Psychological Review
2003, Vol. 110, No. 3, 490–525


0033-295X/03/$12.00 DOI: 10.1037/0033-295X.110.3.490

# Meaning Through Syntax: Language Comprehension and the Reduced Relative Clause Construction

Gail McKoon and Roger Ratcliff
Ohio State University

A new explanation is proposed for a long standing question in psycholinguistics: Why are some reduced relative clauses so difficult to comprehend? It is proposed that the meanings of some verbs like *race* are incompatible with the meaning of the reduced relative clause and that this incompatibility makes sentences like *The horse raced past the barn fell* unacceptable. In support of their hypotheses, the authors show that reduced relatives of *The horse raced past the barn fell* type occur in naturally produced sentences with a near-zero probability, whereas reduced relatives with other verbs occur with a probability of about 1 in 20. The authors also support the hypotheses with a number of psycholinguistic experiments and corpus studies.

Bever's (1970) reduced relative clause sentence *The horse raced past the barn fell* presents a puzzle that has driven a large portion of language processing research for the past 30 years. On first encounter, most English speakers judge the sentence to be ungrammatical. Yet—so it has been claimed—they all come to agree that their first judgment was wrong, that the sentence is grammatical. The puzzle has been why, if the sentence is grammatical, it is so difficult to process. In this article, we take a different tack. We argue that first judgments are correct, that *The horse raced past the barn fell* is not an acceptable English sentence.

In most current views of language comprehension, the goal of research designed to understand sentence processing is a syntactic goal: to model how all the words of a sentence, the nouns, verbs, subjects, objects, and so on, are put together by the processing system into the syntactic structure of the sentence. In contrast, we propose a new view of language comprehension that we label *meaning through syntax* (MTS). The goal of MTS is semantic: to understand how syntactic structures are determined by meaning and how they express meaning.

In this article, we show the value of the MTS view in providing significant new insights into the puzzle of the reduced relative clause, a domain of investigation chosen because of its centrality in the history of psycholinguistics. We make two specific MTS proposals, one for verbs and the other for the reduced relative construction. For verbs, we propose that the parts of verbs' meanings relevant to syntactic structures can be expressed as "event templates." *Break*, for example, has the event template *x* CAUSE (*y*(BECOME BROKEN)), a semantic structure showing the meaningful relations CAUSE and BECOME that connect the entities *x* and *y* into an event of breaking. The template also simultaneously represents syntactic information in that its elements can be linked to syntactic positions in sentences. For example, for *John broke a window*, the causer, *x*, is linked to subject position and the entity that becomes broken, *y*, is linked to direct object position.

For constructions, MTS borrows from current research in construction grammar the notion that syntactic constructions in and of themselves carry meaning, a different meaning for each construction. A passive sentence, for example, is said to convey a different meaning than an active sentence. Pertinent to the research described here, we propose that a reduced relative clause construction of the type *The horse raced past the barn* has its own special meaning, different from the meanings of other syntactic constructions such as transitives, passives, or nonreduced relative clauses.

In the first part of this article, we describe in detail the theoretical claims of MTS: the event template meanings proposed for verbs and the construction meaning proposed for reduced relative clauses. We give linguistic background for each proposal, and then we show how they combine to make *The horse raced past the barn* ungrammatical as a reduced relative clause. In a nutshell, the thesis we defend in this article is that *The horse raced past the barn fell* is ungrammatical because the event template meaning of the verb *race* places the horse in an event structure position that is incompatible with the construction meaning of the reduced relative clause. In the second part of the article, we support this thesis, as well as the proposed event template and reduced relative meanings, with a series of psycholinguistic experiments on sentence comprehension and memory and with corpus data on sentence production.

In the third part of the article, we take up the question of how the MTS view relates to other current views of sentence processing, for reduced relative clauses in particular. One model, Stevenson and Merlo's (1997) parsing model, is similar to the MTS view in suggesting that reduced relatives of the type *The horse raced past the barn* cannot be successfully parsed. With respect to other models, the MTS view is in conflict with one of the major classes, "constraint-based models." In these models, the language processing system is assumed to store and use statistical information about the frequencies with which various linguistic structures occur in

Gail McKoon and Roger Ratcliff, Department of Psychology, Ohio State University.

The theoretical background for this research owes much to Beth Levin; we are grateful for her patient answers to our many questions. We are also grateful to Talke Macfarland, who has added considerable linguistic insight to the ideas presented here.

