option is "smarter" in some ways, but maybe not so cost-effective on a large scale.

The constraint-based approach

In the garden path model, because the parser initially spits out a single preferred structure in the face of ambiguity, it's expected that people will often run into an interpretive dead end, and be forced to reconsider the sentence's meaning; this is an unavoidable by-product of the parser's efficiency. But quite a few researchers have argued that the parser actually shows a lot more subtlety and intelligence in its initial analysis than the garden path model is willing to give it credit for. The core of the argument is that in reality people experience much less confusion in processing sentences than you'd expect if they were blindly computing the "preferred" structure without considering other potentially helpful information.

According to the early version of the garden path account, the parser only looks at the syntactic categories of words, and how they might be assembled according to the language's syntactic building codes, all while following certain preferred building routines. But in that case, *all* sentences that make use of *dispreferred* structures (where more favored, or preferred, ones are available for the same string of words) should be *equally* likely to lead to a garden path effect because the wrong structure is initially built. But this doesn't seem right. For example, compare these two sentences:

The dog walked to the park wagged its tail happily.

The treasure buried in the sand was never found.

Both sentences involve reduced relative clauses. Most people find the first far more difficult to interpret than the second (if your own parser has become benumbed to the distinction as a result of reading too many garden path sentences in the last half hour, try them out on a fresh reader). This became one of the key arguments against the garden path theory, as dissenting researchers began to look for evidence that similar structures resulted in extremely variable garden path effects. One of the key pieces of evidence that such dissenters hoped to find was data showing that some sentences that are saddled with dispreferred structures, such as reduced relative clauses, don't show any measurable garden path effects at all—for example, if we compared the following two sentences, we might not find any difference in reading times for the region marked in bold:

The treasure buried in the sand **was lost forever.**

The treasure that was buried in the sand **was lost forever.**

This would suggest that in the first sentence of the pair, which begins with a potentially ambiguous string of words, readers don't seem to have been misled by the other (supposedly preferred) interpretation, and interpretation goes as smoothly in reading the final portion of the text as if there had been no ambiguity at all.

In one influential paper, researchers led by Maryellen MacDonald (1994) used a witty expository device to make the point that reduced relative clauses don't always cause interpretive snarls. They pulled out many examples of reduced relative clauses from the writings of various psycholinguists, some of whose theories in fact *did* predict that such structures would invariably cause garden path effects. Some of these examples are reproduced in **Box 8.2**, along with a few other examples I dug up myself from various sources.

The main competitor to the garden path theory was the **constraint-based approach**. It parted ways with the garden path account on a number of points. First, constraint-based theorists argued that a broad range of information sources (or constraints) could *simultaneously* affect the parser's early decisions, including semantic or contextual information, which was often thought of as being too subtle or slow to drive initial parsing (specific constraints will be discussed in detail in Section 8.3). Parsing speed, they argued, is not necessarily bought at the expense of subtlety or intelligence in processing. As a result, only a very small percentage of temporary ambiguities might cause any discernible processing difficulties.

Second, they argued that syntactic ambiguity resolution actually looks very much like ambiguity resolution in word recognition. In both cases, multiple interpretations are accessed in parallel, and our minds very quickly try to judge which one is more likely to be correct, based on very partial linguistic evidence along with various sources of information—such as context, or statistical frequencies—that make one interpretation seem more likely to be the right one. These sources of information (or constraints) can serve to either ramp up or suppress the activation of one interpretation relative to the other.

So why does it often seem that we consider only *one* interpretation of a sentence, with the other seeming to be completely inaccessible? The response of constraint-based theorists was that flagrant garden path effects are in fact the results of an unfortunate coincidence: the various sources of information that are available in the ambiguous portion of the sentence just happen to *overwhelmingly* point to the wrong interpretation, causing the parser to suppress the competing correct one, much as a lexical competitor is suppressed over time. By the time readers or hearers get to the disambiguating region, it's too late—they've already pushed the eventually correct alternative beyond accessibility.

constraint-based approach The main competitor to the garden path theory, this approach claims that multiple interpretations of an ambiguous structure are simultaneously evaluated against a broad range of information sources (or constraints) that can affect the parser's early decisions.

BOX 8.2

Not all reduced relatives lead to processing implosions

We've seen how a relative clause like *the horse [that was raced past the barn]* can become ambiguous and potentially impede sentence processing when the clause is "reduced" by removing function words (becoming *the horse raced past the barn*). But are such reduced relatives always a problem for the person processing them?

Maryellen MacDonald and her colleagues (1994) discovered examples of reduced relative clauses in the writings of various psycholinguists, many of whom were proponents of theories predicting that reduced relatives invariably lead to garden path effects. MacDonald's point was not that these psycholinguists are bad writers, oblivious to the fact that their sentences might cause readers to stumble, but (more damningly) that their *theories* are flawed, since these sentences seem to pose no real burden for readers. Here's a sampling of the sentences unearthed by MacDonald et al., with the reduced relative clauses in color type (color was added for this textbook, not by the authors of the original papers):

> Thus, a noun phrase followed by a morphologically ambiguous verb (e.g., "The defendant examined") will be temporarily ambiguous.
>
> Trueswell, Tanenhaus & Garnsey, 1994, p. 287

> Referential information provided by the discourse is of no help.
>
> Britt, Perfetti, Garrod & Rayner, 1992, p. 305

> Recent research reported by Ferreira and Clifton (1986) has demonstrated that syntactic processing is quite independent and that the initial syntactic analysis assigned to a sentence is little influenced by the semantic information already analyzed.
>
> Frazier & Rayner, 1987, pp. 520–521

> In all cases, the examples cited here were not the only reduced relatives in these articles.
>
> MacDonald et al., 1994, p. 678

Of course, the use of reduced relatives is not limited to academic writers. Here are some examples I discovered among the wilds of the popular print media:

> The reward money offered to find Jyrine, a helpless toddler, was a paltry $1,000.
>
> "What happened to my child?" *Essence*, September 2007, p. 224

> The food prepared for the now non-existent media event is donated to homeless shelters.
>
> "911 Coverage of Sept. 11 attacks overshadowed a world of events," *Denver Post*, October 11, 2001

Of the two reduced relatives in the following example, only the first one (italicized) is potentially ambiguous. Neither seems to pose any difficulty for processing:

> ... many Western Europeans believe that a job *offered to an older worker* is a slot taken away from a younger one.
>
> *Foreign Affairs*, May/June 2007, p. 55

This final, lovely example has one reduced relative clause nested inside another (the embedded clause is italicized):

> ... many Russian politicians influenced by more than 40 years of communist propaganda *aimed against the west* finally became good "Homo Sovieticus."
>
> Letter to the editor, *San Francisco Chronicle*, June 6, 1995

In other words, the idea is that a sentence like *The horse raced past the barn fell* starts out activating both interpretations, but since all the available constraints in the ambiguous region are heavily biased toward the simple main clause, the reduced relative reading is too weak to ever become a serious candidate and becomes deactivated. On the other hand, things are a bit different in an easier sentence like *The treasure buried in the sand was never found*. In this case, the constraints are more equivocal, creating a more even balance between the two alternatives. When the disambiguating region is reached, the new information in that region provides the decisive evidence between the alternative meanings (both of which are still active), and readers end up settling on the correct meaning without too much difficulty.

This is really a dramatically different view of the nature of ambiguity effects than the one offered by the garden path model. For the garden path account, comprehension mishaps arise because the inherent limitations on human cognition force the parser to make decisions about structure without considering all the evidence, so these decisions will often be wrong. For the constraint-based model, the parser *has* considered all the evidence and has made a rational prediction based on prior experience with language—it's just that in some cases, the outcome turns out to fly in the face of all reasonable predictions.

Let's turn to testing the predictions made by each model. As I've mentioned, the garden path account predicts that the parser will crash whenever it's forced to build a structure that falls outside of the preferred structure-building routines—it will always build the wrong one first. But the constraint-based account says that garden path effects will be highly variable, even for the so-called dispreferred structures. Processing difficulty may range from devastating to non-existent, depending on the specific words that appear in those structures, or even the specific context of the sentence. This is because the extent to which one reading over another will be favored should depend on the biasing strength of *all* the constraining sources of information together, of which general structural information is but one source. For the garden path theory, only structure should determine the presence of a garden path effect, since this is the only source of information that's considered in the early phases of parsing. Garden path effects, then, should be impervious to other sources of information.

To evaluate the predictions of the garden path and constraint-based theories more concretely, let's look more closely at the varied bouquet of information sources that bear on the resolution of syntactic ambiguity.

WEB ACTIVITY 8.2

When do ambiguities pose processing problems? In this activity, you'll see a number of examples of temporary ambiguities. You'll informally gauge the difficulty of each example—does the degree of processing difficulty seem variable, even for structurally similar sentences? You'll be prompted for potential explanations for your intuitions.

8.3 Variables That Predict the Difficulty of Ambiguous Sentences

Thematic relations associated with verbs

Part of our inherent knowledge of the meanings of verbs involves knowing what kinds of events they describe, including how many and what kinds of participants must be involved, and what roles the various participants play. This is often referred to as knowledge of **thematic relations**. For example, when we hear the verb *bite*, we know that a biting event involves at least two participants—the biter and the bitee. We also know that the biter must be an animate entity in possession of teeth and the ability to close those teeth around something. The bitee, on the other hand, could be another animate entity, or an inanimate object; you can bite your brother, but you can also bite your brother's finger or a piece of toast. A parser that had access to this semantic information (as predicted by the constraint-based account) could make much smarter guesses about the likely structure of ambiguous word strings in at least some cases.

Consider a sentence like *The treasure buried in the sand was lost forever*. The verb *bury* involves at least two participants, one doing the burying, and one being buried. But treasure is a terrible fit as the participant doing the burying. So, as soon as the reader gets to *The treasure buried...*, there's semantic pressure to shift to the alternative reduced relative clause reading, in which the treasure is the entity *being* buried rather than the one doing the burying. This means the reader is less likely to be led down the garden path to an interpretation that later blows up.

thematic relations Knowledge about verbs that captures information about the events they describe, including how many and what kinds of participants are involved in the events, and the roles the various participants play.

But thematic relations aren't quite as helpful in a sentence like *The dog walked to the park wagged its tail happily.* Ultimately, the right interpretation of the sentence refers to an event in which two participants are involved—the dog and the implicit, unmentioned dog walker. Since a passive clause is involved, the dog in this event is the walkee rather than the walker (see Box 8.1 for a primer on passive structures). But when the verb *walk* appears in an active main clause, it also allows walking events in which there is just one participant doing the walking—this participant must have functioning legs and be capable of locomotion. As it happens, the noun *dog* fits perfectly nicely with this requirement, leading to a very plausible interpretation of the ambiguous region of the sentence as a simple main clause in the active voice. Unfortunately, this interpretation turns out to be wrong.

The constraint-based model predicts that when information from thematic relations is strong enough to steer people away from the normally dominant reading of a structure, it should diminish or eliminate a garden path effect, as measured using reading times. In fact, this is what a number of studies have found (e.g., Trueswell et al., 1994).

The syntactic frames of verbs

In Chapter 6, we spent a fair bit of time talking about how individual verbs in English select for specific syntactic frames, and how children have to learn these facts in addition to the general phrase structure rules of their language. You can't just put any verb into a generic verb slot in a sentence structure and expect things to turn out well. Hence, the ungrammaticality of:

*The soldier buried.

*Alice fell the ball.

*Samantha sneezed that Billy was in prison.

*The mom put the cookies.

*The mom put the cookies about the jar.

*Frank said the report.

Some verbs, like *fall,* have to be **intransitive verbs,** with just a subject and no direct object. Some (for example, *buried*) are **transitive verbs,** with a subject and a direct object (a direct object is a noun phrase that denotes one of the participants in the event described by the verb and appears immediately after the verb with no intervening preposition). Some (for example, *put*) are **ditransitive verbs,** calling for both a direct object and an indirect object (which may be introduced by a preposition). Finally, some (like *said*) are **sentential complement verbs,** which introduce a clause rather than a direct object noun phrase. (See **Table 8.1** for some additional examples.)

The garden path model (at least in its original incarnation) claims that all the parser cares about in its first pass is a verb's status as a verb, leaving aside these fancy details about syntactic frames. But constraint-based advocates have countered that the parser has access very early on to information about the syntactic frames that are linked with specific verbs. Access to such information would lead to much "smarter" parsing decisions. Notice that *walk* is allowed to appear with or without a direct object, whereas *bury* absolutely can't get by without one:

Raj walked his dog to the park.

The dog walked.

The dog walked to the park.

intransitive verbs Verbs that occur with a subject but no direct object.

transitive verbs Verbs that take both a subject and a direct object.

ditransitive verbs Verbs that occur with a direct object and an indirect object (which may be introduced by a preposition).

sentential complement verbs Verbs that introduce a clause rather than a direct object noun phrase (NP).

TABLE 8.1 A variety of syntactic frames for verbs

Intransitive verbs *(appear with a subject only)*
The neighbors' dog reeks.
Gerhardt finally relaxed.
Mariah sings beautifully.
The magician vanishes.
The tornado touched down last night.
Transitive verbs *(select for a direct object noun phrase)*
This theory vexes me.
I rarely wear socks.
The author signed his book.
The queen poisoned the courtesan.
The engineer inspected the plans.
Ditransitive verbs *(select for two noun phrases, one of which may appear inside a prepositional phrase)*
Please lend me the money.
Parents should send their teens to boarding school.
Devon presented his fiancée with a ring.
Frances hid the affair from her husband.
The teacher tossed Alisha the ball.
Sentential complement verbs *(select for a clausal unit)*
The workers complained that their bosses harassed them.
I agree that the king should be deposed.
The president assumed his staff would cover his mistake.
The contract stipulated that employees should be paid weekly.
The teacher regretted that she had promised the kids a dollar for every book they read.
Verbs that fall into more than one category[a]
NP-bias verbs: e.g., accept, repeat, advocate, maintain, reveal, print
S-bias verbs: e.g., conclude, decide, promise, worry, prove

[a]Many verbs fall into more than one of the four categories, but may occur much more often in one syntactic frame than another. The verbs given as examples here can appear with either a sentential complement (S) or a direct object noun phrase (NP), but have a bias for one or the other. Try constructing sentences with both kinds of frames for each verb.

The soldier buried the land mine.

*The pirate buried.

*The pirate buried in the sand.

Based on this knowledge, on getting *The treasure buried in the...,* the parser would be able to avoid the normally seductive main clause reading, since this reading lacks the obligatory direct object that *bury* demands. This would leave as viable the reduced relative clause interpretation, in which the syntactic process of passive formation has taken the direct object *the treasure* and stuck it into the subject position of the correct structure.

There's a fair bit of experimental evidence showing that knowledge of syntactic frames plays a central role in ambiguity resolution. What's more, this knowledge about syntactic frames is fairly nuanced, not just limited to what's grammatical or not, but also tuned in to which syntactic frame is most *statistically likely* for a given verb. An especially ingenious case study showing effects of syntactic frame probabilities is presented in Box 8.3.