Correspondence concerning this article should be addressed to Gail McKoon, Department of Psychology, Ohio State University, 216 Lazenby Hall, 1827 Neil Avenue, Columbus, Ohio 43210.

the language. In contrast, by the MTS view, the language processing system knows meanings and rules, but not statistics of usage. Because of this sharp disagreement, the last studies we present test predictions of constraint-based models. The results of the studies are consistent with MTS but not with most current constraint-based models.

The MTS view, at this time, is made up of proposals about representations of information that are important to language processing. It is far from a model of how the representations are used in processing because it generates such a large number of empirical hypotheses about representations that need to be addressed first. The results of the studies presented in this article only begin to tap what needs to be done empirically. We believe the great potential of the MTS view is demonstrated by the fact that neither the variables manipulated in the studies presented here nor the results of the studies would have been anticipated by other current views of the language processing system.

## CORPUS DATA

Throughout this article, we rely heavily on corpus data to support hypotheses generated from MTS. The corpus we use is made up of 180 million words of naturally produced English text: 93 million words of articles from the *New York Times* (*NYT*; 1984–1990), 7.6 million words from fiction and nonfiction adult books, 1.5 million words from fiction and nonfiction children's books, 1 million words of the Brown corpus, 31 million words of transcribed conversations from the "MacNeil–Lehrer News Hour" and "Larry King Live," 6 million words of the *Wall Street Journal* (*WSJ*), and 41 million words of *Time* and *Newsweek* magazines.

The corpus provides sentences that are naturally produced by speakers and writers of English. It thus provides a window to linguistic representations. The data base is large enough to provide thousands of examples of use of a particular verb or construction. For example, there are 4,294 sentences in the corpus that use the verb *race*. Analyzing such large amounts of data is slow and painstaking but, as becomes apparent below, the payoffs in terms of theoretical leverage are great.

Throughout the article, we give many sentences from the corpus as examples of construction or verb usage. Some individual sentences were chosen from the corpus because they make a special point. However, the statistics we report as the results of corpus studies are based either on all tokens in the corpus of the particular verb or construction of interest or on random samples of tokens.

## THE MEANING OF THE REDUCED RELATIVE CLAUSE CONSTRUCTION

A crucial hypothesis of MTS is that syntactic constructions carry meaning independently of the particular words that are expressed in them. We know, for example, that the transitive structure *noun1 verb noun2* means that something or someone (noun1) engages in some activity that in some way affects or involves some other entity (noun2). There are numerous other, more subtle, construction meanings. Fillmore (1968) pointed out that *Bees are swarming in the garden* can mean that they are swarming only in a small part of the garden or that they are swarming in the whole garden, whereas *The garden is swarming with bees* means that they are swarming in the whole garden. Anderson (1971) pointed out that *I loaded the truck with hay* conveys that the truck is completely full of hay, but *I loaded hay onto the truck* does not (and McKoon, Ward, Ratcliff, & Sproat, 1993, showed the psychological consequences of this difference).

The claim that syntactic constructions carry meaning independently of the words expressed in them has a long history in linguistics. Numerous authors (Anderson, 1971; Bolinger, 1968; Borkin, 1985; Fillmore, 1968; Givon, 1985; Goldberg, 1995, 1997; Wierzbicka, 1988) have discussed a wide variety of examples of the differences in meaning that occur when the same words are used in slightly different constructions like the *garden swarming with bees* and *loading the truck with hay* sentences just mentioned. Bolinger (1968) has concluded that "a difference in syntactic form always spells a difference in meaning" (p. 127).

The force of the argument for the independence of construction meaning from word meaning lies in demonstrations that construction meanings often cannot be built up from individual word meanings. To illustrate, consider sentences like these (from Goldberg, 1995, 1997):

*Sally baked her sister a cake.*

*Frank sneezed the tissue off the table.*

*The train screeched into the station.*

*Pat smiled her appreciation.*

In each case, the event denoted by the sentence has meaning that is contributed by none of the words individually. The word *bake*, for example, does not by itself embody the meaning of "cooking something for the purpose of transferring possession of it to someone else"; according to standard dictionaries, *bake* means only "cooking in an oven" or "becoming dry as the result of heat." Likewise for the other sentences: To convey the meaning of the *sneezing the tissue* sentence, the meaning of *sneeze* would have to include "cause to move"; to convey the meaning of the *screech* sentence, the meaning of *screech* would have to include concurrent movement of the entity emitting the screech; and to convey the meaning of the *smile* sentence, the meaning of *smile* would have to include "conveying something by means of smiling." For none of these verbs are the relevant senses included in their dictionary entries, nor is it reasonable that they should be. Instead, so the argument goes, these special senses of the verbs are contributed by the constructions in which the verbs appear. The ditransitive construction, used in the *bake* sentence, means some entity (*Sally*) causes some other entity (*her sister*) to receive something. The caused movement construction, used in the *sneeze* sentence, means that some entity (*Frank*) caused some other entity (*the tissue*) to change location. The *screech* sentence makes use of the concurrent motion construction, and the *smile* sentence makes use of the transitive construction to express something abstract (in this case, "appreciation"). All of these constructions are extremely productive in English; all are used with a large variety of verbs to denote a large variety of events.

The MTS hypothesis we put forward is that reduced relative clause constructions are constructions, like those just mentioned, that convey special meaning. In this article, we focus on "object" reduced relatives, that is, reduced relatives of the type *The horse raced past the barn* in which the head of the clause (*the horse*) is the direct object of the verb in the clause. We imagine that the "subject" reduced relative construction, in which the head is the subject of the clausal verb (e.g., *the horse racing past the barn* . . .) also has meaning, but what that might be is outside the scope of this article. It is object reduced relatives that have been the focus of much psycholinguistic research.

For the object reduced relative construction, we propose that it is used to denote a discourse entity characterized by a particular function. The function is defined in the following way: participation in an event caused by some force or entity external to itself. In other words, a discourse entity denoted by an object reduced relative clause could be represented as $f(x)$, where $f(x) = x$ participates in an event with cause external to $x$.

The relevant causality is that specified by the event template of the verb used in denoting the event, as is discussed in more detail below. For instance, for the reduced relative clause *Cars and trucks abandoned in a terrifying scramble to safety*, the verb used in denoting the event is *abandoned*, and the entity denoted by the reduced relative is "cars and trucks caused by some external force or entity to be abandoned in a terrifying scramble to safety."

There are three critical elements of our hypothesis about the object reduced relative clause construction. First, like the constructions studied by Goldberg (1995, 1997) and others, the construction contributes meaning that is independent of the meanings of the individual words. Specifically, the meaning of the construction is to denote a discourse entity in a particular way. Other syntactic structures in which the same words appear do not have the same meaning. For example, we take the meaning of the reduced relative *cars and trucks abandoned in a terrifying scramble to safety* to be the denotation of one discourse entity ("cars and trucks caused by some external force to be abandoned in a terrifying scramble to safety"). We hypothesize that the meaning of *Cars and trucks were abandoned in a terrifying scramble to safety* is different, at least because it denotes a discourse entity "cars and trucks" separate from the information that they were abandoned in a terrifying scramble to safety, and it probably differs in other ways as well.

Second, and especially relevant to the history of psycholinguistic study of reduced relatives, is the requirement that $x$'s participation in an object reduced relative event be externally caused, where the relevant causality is defined by the event template of the verb used to denote the event. An event that is denoted as internally caused by $x$ cannot appear in the object reduced relative construction, as the construction is defined here; such an event would be inconsistent with the meaning of the construction.

A verb that denotes an internally caused event ascribes responsibility for the event to the entity engaging in it (Levin & Rappaport Hovav, 1995; McKoon & Macfarland, 2000b; Smith, 1970). *Bloom* is a prototypical example—the cause of a flower blooming is internal to the flower. As we shall argue later, *race* is also a verb of internal causality; when a horse engages in the activity of racing, the horse, ultimately, controls whether it races. A more complete background of the linguistic and psycholinguistic motivations for the distinction between external and internal cause verbs is given in the next section. For now, the important point is that the meaning proposed for the object reduced relative construction does not allow events for which the verb denotes control residing in $x$, the entity that is the head of the clause. The implication is that the sentence *The horse raced past the barn fell* is an incorrect use of the reduced relative construction. So long as both the syntactically relevant meanings embodied for verbs in event templates and construction meanings are considered to be part of the grammar, then the claim is that *The horse raced past the barn fell* is ungrammatical.

This claim is similar in an important respect to a proposal made by Stevenson and Merlo (1997). In their syntactic parsing model, the lexical structures of verbs reflect elements of the verbs' meanings. For a motion verb like *race*, its lexical structure does not provide the syntactic elements that are necessary to construct a syntactic structure for a reduced relative clause in which the verb appears. Thus their parser cannot produce a complete syntactic analysis for a reduced relative clause with a verb like *race*. To the extent that this inability is taken to mean that *The horse raced past the barn fell* is ungrammatical, then their proposal and ours are similar (although differing in other important ways, as discussed later).