Frequency-based information

So far, we've been talking about syntactic versus lexical ambiguity as if it were easy to sort ambiguous examples into one bin or the other. And sometimes,

BOX 8.3
Subliminal priming of a verb's syntactic frame

If you work at it, you can develop a reasonable set of intuitions about how sentence processing works by noticing which sentences are difficult, why, and how to make them less so. Nonetheless, there's a limit to how much conscious access we're able to have to the internal workings of our parsers—which is one reason why psycholinguists need to design clever experiments in order for us to truly understand the complex nature of building sentences on the fly.

In fact, one such clever study by John Trueswell and Al Kim (1998) illustrates just how much of our sentence processing can be hidden from consciousness, while at the same time exploring how it all works. The researchers took as their starting point evidence from other studies showing that specific verbs could be biased toward certain syntactic frames, and that these biases in turn could influence the interpretation of garden path sentences. For example, the verbs *accept* and *realize* both allow either a direct object NP or a sentential complement, but *accept* is biased toward the direct object frame while *realize* is biased toward the sentential complement frame. Hence, the two verbs set up very different expectations for how the following ambiguous fragments should continue:

> The photographer accepted the ...
>
> The photographer realized the ...

So if the first fragment perversely continues with a sentential complement as below, contrary to the bias of the verb *accept*, this results in a garden path effect in the disambiguating region (in color):

> The photographer *accepted* the fire could not be put out.

However the same region is read more quickly if it follows the verb *realized*:

> The photographer *realized* the fire could not be put out.

Taking these facts as background, Trueswell and Kim noted that information about syntactic frames is part of the lexical information associated with a word, much as the meaning of a word has to be linked to that specific lexical entry. Given that words can prime related meanings (so seeing *tiger* boosts the activation of the related word *lion*), they wondered whether it's also possible for verbs to boost the accessibility of syntactic frames associated with other words. That is, can seeing *realize* boost the accessibility of the sentential complement frame for *accept*, with the result that readers will be less likely to experience confusion in the problematic example shown above?

To find out, Trueswell and Kim had subjects perform a self-paced reading task on sentences containing potential ambiguity, as in *The photographer accepted the fire could not be put out*. But just as subjects pressed a button right after reading *photographer*, and before they saw the verb *accepted*, a "priming" verb flashed subliminally on the screen for 39 milliseconds—subjects consciously experienced this merely as a brief flicker that happened just before *accepted* was displayed. Half of the time, the priming verb favored a direct object NP frame (for example, *obtained*), while the other half of the time, the priming verb was biased toward a sentential complement frame (*realized*). It turned out that the identity of the secret priming verb shifted the expectations for the syntactic frame, even though the overtly visible verb (*accepted*) stayed the same in both conditions: Subjects spent less time reading the disambiguating region when they'd "seen" *realized* than when they'd "seen" *obtained*. It would appear that the syntactic frame bias associated with one verb can become activated and spread to another verb, in turn influencing the parser's expectations about structure—and that all this can happen without the reader becoming aware of it.

structurally ambiguous word strings really do come down purely to syntax; for example, in a classic attachment ambiguity, such as *I saw an elephant in my pajamas*, it's really just a matter of how the words get grouped together—the meanings and syntactic categories of the words themselves are identical under both readings. But that's not true for all ambiguities, as you might have noticed from Figure 8.2. For instance, where do you sort a sentence like *The government plans to raise taxes were defeated*? Here, the syntactic ambiguity hinges entirely on whether *plans* is taken to be a noun or a verb, and how this *word* is interpreted ultimately drives how the sentence can be structured.

Actually, it's also possible to think of the famous reduced relative clause ambiguity as similarly boiling down to a lexical ambiguity, though a somewhat more subtle one. For a sentence like *The horse raced past the barn fell*, it comes down to how we interpret the word *raced*. Now, with the way I've drawn the tree structures in Figure 8.1, *raced* simply has the syntactic category of verb whether it appears in the main clause or the reduced relative clause reading. But a closer look shows that these *two* verbs are not interchangeable, and are composed of different parts. In the main clause reading, *raced* is a garden-variety example of a verb stem plus a past-tense morpheme. But in the reduced relative clause reading, the verb stem is joined by a past-participle morpheme, not a past-tense morpheme. Unfortunately, for regular verbs of English, these two morphemes look exactly the same—they are homophones, just like the words *bank* and *bank*. Note, for example, that the same form shows up in all of these syntactic environments: *I raced, I have raced, The car was raced*. But you can infer that past-tense verbs are different from past participles by looking at many *irregular* verbs, in which the two types show up in different forms, and appear in different syntactic environments, at least in standard varieties of English (for example: *I drove* but not **I have drove* or **The car was drove by my brother; I have driven*, or *The car was driven*, but not **I driven home*). As a result, when certain *irregular* verbs appear in reduced relative clauses (see the examples below), the italicized portions of these sentences contain no ambiguity whatsoever, even though they have exactly the same syntactic structure as some of the fiercely confusing garden path sentences we've seen:

> *The horse rode past the barn* and fell.
> *The horse ridden past the barn* fell.
>
> *The students gave high marks* to their instructor.
> *The students given high marks* liked their instructor.
>
> *Wheat grew in the Midwest* and was exported in large quantities.
> *Wheat grown in the Midwest* was exported in large quantities.

The idea that many syntactic ambiguities have their roots in lexical ambiguity is a provocative one, because it suggests that perhaps syntactic and lexical processing are not so distinct after all, as suggested by a number of constraint-based theorists. At the very least, factors that influence the interpretation of ambiguous words should have a visible impact on the interpretation of some syntactic ambiguities.

This is the case made by Maryellen MacDonald and her colleagues (1994). They argued that when we look at word-based ambiguity (say, the two meanings of *bank*, or the noun-versus-verb readings of *plan*), we see that the more frequent readings become activated most strongly. This is the pattern that shows up in standard word recognition tasks. If that's so, they argued, shouldn't the relative frequencies of lexical alternatives also play a role in situations where a syntactic ambiguity hinges on a lexical ambiguity? If *plans* is more frequent as a verb than as a noun, shouldn't this make it easier to build the structure that's consistent with the verb reading, since the verb representation will be the more strongly activated one? The same argument can be made for reduced relative clauses that contain past participles. As it happens, some verbs (such as *entertained*) rarely show up as past participles, whereas others (such as *accused*) more commonly do. Thus, we might expect *entertained* to lead to a much stronger bias for the main clause interpretation than *accused*, resulting in a bulkier garden path effect for the first than for the second of these two sentences:

The audience entertained at the gala left in high spirits.

The suspect accused at the crime scene was soon released.

Several researchers (e.g., Trueswell, 1996; MacDonald et al., 1994) have indeed found that the severity of garden path effects for reduced relative clauses is affected by the frequency of past-tense/past-participle readings of ambiguous verb forms. So the outcomes of word recognition processes have a way of leaking into syntactic ambiguity resolution.

We can push the idea of frequency bias a bit further. It seems natural that, when encountering an ambiguous word that could have either Meaning A or Meaning B, the more frequent meaning would be easier to pull out of memory, because it will have left a stronger trace in memory. But even though sentences are *built* rather than pulled from memory, it still stands to reason that more common structures would be easier to build than less common ones, just as a building crew would be more efficient in building a house from a plan that they've used many times before than from one that they've used less frequently. Could some portion of garden path effects be due to different frequencies of use of the alternative structures? For example, are reduced relative sentences so hard in part because their structure is much less common in English than the competing structure of simple past-tense main clause?

Models that factor in how frequently competing structures are used do often seem to give better results than ones that leave this source of information out. To ask this question more directly, though, we can compare similar structures across languages that happen to have different patterns of frequency for these structures. Take the following sentence:

Someone shot the maid of the actress who was standing on the balcony with her husband.

Who was standing on the balcony with her husband? Most English speakers say it was the actress, preferring to attach the relative clause to the nearby noun. But a funny thing happens when the sentence is translated almost word for word into Spanish:

Alguien disparó contra la criada de la actriz que estaba en el balcón con su marido.

The majority of Spanish speakers say that it was the maid who was on the balcony with her husband. This is interesting because, as it happens, relative clauses in Spanish attach more frequently to the first of two nouns, while the reverse is true in English. English and Spanish readers, then, seem to be resolving the ambiguity in a way that reflects their personal experience with these structures (Mitchell & Cuetos, 1991).

The importance of context

Let's return to that famous racing horse, and the ugly sentence in which it appears. What happens when the sentence shows up in a story like this:

Farmer Bill and Farmer John were racing their horses through the field. Farmer Bill rode his horse along the fence, while Farmer John raced his horse past the barn. Suddenly, the horse raced past the barn fell.

Does the last sentence of the story suddenly seem less difficult? If so, why would that be?

Several researchers (e.g., Crain & Steedman, 1985; Altmann & Steedman, 1988) have pointed out that attaching a modifier phrase (such as a relative clause or a prepositional phrase) usually only makes sense when that modifier is needed to pick out one of two or more similar entities—here, to distinguish between two horses. In this story, when the final sentence is encountered, there's pressure to interpret the phrase *raced past the barn* as providing some more information about the horse; without this modifier, the simple noun phrase *the horse* wouldn't provide enough information to distinguish between the two horses, leading to confusion about which one was being referred to. But in a scenario that doesn't involve two horses—for instance, in a story about a single horse, or an out-of-the-blue sentence where the presence of two horses hasn't been established—the use of the modifier is unnatural. Speakers and writers don't generally go around providing extra, unnecessary information in the form of a modifier phrase. Without the proper contextual support for relative clause reading, readers will be more likely to default to the alternative main clause reading.

Reading time studies have shown that the right context can reduce or eliminate garden path effects, even for complex structures such as a reduced relative clause (e.g., Spivey-Knowlton et al., 1993). But the effects of context have been seen most vividly in studies of spoken language in which a visually present context provides a strong constraint for interpretation. In these experiments, researchers have used evidence from eye movements to infer the interpretations that are unfolding in the minds of listeners as they listen to ambiguous language.

As you've seen in Chapter 7, when people look at a visual scene while listening to spoken language, they try to relate the scene in front of them to what they're hearing. Since people look at images that correspond to the words they think they're hearing, this makes it possible to track eye movements to the visual scene as a way of inferring which word(s) they're activating based on spoken input. It takes a bit more ingenuity, but it's also possible to set up visual displays to test between competing syntactic interpretations, as demonstrated by Michael Tanenhaus and his colleagues (1995). For example, look at the scene in **Figure 8.3A** and imagine hearing an instruction like, "Put the apple on the towel into the box."

Where do you think you'd be looking at various points in the sentence? Naturally, upon hearing *apple*, people tend to look at the only apple in the display. But when they hear *towel*, there are now two possibilities. Which towel should they look at? It depends on which sentence structure they're entertaining. If they've attached *on the towel* directly to the verb phrase, they'll interpret the instruction as meaning they should pick up the apple and place it on the empty towel since it makes no sense to put the apple on the towel that it's currently sit-

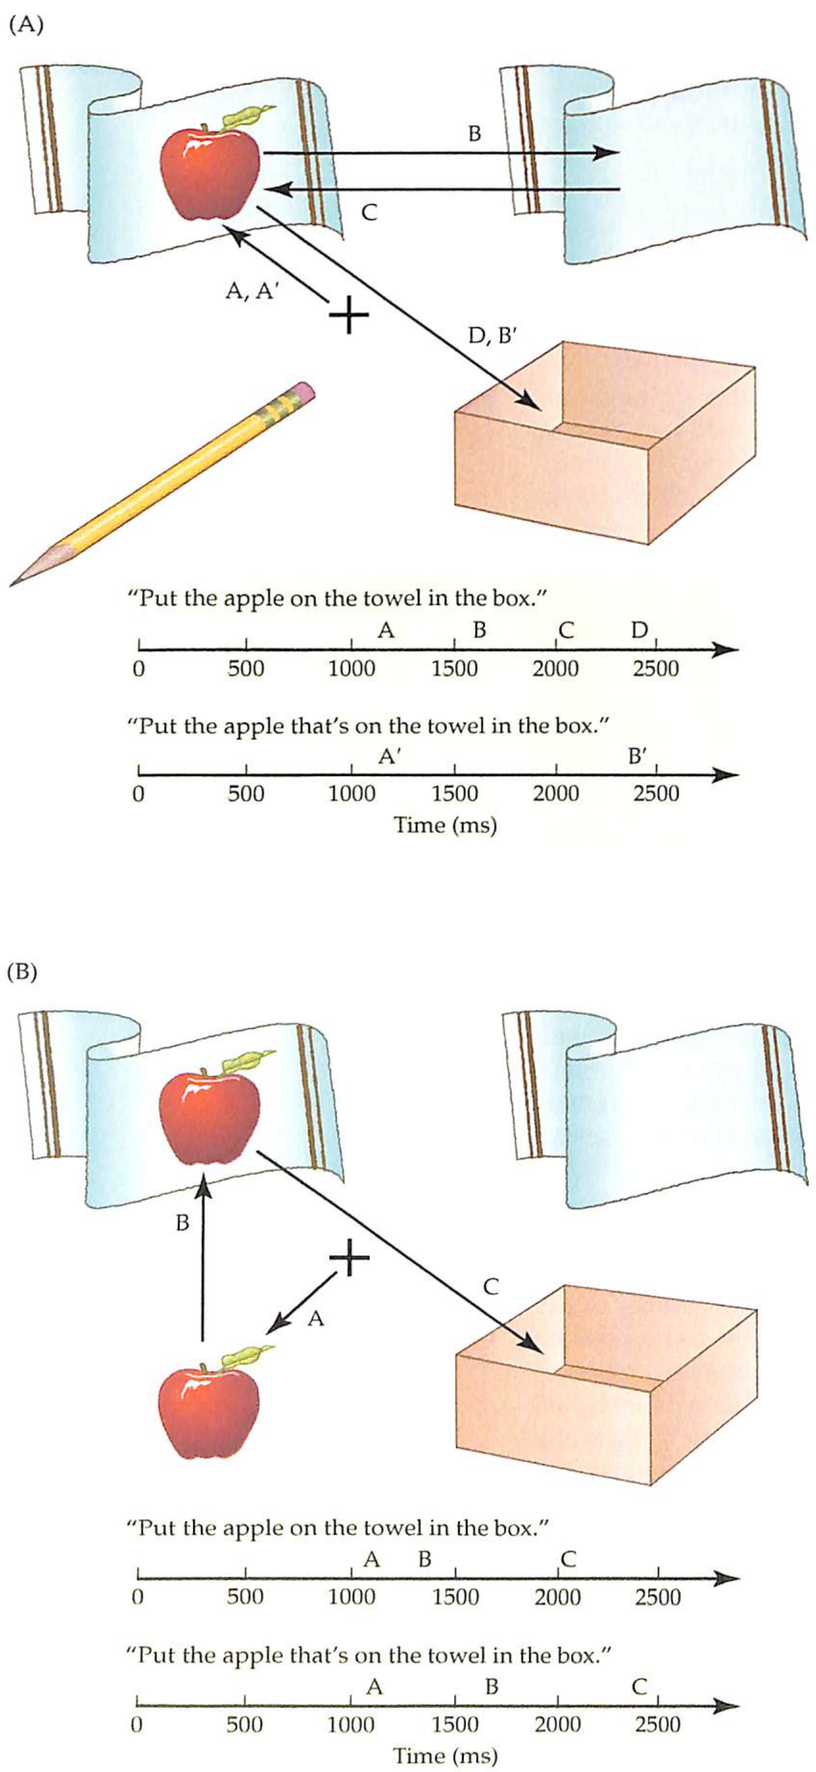

Figure 8.3 Visual displays and eye movement patterns used in the 1995 study by Tanenhaus et al. (A) Visual display with only one referent corresponding to the word *apple*. Letters indicate the typical sequence of eye movements and their timing relative to the spoken instruction. (A′ and B′ correspond to the unambiguous version of the instruction.) (B) Visual display with two referents corresponding to *apple*. Note that for this display, the sequence and timing of typical eye movements (relative to critical words in the speech stream) are the same, regardless of whether the instruction was ambiguous or unambiguous. (Adapted from Tanenhaus et al., 1995.)

ting on. Of course, they'll hit trouble when they get to *into the box*, since they'll already have interpreted the empty towel as the intended destination. On the other hand, if they've attached *on the towel* to the noun *apple* (specifying *which* apple should be moved), then they'll continue looking in the upper left square, which contains the object that corresponds to this noun phrase. They'll then have no trouble integrating *into the box*, since they're still waiting at this point to hear the destination phrase that will attach to the verb phrase. Critically, there should be no need to look at the empty towel, since this object is irrelevant for carrying out the instruction as understood. In other words, it's possible to use eye movements to the irrelevant towel as evidence of a garden path interpretation.