The third important aspect of defining the object reduced relative as a meaningful construction follows from the notion that different constructions (different syntactic forms, in Bolinger's, 1968, terminology) have different meanings. From this point of view, there is no reason to think that nonreduced relative clauses have the same meaning as reduced relative clauses, and there is at least some reason to think that they have different meanings—if they didn't have different functional meanings, why would they both exist in the language? There may be multiple ways in which the meanings of reduced and nonreduced relative clauses are different. The focus of this article is reduced relatives, not nonreduced relatives, and so only an initial and tentative hypothesis is offered about nonreduced relatives.

The hypothesis we offer about nonreduced relative clauses is concerned only with restrictive clauses because those are the structures that have been used in psycholinguistic research for comparison to reduced relative clauses. Restrictive relative clauses are relative clauses that define or limit the entities that are their heads, and they are not delimited by commas or introduced by *which* (only by *that* or *who*; see, e.g., Strunk & White, 1979; Wilson, 1993). Throughout this article, whenever we refer to either reduced or nonreduced relative clauses, we mean restrictive relative clauses.

Given the data presented later, it seems reasonable to suppose that whereas an object reduced relative places a single entry into the discourse ($f(x)$), an object nonreduced relative places two entries into the discourse: an entity $x$ and a fact about $x$, namely the information in the relative clause (which we label $g(x)$). To capture this hypothesized (but perhaps incompletely specified) difference in meaning, an object nonreduced relative clause can be represented as the following:

$$x \; g(x),$$

where $x$ and $g(x)$ are two discourse entries, $x$ corresponding to the head of the clause and $g(x)$ to the information in the clause that tells something about $x$. The consequence of this representation that we take up later in this article is that $x$ and $g(x)$ are separately represented, compared with the representation of a reduced relative, $f(x)$. That is, the entity $x$ and the information about it, $g(x)$, are more easily accessible separately than if they were bound together in a reduced relative.

The MTS notion that reduced and nonreduced relative clauses have different meanings brings up a crucial contrast between the pretheoretic view espoused here and the pretheoretic view that has been commonplace in other studies of relative clauses. Here, the a priori assumption is that a difference in syntactic structure is likely to correspond to a difference in meaning. Hence, *cars and trucks abandoned in a terrifying scramble to safety* and *cars and trucks that were abandoned in a terrifying scramble to safety* are assumed, until it has been shown otherwise, to have different meanings. In contrast, the standard view has been that these two clauses have the same meaning; they differ only in syntactic form.[1]

To sum up this section, the hypothesis about object reduced relative clauses is straightforward: The object reduced relative clause is a construction with its own particular meaning. It places into the discourse an entity defined as having participated in an externally caused event, where the relevant causality is defined by the event template of the verb in the clause. Later, we present corpus studies and psycholinguistic studies to support this hypothesis. Here, it is perhaps worth noting, in passing, what an elegant construction the reduced relative is—paraphrasing its meaning takes many more words than does the construction itself. It fulfills an essential discourse function with great succinctness. It packages for quick and concise entry into the discourse an entity that would otherwise have to be defined in multiple complex ways.

For the remainder of this article, we omit the word *object* from our references to object reduced relative clauses and object nonreduced relative clauses. As has been indicated, all the hypotheses and data that are presented in the article concern object reduced and object nonreduced relatives, never subject reduced or subject nonreduced relatives.

## VERB MEANINGS AND EVENT TEMPLATES

The meaning of the reduced relative clause construction is the first piece of the MTS view about reduced relative clauses. An important component of the reduced relative meaning is the external causality requirement. The notions of external and internal causality, defined with respect to verbs' lexical semantic event templates, are the second relevant piece of the MTS view. In the sections that follow, these notions and their mappings onto various classes of verbs are placed in the context of current lexical semantic theory and are more fully articulated.

### Background

Since the 1980s, the assumption guiding many theories in linguistics has been that a verb's lexical entry contains much, perhaps all, of the information that determines the syntactic structures in which the verb appears. In other words, the assumption is that the verb's meaning determines (at least, in large part) its syntactic behavior.