And evidence for a garden path is exactly what you see when the ambiguous sentence *Put the apple on the towel into the box* is compared with its unambiguous version *Put the apple that's on the towel into the box*. In the ambiguous version, people look at the irrelevant towel more often than in the unambiguous version (see Figure 8.3A).

This method turns out to be especially sensitive to the effects of context, particularly if that context can be set up visually so that hearers don't have to hold it in memory, as they might with a story context. Remember that attaching extra information to the noun is most natural in a situation where it's needed to distinguish between more than one similar entity. It's perfectly easy to visually set up just such a context, as shown in **Figure 8.3B**, where there are now two apples in the scene, only one of which is on a towel. In this visual context, when subjects hear *Put the apple on the towel*, there's a strong incentive to interpret the phrase *on the towel* as specifying which of the two apples is to be moved.

As it happens, the effects of the visual context are powerful; when faced with a scene like the one in Figure 8.3B, upon hearing *Put the apple on the towel into the box*, subjects very rarely look at the empty towel—no

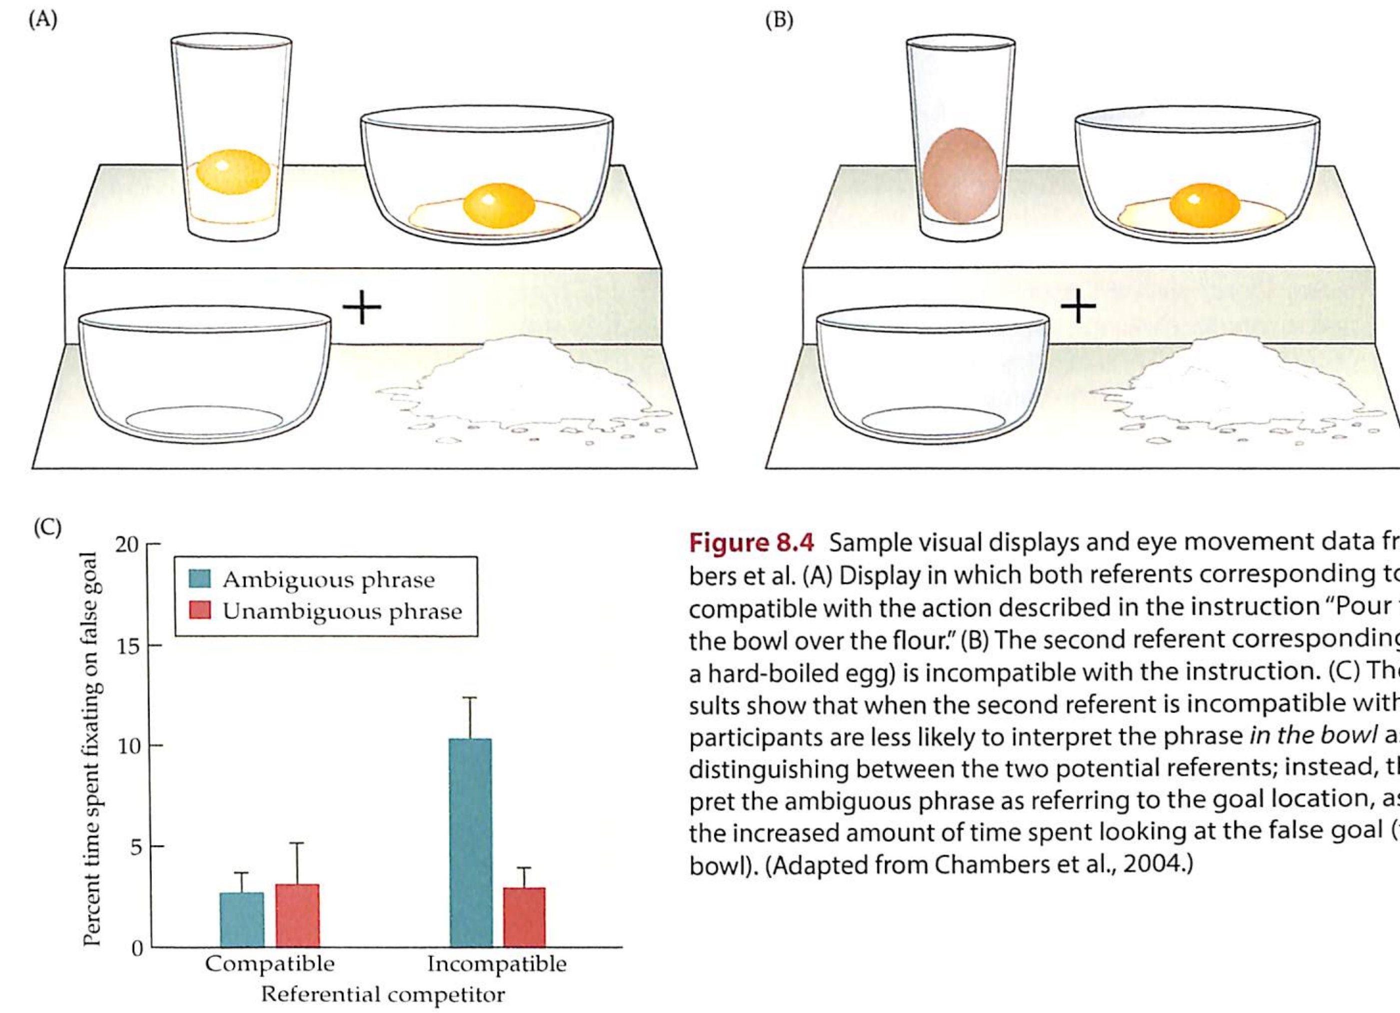

Figure 8.4 Sample visual displays and eye movement data from Chambers et al. (A) Display in which both referents corresponding to *egg* are compatible with the action described in the instruction "Pour the egg in the bowl over the flour." (B) The second referent corresponding to *egg* (i.e., a hard-boiled egg) is incompatible with the instruction. (C) The study's results show that when the second referent is incompatible with the action, participants are less likely to interpret the phrase *in the bowl* as helpfully distinguishing between the two potential referents; instead, they interpret the ambiguous phrase as referring to the goal location, as evident by the increased amount of time spent looking at the false goal (the empty bowl). (Adapted from Chambers et al., 2004.)

more so, in fact, than when they hear the unambiguous version of the instruction. Hence, the eye movement record shows no evidence of a garden path interpretation when the context is appropriate to the normally less preferred structure.

Even more subtle effects of context have been demonstrated by Craig Chambers and his colleagues (2004). In their experiment, they set up displays such as those illustrated in **Figure 8.4**, simulating a cooking environment in which objects needed to be manipulated according to spoken instructions such as, "Pour the egg in the bowl over the flour" (or its unambiguous counterpart, "Pour the egg that's in the bowl over the flour"). You'll notice that in both versions of the visual display, there are two eggs, so the phrase *in the bowl* is potentially helpful in identifying which of the two eggs is the target of the instruction. However, in Figure 8.4B (unlike in Figure 8.4A) the display contains only one egg in liquid form that could possibly be poured—the other egg is still inside its shell. Hence, upon hearing *Pour the egg*, the subject might infer that since the action could only target one of the eggs, there's no need for additional information to specify which of the eggs is under discussion. If hearers are indeed making these inferences based on the visual context, then they might be tempted to misinterpret the ambiguously attached phrase *in the bowl* as designating the location where the egg is to be poured, rather than as further information about the targeted egg. Such an interpretation would be revealed in increased glances to the empty bowl, the "false goal" in the display. In fact, Chambers and his colleagues showed that just such a garden path effect did emerge in the visual displays with a single pourable

BOX 8.4

Doesn't intonation disambiguate spoken language?

Many people assume that in spoken language, unlike in written language, true ambiguity simply doesn't arise; the belief is that while a string of words may be ambiguous, the *way that you say those words* will be different depending on your intended meaning, allowing your hearer to avert any possible misinterpretation. And in many cases, this is true. For example, I'm sure that you can come up with six or seven different ways to utter the word *really*, each of them clearly communicating something different from the others.

But let's look at the issue a little bit more carefully in the context of the garden path sentences that we've come to know and love. Here, the situation is a bit different from your ability to clarify your intended use of the word *really*—in that case, intonation was serving to signal a *general* communicative purpose such as asking a question, expressing incredulity or emphasis, or indicating doubt. In the case of garden path sentences, though, intonation is being asked to signal one *structure* rather than another. Under what conditions would it be possible for intonation to disambiguate structure? (Actually, it would be more accurate to talk about the disambiguating potential of **prosody**, a notion that includes intonation, but also rhythmic aspects of speech such as pausing and lengthening of sounds.)

One possibility might be that particular sentence structures match up with specific prosodic patterns in a one-to-one fashion. But this turns out not to be true. The relationship between syntax and prosody is much more slippery than this. There's a general tendency, for example, to insert pauses or stretch out the last syllable of a phrase just before the boundary between major constituents (you saw evidence of this in Language at Large 6.1, in our discussion of constituents in literary language). For example, ambiguities that revolve around whether a clause is to end or to continue lend themselves especially well to prosodic disambiguation. Consider the following ambiguous fragment:

While the pop star sang the national anthem...

which, as you've seen, could continue like this:

... played in the background

or it could continue like this:

... the people stood respectfully

In the first continuation, the desire to insert a pause is very strong, in keeping with the convention of inserting a comma:

prosody The rhythm, stress, and intonation of a spoken phrase or sentence.

egg (see Figure 8.4C). However, when two pourable eggs appeared in the display, there was no evidence of misinterpretation, with eye movements being similar for both the ambiguous and unambiguous versions of the instruction. This shows an impressive ability to closely integrate the incoming stream of speech with detailed aspects of the visual context. Contrary to the predictions of the garden path model, people seem to be able to recruit context quickly enough to avert garden path effects for a number of syntactic ambiguities.

You might be inclined to wonder just how often the need for such context-sensitive inferences arises in actual spoken language. After all, isn't it true that we can communicate our intended meaning through intonation, even when a given string of words has more than one meaning? This is a common intuition, but while intonation can sometimes help to disambiguate speech, it certainly doesn't sweep away all or even most potential ambiguities (see Box 8.4). Rather, intonation seems to be just one more useful cue among many that we can use intelligently to help us untangle strings of words and their many potential meanings.

WEB ACTIVITY 8.3

Create and recognize "easy" versus "hard" ambiguous sentences In this activity, you'll be guided through a series of steps in which you'll generate sentences that contain a potential temporary ambiguity. Some of these will be likely to cause very severe garden path effects, while others will be likely to cause only minimal disruption. You'll hone your editorial instincts and learn to recognize the potentially troublesome sentences.

BOX 8.4 *(continued)*

While the pop star sang, the national anthem played in the background.

But boundaries between constituents that are smaller than clauses are marked with much subtler prosody—and in fact, whether or not they're marked at all can depend on factors such as the length of the phrases that are involved, so you'd be more likely to mark the boundary here:

The patient told the doctor [$_S$ that he was having some ongoing trouble with his injured knee].

than here:

The patient told the doctor [$_S$ that he smoked].

But even if prosodic patterns don't match up neatly with syntactic structures, it's still possible that speakers would choose to use cues like pauses and syllable lengthening as a courtesy to the hearer, in order to mark the correct way in which the words should be grouped into constituents. This, too, turns out to be not quite true. A number of studies have shown that speakers fail to reliably disambiguate speech for their hearers. For example, David Allbritton and his colleagues (1996) recorded the speech of trained actors as well as psychology undergraduates while they read sentences aloud, and found that *neither* group spontaneously disambiguated sentence structure through prosody. When instructed, the actors at least were able to use prosody to signal their intended meaning, but only when both meanings of the ambiguous sentence were presented to them side by side. This finding is consistent with a number of other similar experiments.

Things get a bit better when speakers and hearers are involved in an interactive task such as a game in which the speaker instructs the hearer to move objects around on a board. For example, Amy Schafer and her colleagues (2000) found that under these circumstances, when speakers became aware of their partners' misinterpretations, they did use prosody somewhat consistently to disambiguate their instructions. But overall, the evidence suggests that in many linguistic settings, it's quite difficult for speakers to anticipate ambiguities that will prove difficult for their hearers and to be helpful in the prosodic patterns that they produce.

Still, some systematic alignment between syntax and prosody does occur, probably as a by-product of the processes that are involved in planning and uttering speech, as you'll see in Chapter 9. Prosody is more helpful for some structures than for others, but there's enough information there to be of some use to hearers for the purposes of disambiguation, as demonstrated by Tanya Kraljic and Susan Brennan (2005). So, while prosody certainly doesn't render spoken language unambiguous, hearers can use it as one more helpful cue, along with the many other sources of information discussed in Section 8.3.

WEB ACTIVITY 8.4

Does prosody disambiguate? In this activity, you'll listen to sound files of speakers uttering ambiguous sentences. You'll be asked to determine whether you can reliably understand the intended meanings of the sentences based on the speaker's prosody.

Humans understand language despite rampant ambiguity

By now, we've compiled quite a list of variables that seem to have the effect of either aggravating or relieving garden path effects. Overall, experimental evidence suggests that the human parser is quite a bit more flexible, intelligent, and less error-prone than was originally envisioned. In order to predict whether a specific sentence will cause a serious derailment for readers or hearers, it's not enough just to know that it contains a potential ambiguity. People are able to combine information from a large variety of informative cues to help resolve an ambiguity. If a number of different cues all conspire to point to the wrong interpretation, readers or listeners will run into a dead end, and have to backtrack in their processing attempts. But situations in which these cues overwhelmingly point to the wrong structure are extremely rare.