In early work, verbal lexical entries contained three kinds of information: a verb's subcategorization frames, its argument structures, and all of the rest of its meaning (see, e.g., Jackendoff, 1972). Information relevant to syntax, what we call "syntactically relevant" information, was represented in the subcategorization frames and the argument structures. Argument structures laid out the semantic elements of the syntactic structures in which a verb could be used and subcategorization frames showed the syntactic placements of the elements into sentences. For example, the argument structure of the verb *put* might have been something like (agent theme location) where the agent is the entity causing the event of putting, the theme is the entity being put, and the location is where it is put. The subcategorization frame would be (NP1) ((V) (NP2) (P NP3)), where NP1 would take the subject position in a sentence, V is *put*, NP2 would take the direct object position, and P heads a prepositional phrase for which NP3 would take the prepositional object position. Many verbs would have multiple argument structures and subcategorization frames, different ones for the different syntactic structures in which they could appear.

In the view just described, syntactic information (in subcategorization frames) was represented separately from semantic information (in argument structures). In later work (e.g., Chomsky, 1986), subcategorization frames were eliminated and a verb's lexical entry is now usually said to contain only semantic information, the semantic information relevant to syntax plus all the other parts of the verb's meaning. The idea is that, because sentence structure "is largely predictable from the semantics of predicates" (Wasow, 1985, p. 202; see also Bates & Goodman, 1997; Grimshaw, 1994; Levin & Rappaport Hovav, 1995; Moens & Steedman, 1988; Pustejovsky, 1991; Tenny, 1994; van Valin & LaPolla, 1997; van Valin & Wilkins, 1996), subcategorization frames are redundant. Currently, one of the important questions for research in lexical semantics is whether sentence structure is entirely predictable from semantic information. The research strategy used here and by a number of linguists (e.g., Rappaport Hovav & Levin, 1998; Tenny, 1994) is to assume that it is, pushing meaning to predict as much of syntax as it possibly can.

With the elimination of subcategorization frames, the question became how to represent the syntactically relevant part of a verb's meaning such that it could explicitly determine syntactic structure. Answers to this question have taken two general forms: The earlier approach used argument role lists like (agent theme location) for *put*; then more recently, when it became clear that role lists were unworkable (as we outline below), predicate decomposition approaches were developed. Argument role lists have been used extensively in psycholinguistic research, so we discuss them briefly here. For MTS, we have adopted a predicate decomposition approach.

In order for argument role lists to represent verbs' syntactically relevant meanings, there are three requirements: a list of the roles, definitions of each of the possible roles, and rules for mapping the roles to syntactic positions. Role lists have run into trouble on all three of these fronts. First, it appears to be impossible to rigorously list the possible roles and define them (e.g., Croft, 1998; Levin & Rappaport Hovav, 1996). Levin (1999) gave examples of the multitude of possibilities for a single noun, *bridge*, as the object of a variety of verbs. For some cases, one or more role labels have been suggested in the literature: *build a bridge* (effected object); *destroy a bridge* (patient or consumed object); *widen a bridge* (patient or incremental theme); *move a bridge* (theme); *wash a bridge* (location or surface); *cross a bridge* (path); *leave a bridge*

[1] We refer to the proposed difference between reduced and nonreduced relatives clauses as a difference in meaning. It might also be described as a difference in discourse function. However, in this article, we follow the literature on text representation (e.g., Kintsch, 1974; McKoon & Ratcliff, 1992) and make no distinction between the types of meanings expressed by individual words and propositions and the types of meanings expressed by larger text units. For present purposes, nothing crucial hinges on the descriptive label applied to the proposed reduced–nonreduced difference.

Table 1
*Event Templates*

| Event template name | Template | Examples |
|---|---|---|
| Activity | *x*(ACT) | *x*(RACE), *x*(HIT) |
| State | *x*(IN STATE) | *x*(LOVE), *x*(EXIST) |
| Simple change of state | *x*(BECOME IN STATE) | *x*(BECOME ARRIVED), *x*(BECOME BLOOMED) |
| Externally caused change of state | $\alpha$ CAUSE (*x*(BECOME IN STATE)) | $\alpha$ CAUSE (*x*(BECOME BROKEN)), $\alpha$ CAUSE (*x*(BECOME FADED)) |

(source); *reach a bridge* (goal); and *see a bridge* (stimulus or object of perception). For others, no particular label suggests itself: *praise a bridge*, *touch a bridge*, *avoid a bridge*, and *imagine a bridge*. In the lexical semantics literature, no single list of roles and their definitions has ever been agreed upon.

The situation is just as problematic when it comes to rules for mapping role labels to syntactic structures. Generally, mapping was done by ordering role labels into hierarchies such that for the roles of a particular verb, the one highest in the hierarchy might be mapped to subject of a sentence, the next highest to direct object, and so on. The problem has been that many different hierarchies have been suggested, varying in their choices of roles and their orderings of roles. Summarizing how hierarchies have been used, Croft (1998) pointed out that there is no underlying motivation for the ranking of roles because a hierarchy of isolated roles "does not take account of their relational nature" (p. 29).

In response to these problems, lexical semanticists have turned to predicate decompositions. Different forms of decompositions have been proposed by a number of theorists, including Hale and Keyser (1993), Rappaport Hovav and Levin (1998), Stevenson and Merlo (1997), Tenny (1994), Pustejovsky (1991), and van Valin and LaPolla (1997). The predicate decomposition approach is relatively new and it is not yet clear whether it can be used to solve all of the problems that plagued role lists, but efforts to date have been encouraging and the data provided later in this article offer further support. The event template decompositions used with the MTS view are most like Levin and Rappaport Hovav's (1995, 1996) decompositions. Table 1 lays out the templates. We first describe them and then elaborate the assumptions behind them.

The templates in Table 1 show decompositional representations of syntactically relevant meanings. The first template, for activity verbs, shows the entity, *x* (the "argument"), that participates in the event and, with the semantic primitive ACT, it shows the activity in which the entity engages. This template represents the syntactically relevant meaning of verbs like *race* and *hit*. The manner of motion verbs that have been important in the study of reduced relative clauses (e.g., *race*, *walk*) are a subclass of activity verbs. They have the *x*(ACT) template, and also they are verbs of internal causality, that is, they carry the specification that the denoted activity is under the control of the entity engaging in the activity. The second template, for states like *love* and *exist*, has the structure that *x* exists in some state. The third template is for simple change of state verbs like *bloom*, *wilt*, and *arrive*. The semantic primitive BECOME specifies a change for the entity *x* to a new state. One subclass of these verbs is the "internally caused change of state" class, verbs like *bloom* and *ferment*. Like the manner of motion verbs, these verbs carry a specification of internal causality: The denoted change of state is under the control of the entity undergoing the change of state. The fourth template is for verbs that are called "externally caused change of state" verbs, verbs like *break*. Their template shows two subevents, a causing event $\alpha$ and a change of state event, connected by the primitive CAUSE. The change of state subevent is the same as the simple change of state (*x*(BECOME IN STATE)). *Alpha* can be any of the four event types, a state, an activity, a simple change of state, or an externally caused change of state (van Valin & LaPolla, 1997).

The four predicate decompositions shown in Table 1 are taken to define four types of events that verbs may denote. A template lays out the structure of an event. It provides a skeleton of meaning upon which a sentence can be built (Tenny, 1994). It shows the placement of arguments in events and subevents, whether and how arguments are affected in the course of an event, and how subevents are connected to each other by causality. States and activities place arguments in situations or events, simple and externally caused changes of state place arguments in events in which the arguments undergo a change, and externally caused changes of state place arguments in two causally related subevents. For some classes and subclasses of verbs, arguments in their template are also marked for causality: For example, manner of motion and internally caused change of state verbs specify causality internal to the entity engaging in the verbal event, and externally caused change of state verbs specify causality external to the entity changing state. For any particular verb, its template is an instantiation of the verb into the appropriate template type, as is shown with the examples in the table, for instance, *x*(HIT) and *x*(BECOME ARRIVED). For all the verbs that share the same template, the assumption is that the syntactically relevant part of their meanings is the same. They differ in all the other parts of their meanings.[2]

In the sections that follow, we further describe and motivate the use of event templates and then present the hypothesized interactions of the templates with the reduced relative clause construction.

[2] Unlike some other lexical decomposition approaches, we do not define the four proposed types of events in aspectual terms. Aspectual classifications sort verbs according to their inherent temporal properties: telic–atelic, punctual–durative, and stative–dynamic. Aspectual properties do not pick out the same classes of verbs as the event templates we define. Also, verbs can switch aspectual classes according to the context in which they are used, whereas we propose that verbs maintain the same basic event template across contexts.

## Event Templates: What They Are and What They Are Not

The syntactically relevant part of a verb's meaning is often a direct and obvious reflection of the other parts of its meaning. For example, *run* has an *x*(ACT) template because it is an activity. However, which template is appropriate for a verb is not always obvious. Levin and Rappaport Hovav (1995) give the verb *blush* as an example. In English and Dutch, *blush* has an activity template. That is, it is construed as an activity verb, as evidenced by the fact that we can say someone blushed *for 10 minutes*, but not someone blushed *in a minute* (meaning it took a minute for the person to change from the nonblushing to the blushing state). In Italian, it appears that *blush* has a different template, a simple change of state template, because the use of *blush* is exactly reversed: Blushing for 10 minutes is an unacceptable use of the verb and blushing in a minute is acceptable (see Levin & Rappaport Hovav, 1995, for further discussion).