This turns out to be a good thing for language comprehension, because as it happens, language is even more inherently ambiguous than psycholinguists might have imagined. The sheer scope of linguistic ambiguity became apparent to computer scientists once they started building computer programs whose goal was to understand sentences produced by humans. They found that for many partial sentences of English, the programs would offer not one or two, but sometimes *thousands* of ways of parsing the sentence in conformity with the rules of English grammar. Most of these massively ambiguous sentence fragments are not at all problematic for humans—we filter out almost all of the grammatically legal structures as nonsensical or enormously unlikely to ever be uttered. But currently, computers are much less good at coping with ambiguity than we are. Attempts to build machines that can understand us have highlighted just how impressive we are at resolving ambiguities in language.

Is it possible to have a language without ambiguity? As far as I know, naturally occurring languages all involve ambiguity, but it is certainly possible to invent an ambiguity-free artificial language, as was done by James Cooke Brown (1960) when he created the language known as Loglan. In Loglan, words never have more than one meaning, and they're combined in ways that resemble logical formulas—these formulas convey conceptual content in a transparent way so that you can never have one string of words that means more than one thing. Brown originally invented Loglan with the idea that it would be interesting to see whether a completely unambiguous language could help its speakers organize their thinking more clearly. But more recently, enthusiasts of Loglan (and a related language known as Lojban) have argued that unambiguous languages could prove to be very useful in the interaction between humans and computers, since ambiguity has such a catastrophic effect on computers' ability to understand language.

Would humans benefit from using an unambiguous language among themselves? Despite the hype from some Loglan and Lojban enthusiasts, it's not clear that they would. As you've seen, people are extremely good at cutting through much of the ambiguity inherent in language. And as discussed in Chapter 7, ambiguity might serve to lessen the demands on language *production*. In Box 7.2, I introduced the idea that ambiguous words allow languages to reuse easy-to-produce bits of language with minimal disruption to comprehension. In the case of syntactic ambiguities, it may be that a language that contains ambiguity allows for a simpler syntactic system, or allows speakers some important flexibility in how they order words and phrases. The notion that ambiguity might offer some important benefits is echoed by author Arika Okrent (2009), who discusses her attempts at learning Lojban. She suggests that "composing a sentence in Lojban is like writing a line of computer code. Choose the wrong function, drop a variable, forget to close a parenthesis, and it doesn't work." As a result, she notes, watching live conversations in Lojban is much like watching people slowly doing long division in their heads.

8.4 Making Predictions

Eyeing the content yet to come

In the previous section, you learned that it's quite rare for ambiguity to be truly disruptive for language processing—the human parser seems to be very adept at coping with the uncertainty introduced by strings of words that are ambiguous. We regularly make use of our experience with language patterns, our knowledge of the world, or our awareness of the specific context to choose the most likely meaning when several are possible. This allows us to process language incrementally and on the fly, even when the fragments we've heard are highly ambiguous.

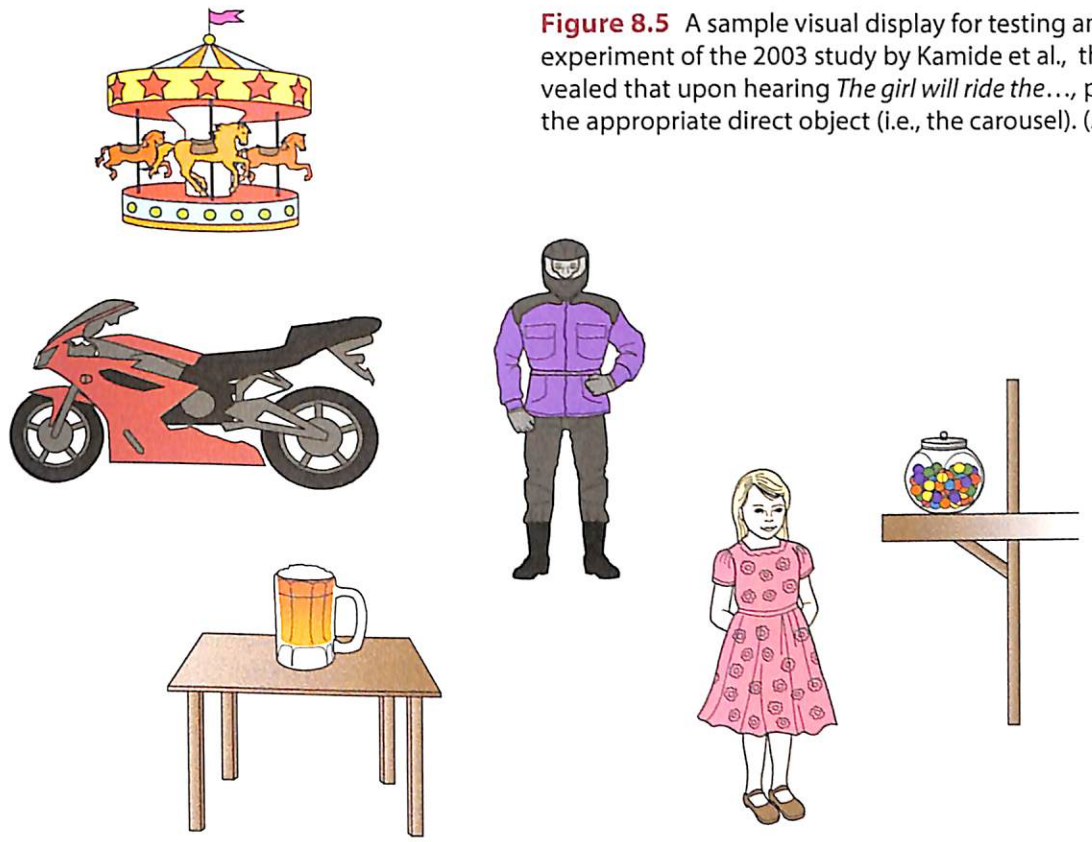

Figure 8.5 A sample visual display for testing anticipatory eye movements. In one experiment of the 2003 study by Kamide et al., the subjects' eye movement data revealed that upon hearing *The girl will ride the*..., participants were already predicting the appropriate direct object (i.e., the carousel). (Adapted from Kamide et al., 2003.)

However, researchers have shown that we go even further than this. Not only do we *choose* among multiple possible structures for incoming fragments of speech, but we can *anticipate* the meanings of sentences, even before we've heard critical portions of these sentence.

Anticipatory, or predictive, language processing has been demonstrated quite clearly in a number of eye-tracking experiments. In one of these studies, Yuki Kamide and her colleagues (2003) showed people visual displays illustrating several characters and various objects, and observed their eye movements as they listened to sentences involving these characters and objects. For example, while viewing a picture such as the one in **Figure 8.5**, subjects would hear one of the following sentences:

The man will ride the motorbike.

The man will taste the beer.

The girl will ride the carousel.

The girl will taste the sweets.

The researchers found evidence that people were anticipating the identity of the direct objects in the sentence. That is, upon hearing *The girl will ride the*..., they were already more likely to be looking at the carousel than at any of the other objects in the display. (In contrast, hearing *The man will taste the*... prompted glances at the picture of the beer.) Listeners' guesses about the upcoming direct objects reflected the *combination* of the previously heard subjects and verbs—that is, people didn't just look at the rideable objects (the carousel *and* the motorbike), or the objects that they think girls like (the carousel and the candy). They looked at whichever object was most likely to be involved in an event involving the already-mentioned subject and verb. This suggests that they had integrated linguistic information about the subject and the verb, used this information to build a partial meaning of the unfolding sentence, and were generating predictions about the most likely way to fill in the sentence's meaning.

In another experiment, Kamide and her colleagues leveraged the grammar of Japanese to show that predictions about meaning are based on a subtle interplay between linguistic and contextual information. Unlike in English, in which the verb usually appears before the direct object, in Japanese the verb typically appears at the end of the sentence after both the subject and its direct and indirect objects. For example, here's the Japanese translation for the sentence *The waitress will merrily bring the hamburger to the customer*:

weitoresu-ga	*kyaku-ni*	*tanosigeni*	*hanbaagaa-o*	*hakobu.*
waitress	customer	merrily	hamburger	bring.

Japanese subjects heard sentences such as these while viewing visual displays (see **Figure 8.6**).

What kind of predictions were listeners making after they'd heard just the first two or three words of the sentence? Since waitresses usually bring people food, it would be sensible for the listeners to assume that this is the direction the sentence was taking, and to throw some anticipatory glances at the hamburger. But waitresses can also do many other things involving customers: they can greet them or tease them, take their orders, recite the menu, and so on. What information could the Japanese hearers use to constrain the set of possibilities?

One potentially useful source of information is the presence of case markers in Japanese—these are grammatical tags that appear on nouns to indicate whether the noun is a subject (in which case it appears with the *-ga* tag), an indirect object (*-ni*), or a direct object (*-o*). In the above sentence, the noun meaning "customer" is tagged with an indirect object marker. This means that the customer could be involved in, say, an event of bringing, which does involve an indirect object; but he couldn't be the target of a teasing event, which involves

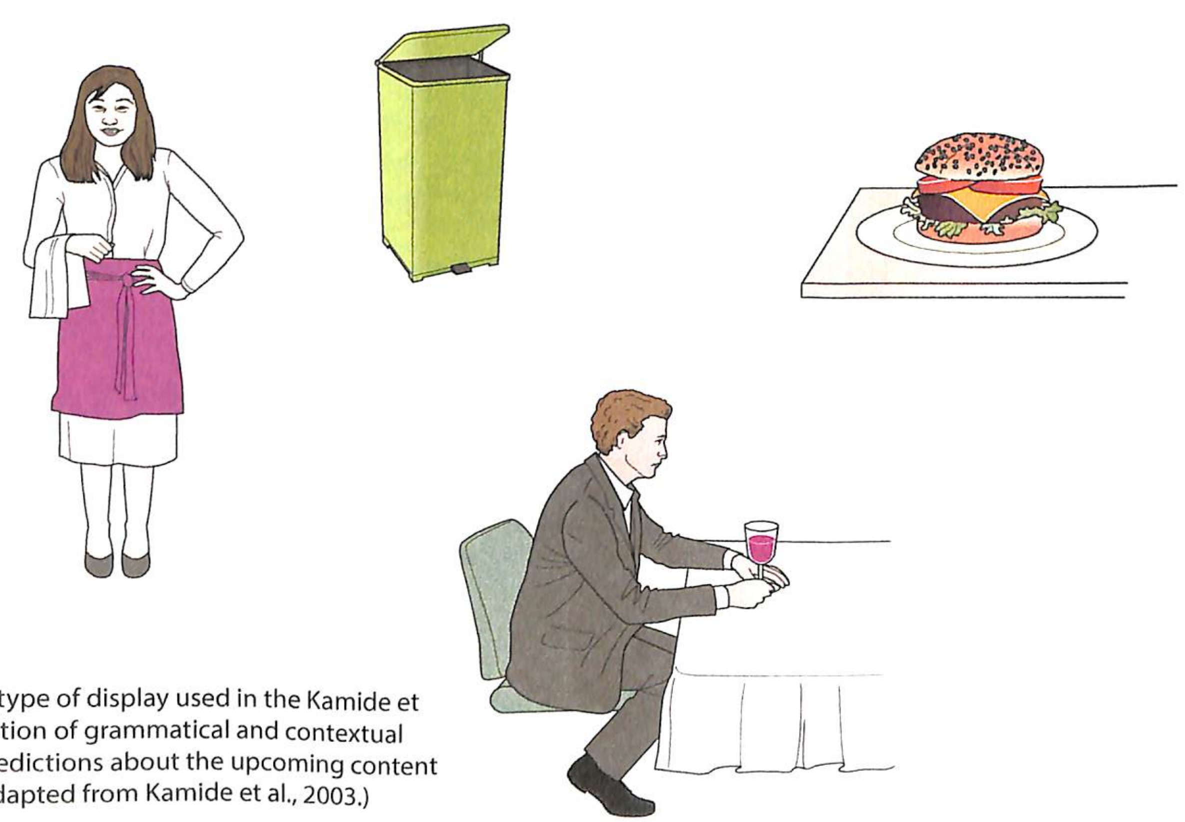

Figure 8.6 A sample of the type of display used in the Kamide et al. study to probe the integration of grammatical and contextual information in generating predictions about the upcoming content in the sentence (see text). (Adapted from Kamide et al., 2003.)

a direct object, but no indirect object. In fact, if the waitress were teasing the customer, we'd have to tag the noun with the direct object tag *-o* rather than *-ni*:

weitoresu-ga	*kyaku-o*	*tanosigeni*	*karakau.*
waitress	customer	merrily	tease.

Is predictive language processing sensitive to such linguistic subtleties? Kamide and her colleagues showed that it is. Upon analyzing their subjects' eye movements, they found that after hearing only the first three words (that is, the equivalent of *waitress customer merrily*), the subjects were more likely to look at the hamburger than at the distractor object (the garbage can)—but only if the word for "customer" was tagged with an indirect object marker, signaling that the word was grammatically consistent with an event in which the waitress brought the customer some food. When the case marker was inconsistent with such an event, the subjects were not lured by the picture of the hamburger.

Brain waves reveal predictive processing

One of the great advantages of the eye-tracking method, as compared with button-press response time measures, is that it allows researchers to get a moment-by-moment view of what's happening in people's minds as they process language. This is because the target behaviors (that is, eye movements) are tracked continuously as speech unfolds, rather than being probed at one particular point in time *after* the presentation of some critical stimulus, as is the case in response time studies. Naturally, this makes eye-tracking a powerful method for studying predictive processing, where we want to see what's happening in the mind *before* people encounter some critical stimulus. But the technique comes with certain limitations as well:

1. You have to be able to visually depict the content of the sentence, which restricts the abstraction and complexity of the stimuli.
2. Eye movements can only reveal how people respond to those aspects of the sentence that are visually depicted, but they're fairly uninformative about whether people are generating predictions about content that isn't present in the display. In a sense, the visual displays may be acting as a sort of multiple choice test situation—we can see whether people are favoring certain hypotheses over others among the options that are provided, but we can't easily tell whether they're generating hypotheses that *aren't* made visually available as options, or whether they would do so in a more open-ended language-processing situation. This is important because presumably, much of daily language processing takes place in less visually constraining contexts.

One way around these limitations is to study brain wave activity while people hear or read language. This approach offers researchers many of the advantages of eye-tracking methods, without the constraints imposed by visual presentation. Like the eye-tracking studies, experiments that measure event-related potentials (ERPs) can tap into moment-by-moment responses to language as it unfolds.

A number of researchers have used ERP techniques to ask: Do people anticipate the actual *words* that will appear downstream, rather than just the general content that's yet to come? In one such study, Katherine DeLong and her colleagues (2005) had subjects read sentences like these:

The day was breezy so the boy went outside to fly a kite.

The day was breezy so the boy went outside to fly an airplane.