The *blush* example illustrates a second point about event templates that cannot be overstressed: Event templates are not intended to directly represent events in the world. Instead they express construals (i.e., interpretations) of events in the world. Apparently, in Italian the verb *blush* denotes a change of state, whereas in English and Dutch, it denotes an activity.

The arguments in event templates (e.g., the *x* in *x*(ACT)) are simply place holders to indicate participants in the events; there is nothing semantic about them in themselves. The meaning of a template lies in the number of arguments it contains, in the template structure that defines the relations among them, and in whether the template is marked for internal or external causality. The primitive relational predicates are tentatively hypothesized to comprise only a small set, CAUSE, BECOME, and ACT. The number of different template types is also thought to be small, perhaps comprising only the four types shown in Table 1: activity, state, simple change of state, and externally caused change of state.

The entities that are represented in a template are only those that are required for syntactic expression. The verb *erode*, for example, has the template *x*(BECOME ERODED). *Erode* is usually used in an intransitive construction, specifying only the single argument *x* (e.g., *The beach eroded*). However, it can also be used in a transitive construction that specifies a second participant in the erosion process, as in *The waves eroded the beach*. But this second participant is not syntactically required and so is not represented in the template (McKoon & Macfarland, 2000b). It is assumed that optional arguments like this are optionally added to a sentence from the parts of *erode*'s meaning that are not specified in the event template.

This is one point on which the MTS view contrasts sharply with the thematic role list approach used in some current psycholinguistic research. Typically, in a role list approach, a verb's lexical entry would have encoded with it all the different possible role lists with which the verb could occur. For example, for *erode*, there would be one role list for sentences like *The beach eroded* and a different role list for sentences like *The waves eroded the beach*. In the MTS approach, the lexical entry for a verb has only a single event template (unless the verb is truly ambiguous), to which optional arguments can be added.

## Linking Rules

The projection of event templates to sentence structures is implemented by linking rules. Here, we illustrate linking rules with two easily explained rules suggested by Levin and Rappaport Hovav (1995). One is the "Immediate Cause" linking rule, which states that the argument denoting the immediate cause of an event denoted by a verb will take the subject position in a sentence. For example, in *John broke a window*, the immediate cause of the breaking event is something John did ($\alpha$) and so $\alpha$ is placed in subject position (where it is referenced by *John*). The rule applies to internally as well as externally caused change of state verbs. For *bloom*, for example, the immediate cause of blooming is something internal to the entity that blooms and so this entity takes subject position.

The second linking rule is the "Directed Change" rule. According to this rule, the argument of a template that corresponds to the entity undergoing a directed change takes direct object position. In *John broke the window*, it is the window that is undergoing a directed change and so it is the direct object. "Directed" change means a change that results in the change of state that is dictated by the verb. *John* probably undergoes some change in the event of breaking a window, but it is not the change dictated by the verb. Likewise, for *John runs*, *John* likely undergoes some change in the course of the activity, but it is not a change that has direction—*John runs* does not necessarily indicate that he actually went anywhere by running.

What is important to note about the linking rules is that they are formulated in terms of the semantic event templates and this means, in turn, that the syntactic positions in a sentence are defined, at least in major part, in terms of the semantic event templates. This is the essence of the MTS view: The event template meaning of the verb is expressed through the syntax of the sentence. For example, for *John broke the window*, *John* being the subject of the sentence conveys that something John did is the immediate cause of the breaking event and *the window* being the direct object conveys that the window changes state as a result of the event.

## Intuitions and Generalizations

In linguistics, one goal for event templates, and predicate decompositions in general, is to contribute to an explanation of what language users know about syntax–semantics interactions and how they can know it. Language users have strong intuitions about how individual verbs can be used and these intuitions generalize across classes of verbs.