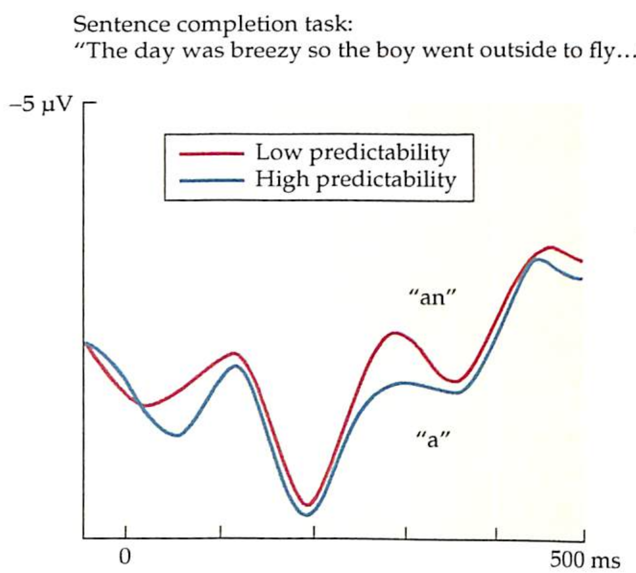

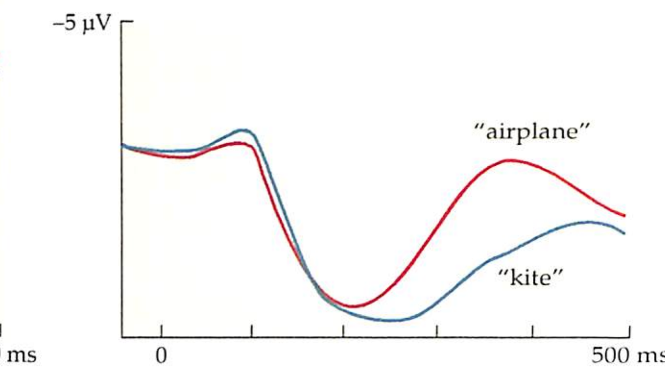

Figure 8.7 ERP data from the study by DeLong et al. (2005). This graph plots the ERP waveforms at the midline central (vertex) recording site for words that were determined to be of high versus low predictability based on a separate sentence completion task. Note that negative values are plotted upward. Both articles and nouns that were low in predictability showed significantly greater negativity between 200 and 500 ms after the onset of the word (N400) compared with articles and nouns of high predictability.

If people are generating predictions about upcoming words, they're probably more likely to come up with *kite* than with *airplane*. If you remember the discussion of N400 brain waves back in Chapter 3, unexpected words result in heightened N400 activity compared with more predictable words. (This is the case regardless of whether "predictability" reflects the semantic plausibility/anomaly of a word in the context, or simply a contrast between a common versus rare word.) Hence, subjects should display more energetic N400 activity upon reading the word *airplane* at the end of the sentence.

This is exactly what DeLong and her colleagues found (**Figure 8.7**). But more importantly, they found heightened N400 activity for the "airplane" sentences at the *previous* word, the article *an*. Notice that in the sentences above, the more predictable word *kite* should be preceded by the article *a*, following the English rule of phonology that says the article *an* only appears before vowel-initial words. By cleverly setting up a contrast in the articles that could appear with predictable and unpredictable words, the researchers were able to probe for clear evidence of word anticipation. The earliness of the N400 effect suggests that even before getting to the word *airplane*, subjects must have generated some expectations about the specific *shape* of the upcoming word; otherwise they wouldn't have been surprised by the occurrence of the article *an*.

Subsequent ERP studies support the conclusions of DeLong and her colleagues: people appear to generate some predictions about the form as well as the content of upcoming material. Moreover, their predictions show sensitivity to a variety of information sources. In the study that I've just described, subjects were presumably relying on their general knowledge of the world in anticipating that boys are more likely to fly kites than airplanes on a breezy day. But another study showed that people's predictions relied more on the specific story context than on real-world plausibility—for example, in reading a story about an amorous peanut, they showed *less* N400 activity (reflecting greater expectation) for the sentence *The peanut was in love* than for *The peanut was salted* (Nieuwland & Van Berkum, 2006). Moreover, the identity of the speaker can also affect a hearer's expectations. In a study led by Jos Van Berkum (2008), subjects exhibited an N400 effect upon encountering a word that would be unexpected from a particular speaker. For instance, if someone

with an upper-class accent uttered, "I have a large tattoo on my back," subjects showed heightened N400 activity at the word *tattoo*; in fact, this effect occurred very early in the utterance of the word, leading the authors to argue that hearers must have been generating expectations of likely upcoming words even before the speaker began to say *tattoo*.

Predictions in language and hockey: A matter of expertise

Together, the eye-tracking and ERP studies reveal an extremely active and adaptive language-processing system. Instead of waiting for the speech signal to trigger the process of word recognition and subsequent structure-building, we quickly pull together all the information that's available to us at any point in time, and generate expectations about the sentence's upcoming content and form. Naturally, these preliminary expectations might turn out to be wrong, in which case we have to be able to quickly revise them.

Our ability to generate and revise predictions appears to be a critical aspect of smooth and efficient language processing. Increasingly, psycholinguists see language processing—whether or not it involves ambiguity—as the business of making predictions about a sentence's future trajectory, much as you might make predictions about the trajectory of objects moving in physical space. In some cases, this means making predictions about how an ambiguous structure will pan out. But it can also involve generating expectations about the type of event that is being described, about whether a just-uttered word like *the* will be followed by an adjective or directly by a noun, or even stylistic expectations about whether the speaker is likely to refer to a youngster by using the word *child* or the more informal *kid* in a particular social situation.

Some researchers have argued that the idea of predictability in language can offer a unified approach to understanding the processing of both ambiguous and unambiguous sentences. For example, John Hale (2001) has used the term **surprisal** to capture this notion, arguing that processing difficulty is a reflection of whether a sentence's continuation is highly predicable, and whether the unfolding sentence ends up conforming to the most likely outcome. Surprisal is lowest when the probability that a sentence will continue in a particular way is very high, and the sentence does in fact continue in the most likely manner. The first two examples below (only one of which involves a syntactic ambiguity) illustrate a low level of surprisal upon encountering the words in bold—and hence, these words should be very easy to process:

The horse raced past the barn **and fell**.

Lin likes to eat bacon with her **eggs**.

Surprisal is highest (leading to greatest processing difficulty) when the sentence actually continues in a way that was extremely *improbable*, as shown by the words in bold in the next two examples:

The horse raced past the barn **fell**.

Lin likes to eat bacon with her **cheesecake**.

Intermediate levels of surprisal occur when the unfolding sentence is reasonably consistent with a number of possible outcomes, one of which turns out to be realized, as illustrated by the words in bold in the final two examples below:

The landmine buried in the sand **exploded**.

Lin bought some **potatoes**.

A great deal of knowledge and skill goes into being able to make linguistic predictions. In the physical world, your ability to make predictions about

surprisal A measure that's inversely related to the statistical predictability of an event such as a particular continuation of a sentence. Processing difficulty is thought to reflect the degree of surprisal at specific points in the sentence, so that less predictable continuations result in greater processing difficulty.

where a moving object will land depends on your experience with moving objects. In much the same way, making linguistic predictions is dependent on your experience with language—with its grammar, with its common words and structures, and with how people tend to wield it in a variety of situations. Now, predicting an object's motion trajectory is usually far simpler than predicting the path of a sentence. But to push the analogy a bit further, think of language processing as similar to the skills that are required for playing hockey at a very high level, and predicting the movement of the puck. According to author Adam Gopnik (2011), hockey requires, above all, the ability to "see the ice":

> … the ability to grasp a changing whole and anticipate its next stage. It's the ability to make quick decisions, to size up all the relationships in a fast-changing array and understand them. A related notion is that of situational awareness: a heightened consciousness of your surroundings and both the intentions of the people around you and their anticipated actions.
>
> Well, hockey, obviously, which is played at incredibly high speed, reveals and rewards situational and spatial intelligence at a degree of difficulty that no other sport possesses. So much so, that the greatest of all hockey players, Wayne Gretzky, had, besides his other significant skills as a fine-edge skater, almost nothing else that he's specifically good at. That is his gift—the gift of spatial and situational intelligence: knowing what's going to happen in three seconds, anticipating the pattern approaching by seeing the pattern now, sussing out the goalie's next decision and Jari Kurri's or Luc Robitaille's eventual trajectory in what would be a single glance if a glance were even taken. Gretzky is the extreme expression of the common skill the game demands.

To many psycholinguists, this is as apt a metaphor as there is for the "common skill" that is demanded by everyday language processing. Perhaps if all of us spent as much time playing hockey as we did using language, we too would have the expertise of a Wayne Gretzky.

Like hockey, making language-based predictions does appear to be an acquired skill. Arielle Borovsky and her colleagues (2012) set out to learn whether children, like adults, make predictive eye movements, in a study very similar to the eye-tracking experiment by Kamide et al. (2003). They tested adults and children age 3 to 10 years, presenting their subjects with spoken sentences such as, "The pirate hides the treasure," along with visual displays of various objects related to the sentences. They measured the point in time at which a subject began to look more often at the picture of the treasure than at the other objects in the display, in anticipation of the direct object.

As in the earlier study by Kamide and colleagues, adults in the Borovsky et al. study regularly made predictive eye movements very soon after hearing the beginning of the verb. Children did too, although their eye movements were somewhat more sluggish than the adults', suggesting that the predictions about the upcoming sentence content were not being generated as quickly. Of special interest, though, was the fact that children with relatively large vocabularies launched speedier anticipatory looks than their peers with smaller vocabularies. Knowledge of language, more than age itself, was correlated with the speed of their predictions. And it might surprise you to learn that, even among the adults, vocabulary size was correlated with the speed of predictive processing. In the language arena, as in hockey, it appears that some people are more expert than others at anticipating what is going to happen next—though much more research still needs to be done to understand the source of the variability of this skill. And, alas, as with hockey, aging may result in a decline in the ability to anticipate upcoming words, as discovered in ERP studies by Kara Federmeier (2007).

WEB ACTIVITY 8.5

Linguistic predictions In this activity, you'll explore a range of language phenomena in which expectations about upcoming structure or content could potentially guide language processing.

8.5 When Memory Fails

When are complex sentences difficult?

Let's take another look at the difficult sentences introduced at the beginning of this chapter. As you've seen, ambiguity is a common culprit when it comes to creating havoc for language comprehension. But sentences can be hard to process for other reasons too. Several of the most difficult sentences from the introductory section could not readily be pinned on an ambiguity. For example:

The mouse the cat chased keeled over.

The administrator who the intern who the nurse supervised had accused fudged the medical reports.

It's easy to see that these sentences involve more than one clause or sentence unit, and encode multiple events (for instance, in the first sentence, there's an event of the cat chasing the mouse, and an event in which the mouse keels over). But that alone can't explain the difficulty, because it's certainly possible to express these same events in a sentence without straining the processing system:

When the cat chased the mouse, the mouse keeled over.

The mouse keeled over when the cat chased it.

And oddly, if we want to say that it was the cat that keeled over, rather than the mouse, the next sentence seems to be considerably easier to understand:

The cat that chased the mouse keeled over.

So, it's not just that the sentence is generally complex or the ideas inherently difficult to understand; it's that there's something taxing about the particular structure with which those ideas are expressed. For instance, the first sentence of the following pair would typically be read faster than the second, even though they're made up of exactly the same words:

The senator who spotted the reporter shouted.

The senator who the reporter spotted shouted.

What's the key difference? Both contain a relative clause attached to the subject of the main clause, *the senator*. But in the first sentence of the pair, *the senator* is also linked to the subject role of the embedded verb *spotted*. (Notice what happens in the question/answer pair: *Who spotted the reporter? The senator spotted the reporter.*) This type of relative clause is called a subject relative clause. In contrast, in the second sentence, *the senator* is linked to the direct object role of the embedded verb *spotted*. (Notice: *Who did the reporter spot? The reporter spotted the senator*). This yields an object relative clause. As I did in Chapter 6, I'll notate the relationships by using the same color for the phrase *the senator* and the underscore marking the position in the sentence that it links to:

The senator who ___ **spotted** the reporter shouted.

The senator who the reporter **spotted** ___ shouted.

So: why are object relative clauses harder to process than subject relative clauses?

Memory failure or faulty predictions?

As with the ambiguous sentences discussed in Section 8.2, there are two competing classes of explanations for the difficulty: one that highlights the inherent limitations of the parser, and another that emphasizes its tendency to predict structure. According to the first account (a version of which is articulated in great detail by Ted Gibson, 1998), the trouble crops up because complex sentences involve holding some parts of a sentence in memory until they can be related to other parts of the sentence—as the distance between these related parts gets bigger, so do the demands on memory and the prospects of processing failure. And, in some cases, multiple dependencies have to be held in memory at the same time, adding even more strain on memory.

In a subject relative clause, an NP (*the senator*) has to be integrated with the embedded verb (*spotted*), as well as the main clause verb (*shouted*), leading to *two* dependencies:

The senator who ___ **spotted** the reporter **shouted**.

This turns out to be fairly easy: at *spotted*, there's no intervening NP to get in the way of the integration of the two components, and by the time *the reporter* is reached, the first dependency has already been resolved. Things are a bit different with the object relative version:

The senator who the reporter **spotted** ___ **shouted**.

Readers have to hold *the senator* in memory while encountering *the reporter*, and it's the intrusion of this second NP that makes the integration with *spotted* that much harder. The sentence gets harder still if you add yet another intervening NP:

The senator who the reporter next to the president **spotted** ___ **shouted**.

Second, while the intervening phrase *the reporter* is being encountered and stored, neither of the two dependencies involving *the senator* has yet been resolved, adding to the memory burden. You can see the effect of piling up the number of unresolved dependencies by taking the original object relative sentence and adding yet another object relative clause:

The senator who the reporter who the president detested ___ spotted ___ shouted.

For most people, this sentence disintegrates into utterly unrecognizable word mush. But again, there's nothing difficult about the *ideas* expressed in the sentence, as we can see if we unravel the sentence and present the same content in a way that avoids the memory crash:

The president detested the reporter who ___ spotted the senator who ___ shouted.

There's yet a third memory-related issue that may play a role in these horrendous sentences: the semantic similarity of the nouns *reporter* and *senator*. Some researchers have pointed out that it's a general property of memory that similar items tend to blur together and become harder to retrieve separately. It's possible that when complex sentences involve very similar nouns that have to be held in memory at the same time, this might make it hard to retrieve the right noun in the right place in the sentence. The experience might be analogous to reading a long Russian novel populated by a large cast of characters, all of whose names sound alike—it becomes onerous to keep track of who did what to whom. Indeed, readers often seem to fare better with object relative sentences that contain very dissimilar nouns. Compare:

The senator who the president summoned arrived.

The helicopter that the president summoned arrived.

These explanations all derive from the very reasonable assumption that the process of building syntactic structure taps working memory resources, which, sadly, are known to be finite. But not all researchers focus on memory limitations in explaining the difficulty of object relative clauses. Some have suggested that the difficulty is a by-product of the parser's tendency to predict upcoming structure based on previous experiences with language; in this case, experience with the statistical likelihood of different structures leads the parser astray. In a sense, you might think of the subject-relative/object-relative asymmetry as simply another kind of very transient garden path effect. To see this, have several people you know (but not from this class) complete the following fragment:

The beauty queen who ...

How many of them produced an object relative clause, as opposed to a subject relative clause? Based on the statistical patterns of English, chances are that the subject relatives came up more frequently, with most people providing a verb rather than a second NP as the next element to continue the fragment. In that case, according to an experience-based account, they'd have trouble reading sentences with object relative clauses (for example, *The beauty queen who the judge disqualified broke into tears*) not because they'd drained some finite processing resources, but because their perfectly intelligent and reality-based expectations about upcoming structure were not met once they ran into the more unusual sentence structure.

Both of these accounts are well grounded in reasonable assumptions about language processing, assumptions for which there's a fair bit of independent evidence. Luckily, there are ways to pull apart the two explanations—for instance, they make somewhat different predictions about *where* in the object relative clause readers should experience the greatest difficulty. For experience-based accounts, this should happen right upon encountering the second NP, where expectations of a subject relative structure are confounded (as in *The beauty queen who* ***the judge*** *disqualified broke into tears*.) For memory-based accounts, the trouble spot should be at the embedded verb, where the noun has to be integrated with the verb, over an intervening NP (*The beauty queen who the judge* ***disqualified*** *broke into tears*.) These are subtle points, but it seems likely at this stage of data gathering that both explanations correctly capture portions of the difficulty with object relative clauses.

LANGUAGE AT LARGE 8.2

Straining the parser for literary effect

What makes for good writing? A sensible editorial response might be: writing that puts as little stress as possible on the language processing system. If that's the case, no one seems to have told the novelist Henry James, famous for his convoluted syntax. Here are some typical Jamesian sentences from his 1903 novel *The Ambassadors*:

> *One of the features of the restless afternoon passed by him after Mrs. Peacock's visit was an hour spent, shortly before dinner with Maria Gostrey, whom of late, in spite of so sustained a call on his attention from other quarters, he had by no means neglected.*
>
> *Chad offered him, as always, a welcome in which the cordial and the formal—so far as the formal was the respectful—handsomely met; and after he had expressed a hope that he would let him put him up for the night, Strether was in full possession of the key, as it might have been called, to what had lately happened.*

Why would a writer make a conscious choice to create sentences that wind their way through multiple embeddings, stretch syntactic dependencies over dizzying distances, and follow bizarre twists of structure, all of which make the prose harder to read? Author and critic Zadie Smith (2009) suggests that the technique was part of James's attempt to cultivate a more acute consciousness in his reader, that his syntactic choices were "intended to make you aware, to break the rhythm that excludes thinking."

Like most literary writers, James (and Smith) likely relied on intuitions about language. But from a scientific perspective, the idea's not crazy. It seems strange and counterintuitive, but a number of studies suggest that when information is *too* easy to process fluently, people are somewhat prone to thinking *less* deeply, to falling back more readily on fast-but-dumb cognitive heuristics.

One intriguing example comes from a study led by Adam Alter (2007), in which subjects had to answer math problems like the one you saw earlier: In a lake there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half the lake? Many people mistakenly answered 24, but they were less likely to do so if they'd read the problem in a nearly illegible font. One interpretation of this intriguing finding is that the font made the problems *feel* harder, thereby kick-starting a more careful mode of thinking. (Or, if you will, the hard-to-read font had the effect of "breaking the rhythm that excludes thinking.")

There are other fascinating cases that might be part of the same phenomenon. One of these is a study by Boaz Keysar and colleagues (2012) in which bilingual subjects were confronted with what's known as "the Asian disease problem," first studied by Amos Tversky and Daniel Kahneman (1974):

> A new disease, recently emerged in Asia, has begun to spread. In the U.S., without medicine, 600,000 people will die from it. Two new medicines have been developed, but only one of them can be manufactured and distributed. You must choose which medicine to use.
>
> If you choose medicine A, 200,000 people will be saved.
>
> If you choose medicine B, there is a 33.33% chance that 600,000 people will be saved, and a 66.66% chance that no one will be saved.
>
> Which medicine do you choose?

The interesting finding from this experiment is that people's choices tend to depend heavily on the wording of the problem, a phenomenon known as the **framing effect**. When the problem is worded as it is above, people are more likely to choose medicine A. But they switch and are more likely to choose medicine B if they are informed that the outcome of medicine A is that *400,000 people will die*, but that with medicine B there is a *33.33% chance that no one will die and a 66.66% chance that 600,000 people will die*.

The careful reader will have noticed that this scenario is exactly the same as the earlier one, so it seems somewhat irrational that people would have such a different response to it. Keysar and colleagues found that their bilingual subjects showed the usual framing effects if they heard the problem in their native language. But when the problem was presented in their second language—which required more effort to process—the bias introduced by the wording vanished, and their subjects showed a more "rational" pattern of responding to both versions in the same way. Again, one possible interpretation is that the extra processing challenge of reading the problem in a foreign language triggered a more thoughtful processing mode.

framing effect A phenomenon in which decisions or preferences regarding two identical outcomes are observed to be dramatically different, depending on the wording of the outcomes.

LANGUAGE AT LARGE 8.2 *(continued)*

Whether or not it leads to deeper thinking, there's no doubt that reading Henry James feels more strenuous than reading the easy, breezy style of much contemporary popular writing. And this sense of effort can be leveraged for literary effect. One exceptional example can be found in "The Depressed Person," a short story by David Foster Wallace (2007). I'll leave it to you to identify the various structures and long dependencies that strain the parser. Add to that the actual content, and the effect is a passage that feels as exhausting and overwhelming to read as depression is itself:

> *The friends whom the depressed person reached out to for support and tried to open up to and share at least some contextual shape of her unceasing psychic agony and feelings of isolation with numbered around half a dozen and underwent a certain amount of rotation. The depressed person's therapist—who had earned both a terminal graduate degree and a medical degree, and who was the self-professed exponent of a school of therapy which stressed the cultivation and regular use of a supportive peer community in any endogenously depressed adult's journey toward self-healing—referred to these friends as the depressed person's Support System. … The excruciating feelings of shame and inadequacy which the depressed person experienced about calling supportive members of her Support System long-distance late at night and burdening them with her clumsy attempts to articulate at least the overall context of her emotional agony were an issue on which the depressed person and her therapist were doing a great deal of work in their time together.*

8.6 Variable Minds

Individual differences in memory span

So far, I've been using expressions like "*the* parser" and "*the* language processing system" as if sentence comprehension involved something like a software program, and each one of us had exactly the same copy of it installed in our heads. Many language researchers will admit that when you look more closely at individuals' data, you might see quite a bit of variation in how different people respond—for example, some people might display enormous garden path effects for certain sentences, while others might show subtle or nonexistent slowdowns. But until fairly recently, this kind of variability was treated as unsystematic experimental "noise," not as an integral part of building theories of language comprehension. Building and testing theories has typically focused on the commonalities that can be found across large numbers of language users, not their differences.

But differences among individuals are becoming a very interesting part of the language-processing story. It's obvious that there are some meaningful cognitive differences across individuals. We all know people who have phenomenal memories, and others who are hopelessly absent-minded; some have an extremely analytical style of thinking, while other are more creative and impulsive in their thought processes. These are stable differences that characterize people over long spans of time. Far from being "noise" or transient variations that might be there today but gone tomorrow, they seem to play an important part in organizing people's experiences and the ways in which they process information. So we'd like to know whether different cognitive profiles also have an impact on how people organize and process linguistic information.

The connection between individual cognitive profiles and language processing was first made with regard to claims about memory. As we've seen in the previous section, in order to explain why some sentences are hard to understand, a number of researchers have pointed the finger at working memory limitations. And that's interesting, because it seems that some people are blessed with longer memory spans than others, and that these are stable indi-

reading span test A behavioral test intended to measure an individual's verbal working memory. The test involves having the individual read a sequence of sentences while holding the last word of each sentence in memory. The number of words successfully remembered corresponds to that individual's memory span.

vidual traits that can be measured by standard tests. So, if the difficulty inherent in certain kinds of sentences comes from the fact that these sentences bump up against limits on memory capacity, then we should be able to predict, based on someone's working memory span, whether their interpretation of such sentences is especially prone to crashing—and maybe even whether they experience the prose of Henry James as pleasantly stimulating or as a depressing slog.

In an important paper, Marcel Just and Patricia Carpenter (1992) looked at several aspects of language processing in subjects with high memory spans versus subjects with lower spans. They measured memory span using a particular test known as the **reading span test** (developed by Meredith Daneman and Patricia Carpenter, 1980). This task, which is intended to mimic the memory pressures on the parsing system, requires subjects to keep certain words active in memory while reading and understanding text. You try it—read the following sentences, and as you do, be sure to remember the last word in each sentence:

> Despite their persistent disagreements, they managed to agree how to best educate their child.
>
> At last, she spotted the rowboat coming across the bay, tossed about on the tall waves.
>
> Long periods of unstructured reading and thinking sometimes lead to the most fertile ideas.

Quiz time: Without looking back, can you remember the last word of every sentence? If you could, the reading span test would continue, inflicting upon you higher and higher numbers of sentences until you failed to recall all of their last words. This breaking point would reflect the limits of your memory span on this particular test.

Just and Carpenter argued that there are several consequences to having a lower memory span. The most obvious one is that it becomes very hard to understand sentences that make heavy demands on verbal working memory. They reported evidence showing that low-span subjects were especially slow at reading object relative clauses (which have been argued by some researchers to be real memory-hogging sentences) compared to the less taxing subject relative clauses. On the other hand, the higher-span subjects experienced a much more subtle slowdown for sentences with object relative clauses, presumably because their extra processing capacity could accommodate these structures fairly easily.

It's also possible that low-span subjects approach ambiguity resolution in a different way than higher-span subjects do. In Section 8.3, I summarized a variety of sources of information that might be helpful in the disambiguation process. Just and Carpenter argued that people with higher memory spans are better able to juggle all these sources of information at once, whereas people with lower memory spans might have to constrain the amount of information they take into consideration while resolving ambiguities. For instance, they might not as readily take into account the semantic plausibility of competing interpretations, and they might not be able to hold multiple interpretations open for long periods of time. The end result might be that they commit early to one single interpretation, based on a very limited amount of information—in other words, low-span subjects might show exactly the kind of processing style that's been described by the garden path model. Thus, according to Just and Carpenter, the garden path model had it partly right: Resource limitations *do* result in a parser that builds a single structure without considering all the options—but this isn't built into the *architecture*, or mind design, of the parser. It's just a

by-product of trying to understand sentences using a narrower memory span. Those who have the luxury of spending mental resources on a parsing style that considers more information take advantage of the opportunity to do so.

You'll remember, though, that not all researchers agree that object relative clauses are difficult because they stretch memory resources too thin. The experience-based argument was that these sentences are hard simply because they don't occur very frequently, hence they defy the word-by-word predictions that parsers make over the course of a sentence as it unspools. Notice that in this debate, individual differences become an important source of evidence that can shed some light on very general questions about the nature of parsing. Namely, if it turns out that we can systematically predict how hard it is for individuals to parse object relative clauses by looking at their memory spans, then this looks like pretty strong support for the view that such sentences do in fact draw heavily on memory resources, and that this is what makes them potentially difficult. In other words, the different parsing outcomes of different individuals can tell us something about the language processing system.

How have researchers from the experience-based camp responded to the evidence linking memory span with processing difficulties? By casting doubt that the reading span test is a valid measure of pure memory in the first place. Maryellen MacDonald and Morten Christiansen (2002) have argued that some people do better on the reading span test not because they have a more spacious memory, but because they are better *readers*. Better readers, they claim, are better able to read and understand sentences efficiently—they can do so while consuming less in the way of processing resources, which simply makes them *appear* to have more processing capacity. And how do people get to be better readers? Generally by spending more time immersed in written language than their peers. Think about any type of athletic training, say, training to run a marathon. We're used to thinking of training as having the result of extending an athlete's capacity (her strength or endurance). But one of the most important aspects of training is that the athlete learns to use her body more efficiently to achieve the same movements. Someone who runs 5 or 10 miles a day certainly has more strength and endurance than I do—but if both of us go for a run, it's also true that I will simply be working harder to run at the same speed. According to Christiansen and MacDonald, high performance on the reading span test merely reflects the fact that the so-called high-span subjects don't have to work as hard to process the same sentences. The memory test, then, is just a stand-in for the amount of reading "training" a person has had. And why would those with more experience at reading have an easier time with unusual structures like object relative clauses? Possibly because they've encountered these unusual structures many more times than people who read less (especially given that written language is far more likely to use unusual, infrequent structures than spoken language). Under this view, the way to get better at coping with the intricate syntax of Henry James is not to do mental calisthenics that extend memory capacity—rather, it's simply to read more.

Ultimately, this debate about individual differences and working memory hinges on being able to answer two yet-to-be-resolved queries: (1) How much does the reading span test really capture about memory capacity (and how does it relate to other, possibly "purer" tests of memory span, and do these also correlate with difficulties in processing complex syntax)? (2) To what extent does experience with language affect performance on both memory tests and tests of language processing?

WEB ACTIVITY 8.6

Test yourself on various memory tasks To what extent do you think these tests do or don't tap into the same kinds of memory resources that are needed to process complicated sentence structures?

cognitive control Also known as **executive function**. The goal-directed cognitive processes responsible for directing attention and supervising behavioral responses to stimuli.

Stroop test Behavioral test in which subjects are required to name the color of the font that a word appears in while ignoring the (possibly conflicting) meaning of the word.

Differences in cognitive control

More recently, researchers have become interested in another source of individual variation and its relation to language processing, namely, the ability to ignore and suppress irrelevant information. So far, the picture that I've painted of the language-processing landscape is one of rampant, cutthroat competition among linguistic representations, all clamoring for attention. Often the difference between successful understanding and utter cognitive mush comes down to being able to focus on the right representations while ignoring the others. This seems to be part of a general skill called **cognitive control** (also referred to as **executive function**)—you might think of cognitive control as the managerial aspect of cognition, directing attention depending on specific mental goals and supervising behavioral responses to stimuli. You need it for much more than just language, and without it, it would be impossible to function competently in a complex environment; you wouldn't be able to do things such as drive, follow a conversation in a crowded room, avoid swearing at your boss, or do any one of a thousand daily things that require you to ignore information or impulses that get in the way of your specific goals. As you might have noticed, the capacity for cognitive control seems to vary quite a bit among individuals. Some are able to push away intruding information or impulses very effectively, while others are more likely to be lured by them.

At one extreme of the spectrum lie those individuals who, because of a stroke or brain injury, have clearly identifiable brain damage to the prefrontal cortex area of the brain, an area that seems to be responsible for exercising cognitive control. Patients with damage to one specific area of the brain's frontal lobes—the left inferior frontal gyrus (LIFG)—have a hard time resolving conflict between competing representations. One way that this shows up is in their performance on a **Stroop test**, where they are told to name the color of the font that a word appears in while ignoring the meaning of the word. For example, subjects would have to identify the font color for words that appear like this: BLUE, CHAIR, YELLOW, and so on. Needless to say, it's not hard to name the font color for BLUE or even CHAIR. But to give the correct answer "red" for the word YELLOW means having to suppress the interfering meaning of the word. People are slower and make more mistakes when the font color and meaning mismatch in this way, and patients with damage to the LIFG are especially slow and error prone.

WEB ACTIVITY 8.7

Measuring cognitive control A number of different tests can be used to measure cognitive control. This activity provides a demonstration of several of them.

Language processing is often a lot like a Stroop test. Think about what you have to do to understand a garden path sentence, for example, or the meaning of the word *ball* in a sentence like *She decided not to attend the ball*. In both cases, *irrelevant* meanings are activated, and have to ultimately be squashed. There's a growing body of research that suggests all of these examples are related to cognitive control abilities. For instance, patients with LIFG damage have a hard time settling on the less frequent meaning of an ambiguous word, or recovering from garden path sentences. Among a normal population, brain-imaging studies have shown that the LIFG is especially active in just such situations, when people have to override a strong bias toward an interpretation that turns out to be wrong. And again, among people who are considered to be in the normal range, the degree to which it's hard to retrieve the subordinate meaning of a word or to recover from a garden path sentence seems to be systematically related to some measures of cognitive control (for a summary of these findings, see Novick, Trueswell & Thompson-Schill, 2005). All of this points to there being some consistent aspect of a person's cognitive profile that affects the ways

in which she experiences language and its tremendous potential for ambiguity. To some extent at least, a person's degree of cognitive control is shaped by experience—an intriguing line of research suggests that bilingual and multilingual people develop this skill more robustly than monolinguals (see Box 8.5).

As luck would have it, human brains go through a good deal of their lives being somewhat short of managerial personnel. The prefrontal cortex is one of the last of the brain regions to fully mature, which explains why teenagers who are

BOX 8.5
Bilingualism and cognitive control

In Chapter 7, you learned that people who know more than one language mix them together in their minds, rather than storing and accessing each of them separately. This intermixing has some dramatic consequences for language processing. In Box 7.3 you saw that Russian-English bilinguals experienced competition from Russian-sounding words that sounded similar to the English words that they were hearing. This suggests that, because bilinguals carry around an extra vocabulary, they regularly experience more competition than monolinguals. Increased competition can make language processing slower and harder, and there's evidence that bilinguals show less efficient word recognition compared with monolinguals (Rogers et al., 2006). They also experience interference across languages in speaking as well as understanding language, and this is reflected in poorer performance in picture-naming tasks, even when the bilinguals are speaking in their first and dominant tongue (Ivanova & Costa, 2008).

But increasingly, researchers are finding an upside to bilingualism. It appears that as a result of the regular exercise of wrestling with a greater degree of competition, bilinguals build up more muscular cognitive control abilities. These benefits can be seen throughout the life span.

At the *very* young end of the life span, Ágnes Kovács and Jacques Mehler (2009) tested 7-month-old babies who were being raised in monolingual and bilingual households. They used a variant of a classic cognitive test known as the "A-not-B test," in which infants see an object being hidden in one location, and then watch as it's moved to a different location. In order to find the object in its new location, babies have to suppress their knowledge of the first location. This is a very basic cognitive control task that babies only reliably accomplish by the age of 18 months. Kovács and Mehler used a simplified version of the task in which children were first trained to look at one side of a screen, and then had to learn to look at the other side before being rewarded with a nifty visual treat. The babies from bilingual households were better at redirecting their gaze to the new location, while the monolingual babies were more likely to get stuck in the pattern of looking at the old location. This is quite a dramatic demonstration, as babies at this age have yet to utter their first words!

Preschoolers and school-age children who are bilingual continue to show enhanced performance on tests of cognitive control, as do adults at all ages. In fact, the bilingual advantage seems to *increase* in magnitude as adults advance into old age—perhaps this is a reflection of the fact that, just as it's one of the latest systems to mature, cognitive control is also one of the first to decline as we age. (Note to you in your 20s or 30s: be sure to make good use of your peak performance years!)

Ellen Bialystok and her colleagues (2004) measured the performance of adults age 30 to 80 on a cognitive interference task known as the "Simon task." In this task, a red or green square pops up on the screen, and people have to press a button on the right to indicate a green square, or a button on the left to indicate a red square. The control condition of this test involves no spatial interference; the square simply shows up in the neutral location at the center of the screen. But in the "Simon" condition, the squares appear on either the left or right side of the screen. Sometimes the square is on the side that's congruent with the correct button press (a red square appears on the left, requiring the subject to press the left button), while at other times the square appears on the side that's incongruent with the button press (a red square appears on the right side, requiring the subject to press the left button). In the incongruent trials, subjects have to ignore the square's location in order to avoid its interference with their response. People usually respond more slowly to the incongruent trials than to the congruent ones, and the difference between them is taken as a measure of the degree of interference they experience.

BOX 8.5 *(continued)*

Figure 8.8 presents the responses of subjects grouped by age. **Figure 8.8A** shows the control condition in which there's no interference; you can see that the subjects show overall slower response times beginning in their 50s, but that the performance of bilingual and monolingual subjects is identical at all ages. But a different picture emerges from **Figure 8.8B**, which plots the degree of interference on the Simon trials. Here, the bilinguals outperform the monolinguals at all ages, and the difference becomes more pronounced as the subjects enter their 60s (and continues to grow).

All of this seems to have some important clinical implications for people suffering from Alzheimer's disease. Along with memory loss, people with Alzheimer's experience a dramatic loss of cognitive control. But bilingualism appears to offer some protection against the effects of the disease. A number of studies across quite different populations (e.g., Alladi et al., 2013; Bialystok et al., 2007) have shown that bilingualism delays the progression of symptoms, so even when they're at a relatively advanced stage of the disease, bilingual patients typically function at a higher level than their monolingual counterparts.

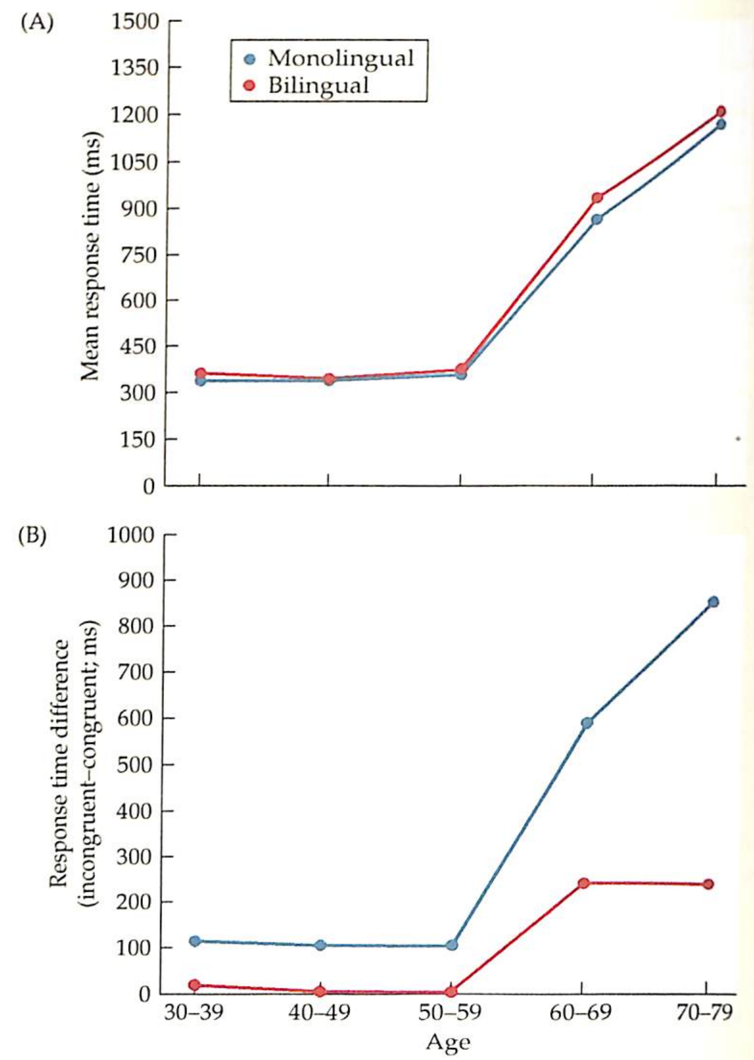

Figure 8.8 Mean response times for subjects by age. (A) Response times are plotted for the control condition in which there's no interference. (B) The degree of interference on Simon trials, calculated by subtracting the response times for congruent trials from the (longer) response times for incongruent trials. (Adapted from Bialystok et al., 2004.)

perfectly capable of mastering calculus or higher-order logic can nonetheless torment their parents with their inability to suppress certain impulses or inhibitions. These facts of neuroanatomy have been used by experts to argue that criminal courts should not apply the same sentencing criteria to adolescents as they do to adults. Now, it won't land them in jail, but you might expect that youngsters would also experience some trouble with linguistic ambiguity, and they do. John Trueswell and his colleagues (1999) tracked the eye movements of 5-year-olds listening to garden path sentences such as, "Put the frog on the napkin into the box." There's a temptation to treat *on the napkin* as attaching to the verb, and hence understand it as the intended destination for the frog, rather than as modifying the noun (as in *the frog that's on the napkin*; see **Figure 8.9**).

When confronted with such sentences, adults usually experience a brief garden path effect, as evident from their eye movements, but then quickly recover and perform the correct action. But in what the researchers cutely dubbed the "kindergarten-path effect," 5-year-olds experienced

WEB ACTIVITY 8.8

The kindergarten-path effect In this activity, you'll see some video recordings of subjects' eye movements and actions in response to ambiguity. You'll be able to see how adults and children differ in how they resolve confusing garden path sentences.

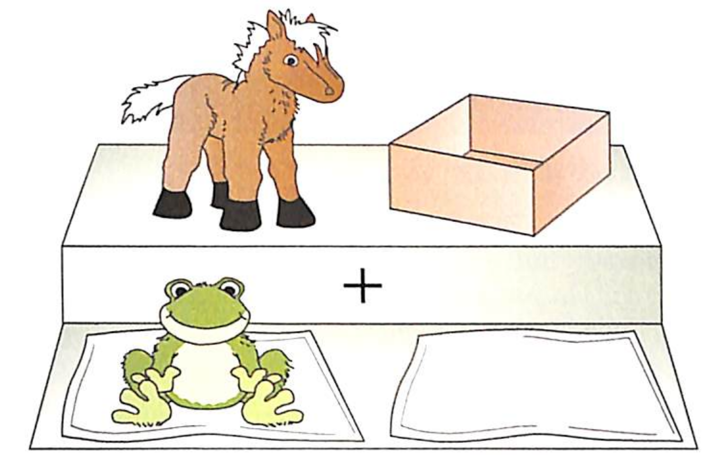

Figure 8.9 An example of a visual display from a 1999 eye-tracking study conducted by Trueswell et al. The subjects were children, who viewed the display while being given ambiguous and unambiguous versions of the instruction "Put the frog (that's) on the napkin in the box." With the ambiguous instruction, in addition to showing evidence of confusion in their eye movements, children often performed incorrect actions, such as hopping the frog onto the napkin and then into the box. This suggests that they were unable to completely suppress the incorrect interpretation of the ambiguous sentence. (Adapted from Trueswell et al., 1999.)

more devastating consequences of the ambiguity, performing the right action *less than half* of the time—for example, they might hop the frog onto the napkin, and then into the box, as if they were blending both interpretations rather than suppressing the incorrect one in favor of the right one.

In this chapter, I've pointed out how an eager, incremental style of language processing leads to tremendous potential for ambiguity—an ambiguity that usually creates fairly minor processing wrinkles for most adults. But for kids and their immature prefrontal cortexes, the processing traps laid by potential ambiguity aren't limited to exotic garden path sentences; such traps can also turn up in the mundane course of ordinary word recognition.

Remember that all words—even those that don't have a homophonous match—are temporarily ambiguous, so in the first few hundred milliseconds of the word *log*, both *log* and *lock* become activated (among others). In Chapter 7, we saw that evidence of such eager activation of word prospects can be seen in the pattern of subjects' eye movements to a visual scene—for example, when hearing *log*, people often look at a key because of its relationship to the cohort competitor *lock*. But for normal adults, such eye movements are usually extremely fleeting, quickly dropping down to baseline right after enough phonetic information has rolled in to identify the target word. And subjects are almost never aware of these spurious eye movements, or that they briefly flirted with interpreting *log* as *lock*. The competition seems to linger on a bit more for children: Yi Ting Huang and Jesse Snedeker (2010) found that in the same situation, 5-year-olds tended to persist in looking at the key for some time after hearing the disambiguating final consonant of the word *log*, and that they sometimes chose the key after being asked to "pick up the log." Perhaps once *lock* has been activated, and hence its related word *key*, the irrelevant representations continued to reverberate.

More relaxed cognitive control, though, may not be all bad. Some researchers (for example, Sharon Thompson-Schill and colleagues, 2009) have actually suggested that on balance, a slow-maturing prefrontal cortex may have evolved because of its overall cognitive advantages: Less cognitive control might allow children to learn language by easily generalizing rules to new situations. And less top-down cognitive management may well help anyone think more creatively. When it comes to language, a bit less cognitive discipline could even lead to a deeper enjoyment of the aesthetic potential of language—I'm thinking of Virginia Woolf from Chapter 7, and her inability to hear the words "Passing Russell Square" without hearing the rustling sound of long skirts on floors. Who knows? Perhaps her keen awareness of the rich ambiguities inherent in all words may have had its roots in an artistic temperament and its failure to squash down all sorts of irrelevant but wonderfully interesting information.

GO TO
sites.sinauer.com/languageinmind
for **web activities, further readings, research updates, new essays,** and other features

LANGUAGE AT LARGE 8.3

A psycholinguist walks into a bar...

Time for some jokes.

A man and a friend are playing golf at their local golf course. One of the guys is about to chip onto the green when he sees a long funeral procession on the road next to the course. He stops in mid-swing, takes off his golf cap, closes his eyes, and bows down in prayer. His friend says: "Wow, that is the most thoughtful and touching thing I have ever seen. You truly are a kind man." The man then replies: "Yeah, well we were married 35 years."

Two weasels are sitting on a barstool. One starts to insult the other one. He screams, "I slept with your mother!" The bar gets quiet as everyone listens to see what the other weasel will do. The first weasel yells again, "I SLEPT WITH YOUR MOTHER!" The other says, "Go home, Dad, you're drunk."

I want to die peacefully in my sleep like my grandfather. Not screaming in terror like his passengers.

A man put on a clean pair of socks every day of the week. By Friday he could hardly get his shoes on.

If you wanted to get analytical about humor (and why in the world wouldn't you?), you might notice that these jokes create the eerily familiar cognitive sensation of leading you down one train of thought for a while and then, at the punch line, yanking you into a completely different mental frame of reference so that you have to reinterpret the entire situation in an unexpected way. You might think of it as the garden path approach to humor, with the punch line serving as the "disambiguating region." Scholars who study humor often argue that such incongruous shifts are an important part of what makes something funny. The philosopher Immanuel Kant put it this way: "Laughter is an affection arising from a strained expectation being suddenly reduced to nothing." There's a striking parallel between the "strained expectation" you might feel in reading a garden path sentence, and the brief bewilderment that occurs before the *Aha!* moment in the joke.

For many jokes, the humor goes beyond merely being analogous to linguistic ambiguity—it *relies* on it:

A man walks into a restaurant and growls at the maitre d': "Do you serve crabs here?" The maitre d' responds, "We serve everyone here. Have a seat, sir."

What do you call a fly with no wings? A walk.

There are two fish in a tank. The first fish says to the second fish, "How the hell do we drive this thing?"

Why did Ludwig van Beethoven kill his two ducks? They wouldn't stop saying "Bach! Bach!"

Given that so many jokes hinge on recognizing and resolving conflicting interpretations, whether linguistic or otherwise, it wouldn't be surprising to find that there's some overlap in the neural machinery that's involved in interpreting jokes and resolving ambiguities. We've seen that the frontal lobes of the brain play an important role in refereeing between competing interpretations, with a starring role for the region known as the left inferior frontal gyrus (LIFG). A brain-imaging study by Tristan Bekinschtein and colleagues (2011) found that this region was more active for jokes that rely on competing meanings than for ones that rely on absurdist thoughts or images:

Why did Cleopatra bathe in milk? Because she couldn't find a cow tall enough for a shower.

But ambiguity alone isn't enough to tickle the funny bone. If it were, this would also be a joke:

What was the problem with the other coat? It was hard to put on with the paint roller.

Jokes rely on ambiguity, plus a little magic. With run-of-the-mill ambiguity, the goal of the language processing system is to identify the correct interpretation and quickly and efficiently squash any competing ones. The magic in jokes probably has to do with the fact that the whole point is *not* to squash the competing meaning, but to keep both interpretations alive in order to enjoy the tension between them—the more interesting the tension, the better the joke. (And it's precisely this tension that makes some unintended "crash blossoms" hilarious.) Becoming conscious of the reverberations between meanings is even more cognitively complex than basic ambiguity resolution. And indeed, the brain imaging study by Bekinschtein et al. found that the LIFG was especially active when people listened to jokes

LANGUAGE AT LARGE 8.3 *(continued)*

that played with ambiguity, when compared with non-funny sentences involving ambiguous words.

So there's a great deal of enjoyment to be had from the fact that our minds activate multiple meanings at once, whether you're a Virginia Woolf type who savors the useless multiplicity of the meanings of words, or just someone who likes to hear a good joke. If you need a good laugh, you might try visiting the LaughLab (http://richardwiseman.wordpress.com/books/psychology-of-humour/), a website in which psychologist Richard Wiseman documents his search for the funniest joke in the world. You'll find many jokes there that rely on language processing working the way it does. Among the submissions, you'll find one of my favorites. It hinges not on a garden path ambiguity, or even a pun due to outright lexical ambiguity, but on the more subtle lexical neighborhood activation effect, in which words that sound much like their targets are also accessed in the mind:

The monks were in the monastery copying those beautiful illuminated manuscripts. One young monk suggested that, since they'd been copying copies, it might be time to go back to the original and make sure that their copies were correct. The abbot agreed and sent the monk down into the cellar to examine the original. The monk was gone for a long time, and finally the abbot went to look for him. He found the monk in tears and asked what was wrong. Through his tears, the monk blurted out, "The word was *celebrate*!"

DIGGING DEEPER

Knowledge versus processing

One of the foundational ideas of modern linguistics is the distinction between *knowledge* of language and the *implementation* of this knowledge through processes whose end result is the actual comprehension or production of language. The idea is based on the very intuitive notion that we can have the knowledge required to do something, but all sorts of other factors in the end determine whether that knowledge is successfully carried out—presumably, you *know* how to multiply 24 and 37, but whether you manage to compute the answer in your head will be affected by a variety of things ranging from drunkenness to sleep deprivation, acute relationship troubles, a short memory span, someone yelling "952!" in your ear, and so on. In linguistics, *knowledge* of language—that is, the representations of words, the knowledge of sound patterns, understanding the options for how words can be formed from morphemes, grasping the rules of syntax, and all of that—is called *linguistic competence*. On the other hand, the mechanisms for recognizing or producing words, for parsing strings of words into meaningful structures, or figuring out, in the heat of the moment, how to assemble words into a sentence to express an idea, all fall into the domain of *linguistic performance*.

This separation has been implicit in all of our discussions on parsing so far—for instance, I've talked about how the rules of syntax are like a building code that *constrains* the construction of a sentence, but what actually gets built is determined by something else, namely "the parser." This seems to be a very useful distinction, because it allows us to explain why interpretation might fail even though a sentence seems to abide by everything we know about the patterns of our language. Consider, for example, the following sentences, all of which register as "bad" to the average English speaker:

The mouse the cat the dog bit chased ran.

The authors celebrated in the garden were drunk.

The three chairs is beside the table.

Her ate the cake.

The first example is an interpretation nightmare—most people just can't grasp what it means unless the structure of the sentence is made apparent (and it gets a lot easier if you transform the sentence to *The mouse that the cat that the dog bit chased ran*), in which case they might be able to squint and see that yes, it's a legitimate sentence after all. The second example is a garden-variety garden path sentence involving a reduced relative clause—with that knowledge, you should now be able to figure out what it means and agree that it's a legitimate sentence. On the other hand, the third and fourth examples are perfectly easy to understand, but they're just "wrong," and it's easy to point to what's wrong with them: in the third sentence, the subject and verb have to agree in number, so *is* should be *are*; in the fourth example, an object pronoun has been used instead of a subject pronoun, so *her* should be *she*. When foreigners are learning to speak a new language, we attribute many of their errors to a non-native *competence*, or a failure to have fully learned the patterns and structures of the language. The errors of native speakers though, whether slips of the tongue or lapses in understanding of convoluted but perfectly grammatical sentences, are normally said to be *performance* errors.

That seems reasonable enough, but the distinction gets murky pretty quickly. Take this common example:

Who did she meet?

Whom did she meet?

Is either of these "wrong"? It depends who(m) you ask! Until quite recently, the first sentence would have sounded just as bad to speakers of English as *Her ate the cake*, and for the same reason: namely, that the *wh*-pronoun bears the wrong *case marking*—it's marked as a subject form when it needs to be related to the object of *meet*. But languages evolve, and English is in the process of dropping the subject/object distinction on these *wh*- words. To many younger English speakers, *Whom did you meet?* sounds as archaic as *I bid thee good-bye*. In fact, many undergraduates I teach have so completely lost the subject/object distinction for *wh*- pronouns that they can't reliably tell which sounds more "right": *Whom did she meet?* Or *Whom ate the cake?*—a fact that might horrify some of their grandparents. (If you can't tell, either, I encourage you to admit this to someone a couple generations older than you and observe the effect.)

So, what are we to do with a situation in which English speakers disagree about which form is right? (And we're talking here about disagreement among monolinguals who have learned English from the day they were born, not foreign learners of English.) And, especially, what do we do with the speaker who thinks that either *who* or *whom* is perfectly good as a direct object of a verb? Is this a matter of competence? And if so, who are the fully competent speakers? Or should this be classified as a matter of performance—of having trouble understanding the sentences that are deemed bad? The answer to that depends on exactly why responses are varied. Is it because people's knowledge of these forms varies (especially as a function of age), or is it that younger and older people differ in their abilities to understand these *wh*- questions?

Most students I've asked say it's about knowledge. It seems obvious, they argue, that people will learn and prefer whichever form they're exposed to most often. Each subsequent generation has been using fewer *whom*s and more *who*s, so that eventually, children learning English from young parents will very rarely if ever encounter a *whom*. Since learning the syntactic patterns of language is sensitive to the statistical distributions of those patterns, it makes sense that those kids who've heard *whom* quite often will think of that form as more "right" than those who've never heard it. So it also seems to make sense that one's competence, or knowledge of language, shouldn't just have all-or-none rules, but that perhaps these rules should be graded in a way that reflects the amount of experience they've had with specific linguistic patterns.

Again, all this seems reasonable enough. But here's where the potential blurring of competence and performance comes in. Remember that when we talked about the longer reading times for object relatives as opposed to subject relatives, this was all done in the context of talking about object relatives causing *processing* difficulty. And that seems right, if the difficulty really does boil down to issues with working memory constraints, as suggested by a number of researchers. But one of the alternative explanations for the processing trouble was that object relatives are simply less frequent than subject relatives, and hence less expected by a parser whose job it is to predict the form of the unfolding speech, whether it's a single word or a complex sentence. This means that frequency is reflected *both* in knowledge of language (so less-frequent sentences can be more grammatical than highly frequent ones) and in processing efficiency (less-frequent sentences are more difficult to process than more-frequent ones).

You might see where this is going. In the eyes of some researchers, especially those who are immersed in experience-based theories, there are no meaningful differences at all between knowledge and processing. Knowledge of language simply boils down to the accumulation of those routines that have been used to interpret or produce forms. Rather than having a separate and static set of rules that you consult like a building code during the process of building, the "rules" of syntax are more like a series of paths you've had to travel in the past in order to understand language: the more heavily traveled paths are the "rules" that generate more frequent structures. And, just like a familiar path home, it will both be easier to travel quickly, and feel more "right" than an unfamiliar one.

Still, for many researchers, the distinction between competence and performance remains a very important one. This is especially so for those scholars who've argued that much of the knowledge of language is innate. These researchers very explicitly reject a purely experience-based view of linguistic knowledge; they claim that much of what kids know about language *couldn't* have been learned through experience without some extra set of constraints or defaults, simply because their accuracy with certain kinds of structures goes far beyond what they could possibly have

encountered as linguistic input (recall the arguments from the "poverty of the stimulus," as discussed in Section 6.4). So linguistic knowledge has to amount to more than just well-worn processing paths or routines.

As we've seen in Chapter 6, one of the strongest arguments that's made in support of the idea of innate linguistic knowledge is the fact that when you look across languages, there tend to be certain commonalities in structures or rules, so languages seem to be constrained in terms of the kinds of structures or patterns they allow. Small kids, the argument goes, show an uncanny ability to avoid just those structures that tend to be universally ruled out, presumably because they are innately guided toward jumping to the right conclusions about structure.

Let's go back to an example of a potentially universal constraint on the structure of *wh*- questions, which we first saw in Chapter 6. Remember that *wh*- words have to be linked to some position in the sentence, and that this relationship can span quite a large distance, and involve a variety of different syntactic slots:

What did Felicia claim that she saw her boyfriend eating spaghetti with __?

Who does Joey think he should be studying hard in order to impress__?

When did Britney say her lawyer would set up a hearing with the judge __?

But, as we saw, certain long-distance relationships seem to be forbidden even when they look superficially very similar to *wh*- relationships that never raise anyone's eyebrows. Compare, for example, the following set of contrasts (remember that the asterisk conventionally indicates unacceptability):

What do you think that Joey bought __?

* **What** do you wonder whether Joey bought __?

What did you claim that John bought __?

* **What** did you make the claim that John bought __?

What do you think that Josh left __ at my house ?

* **What** would you laugh if Josh left __ at my house?

Who did Betsy invite Silke to come with __?

* **Who** did Betsy invite Silke and __?

Who did Weyman believe __ had kicked Cecilia?

* **Who** did Weyman believe the rumor that __ had kicked Cecilia?

What did Stuart wonder if Myra made __?

* **What** did Stuart wonder when Myra made __?

These prohibitions on structure, often called *wh*- island constraints, are apparently echoed in many languages. And, as discussed in Chapter 6, kids seem to make surprisingly few errors with them, showing that they are aware of these restrictions on *wh*- questions. All of this is hard to explain under a purely experience-based account.

But there's another possible explanation that doesn't resort to saying that kids are genetically predetermined to favor certain *wh*- question structures over others. Some researchers have suggested that the ugliness of the above *wh*- questions for English speakers has nothing to do with speakers' knowledge or *competence*, but rather, that this is a matter of *performance*: such sentences, they claim, create difficulties for the parsing system. Hence, we should think of them more along the lines of horrendous sentences like *The horse raced past the barn fell*, or *The mouse the cat the dog bit chased ran*. Specifically, Robert Kluender and Marta Kutas (1993), along with Philip Hofmeister and Ivan Sag (2010) have proposed that *wh*- island constraints reflect limitations on working memory during sentence processing. If that's true, the mystery about the apparent universality of *wh*- question structures evaporates very quickly. And for that matter, so does the mystery of how kids know to avoid certain *wh*- structures: it's simply that those structures are too hard to compute.

One hint that a performance-based account of *wh*- islands might be on the right track is the fact that people don't seem to be absolutely uniform in condemning all *wh*- island violations as unacceptable (you may have quite a lot of variation in judgments in your own class, for example). This is surprising if the sentences are ruled out by universal, genetically programmed constraints, but much less so if they are grammatically well-formed but simply create processing difficulty. After all, there appear to be some systematic individual differences that correlate with how much difficulty people experience on certain garden path sentences or long-distance dependencies. If one could find that certain traits that are *known* to be implicated in processing difficulty also predict the variation in whether *wh*- island violations sound "bad," this would deflate one of the significant arguments for innate universal language-learning constraints.

Such evidence has been hard to find. In one paper, Jon Sprouse and colleagues (2012) tested subjects on a number of working-memory tasks, and found no relationship at all between memory performance and the degree to which these subjects rejected *wh*- island violations as bad sentences. On the other hand, Hofmeister et al. (2012) have complained that not much can be concluded from the study by Sprouse and colleagues until they first establish that the specific memory tests that they chose *do* reliably predict acceptability judgments for sentences that are already known to be unacceptable to subjects for processing reasons—for example, sentences like *The mouse the cat the dog bit chased ran*. Without this preliminary control in place, it's hard to know whether their particular method really bears on the question of competence versus performance at all.

In any case, given the juicy implications of the issue, there is bound to be a good deal more heated debate about whether to classify these iconic *wh*- island violations as a matter of knowledge, or a matter of processing.

PROJECT

Aside from approaching the question from an individual differences perspective, can you think of another way to test whether *wh*- island violations reflect competence or performance? Propose and, if possible, carry out an experiment to find out. Be sure to think about the appropriate control sentences you would want to use, and to establish that your method is capable of capturing the distinction.

9 Speaking: From Planning to Articulation

Speaking can get you into trouble. Ask a friend of mine, who is married to a woman named Carmen. A frosty weekend in the marital household once ensued after he slipped up during a conversation and addressed her as *Susan*—the name of his previous spouse. Oops. Carmen took this as an ominous sign that he was not yet over his ex-wife. To make matters worse, my friend's entire family, despite having been warned of the dire consequences, has managed to make exactly the same slip in poor Carmen's presence at one time or another.

It gets worse. You can lose your job for misspeaking. In 2006, radio announcer Dave Lenihan was fired for dropping a racial slur while discussing the prospect of Condoleezza Rice as commissioner for the National Football League. Lenihan said, "She's got the patent résumé of somebody that has serious skill. She loves football, she's African-American, which would be kind of a big coon. A big coon. Oh my God—I totally, totally, totally, totally am sorry for that. I didn't mean that."

Listeners pounced, demanding his resignation. Lenihan claimed that he was aiming to articulate *coup* and mispronounced the word. No matter. In the ensuing kerfuffle, many commentators expressed the belief that uttering the word (whether intentionally or not) reflected racist attitudes (whether consciously held or not) on the part of Dave Lenihan. The radio station manager agreed, and within 20 minutes he was on the air to announce that Lenihan was being tossed from the show. The manager said he believed the offending word was an inadvertent slip of the tongue, but was nonetheless "unacceptable, reprehensible, and unforgivable."