To better explain what verb templates are designed to accomplish, we give as an example the contrasts between two classes of verbs, "*wipe*" verbs and "*remove*" verbs. Consider the prototypical verbs of these classes: Dictionaries typically define *wipe* as a kind of movement that removes some substance from a surface and *remove* as taking some substance away from where it was. The two definitions are much alike, and that is in accord with first intuitions about their meanings by users of the two verbs. But users of these verbs have more, quite subtle, intuitions about the differing linguistic environments in which the two verbs can be acceptably used. Tokens of *wipe* from our corpus include the following:

1. To wipe some object: *Kate Johnson, as a less than*

*innocent servant, strokes and wipes the shin of her mistress, Cathy McCann.*

2. To wipe some object or substance from some place: *The Treasury wipes the obligation from its books . . .*

3. To wipe some object or substance off some place: . . . *the White House wiped its critics off page one of all the newspapers.*

4. To wipe away some object or substance: *The volunteers were here to wipe away any fingermarks clouding the memorial's polished black granite slabs.*

5. To wipe some object into some condition: *Continental Bank's officials say their approach would wipe the slate clean . . .*

6. To wipe some object along some path: . . . *men insisting on wiping a rag across motorists' windshields congregate at stoplights all across the city . . .*

7. To wipe some substance onto a surface: *You can . . . then wipe a light coat of stain over the whole area to help conceal the patch.*

What differences are there between *wipe* and *remove*? Of the seven different constructions above, *remove* can appear in the first and the second and perhaps the third and fourth, but not the others. One cannot, for example, remove a slate clean or remove a rag across a windshield, and *remove* cannot be used to add a stain or any other substance onto a surface.

Not only are the differences between *wipe* and *remove* considerable, the same differences are found between other verbs of their classes. *Sweep*, *dust*, *erase*, *rinse*, *rub*, *scrub*, *strip*, and *wash* behave like *wipe*, whereas *delete*, *dislodge*, *dismiss*, *eject*, *evict*, *omit*, and *sever* behave like *remove* (see Levin, 1993).

If the constructions in which a verb can appear are determined, in significant part, by the features of verbs' meanings that are captured in event templates, then differences in the possible constructions for the verbs of the two classes should be explainable from their templates. This is, in fact, the case.

The *wipe* verbs are assumed to have a simple activity template, *x*(ACT), and the *remove* verbs to have an externally caused change of state template, $\alpha$ CAUSE (*x*(BECOME IN STATE)) (see Rappaport Hovav & Levin, 1998). Intuitively, *wipe* verbs are used to construe events as being about the manner in which an event is carried out, whereas *remove* verbs are used to construe events as being about the result or end state of the event. With the *x*(WIPE) template, nothing beyond the activity itself is required in a sentence. A location for the activity (*her mistress's shin*), a substance to be removed (*the Treasury's obligation*), or a path (*across windshields*) can be added, but these are not required by the template and they come from the parts of the verb's meaning that are not expressed in the template. Also, changes of location and changes of state can be added. In Examples 2, 3, and 4, the Treasury obligation, the White House critics, and the fingermarks all change location, and in Example 5, the slate becomes clean. Perhaps the most compelling evidence for the simple *x*(WIPE) template is that *wipe* can actually be used with the opposite of a *remove* meaning: In Example 7, the change of state is that stain is added, not removed.

For *remove*, on the other hand, its template, $\alpha$ CAUSE (*x*(BECOME REMOVED)), already expresses a change of state and another cannot be added (because an event is prohibited from having more than one change of state; see Tenny, 1994). For example, it cannot be said that *the slate was removed clean.* Constructions like those in Example 2 above and perhaps those in Examples 3 and 4 are possible because removing something "away" or removing something "off" of some surface provides further specification of the removal change of state, not addition of a different change of state. Also, because the template specifies the "removed" change of state, the verb can only be used for taking a substance or object away, not for adding it.

The *wipe–remove* examples illustrate the sorts of intuitions and generalizations across classes of verbs that can be explained by event templates as representations of syntactically relevant meaning. Below, we show similar sorts of regularities for the classes of verbs relevant to reduced relative clauses.

## VERB TEMPLATES AND THE REDUCED RELATIVE CLAUSE CONSTRUCTION

At this point, we have introduced the meaning of the reduced relative construction and the event template meanings of verbs. The crux of the MTS view of reduced relatives is that a reduced relative is grammatical only if the template meaning of the verb in the clause places the entity that is head of the clause in an event structure position that is compatible with the reduced relative's construction meaning. The reduced relative construction is defined as denoting an entity that participates in an externally caused event. In the reduced relative *snowflakes sent by the good witch*, the head, *snowflakes*, participates in an event of sending. The sending event is controlled by the witch and so control is external to the snowflakes. According to MTS, the reduced relative construction allows only verbs that provide a template structure in which causality resides in some entity other than the one that is the head of the reduced relative.

There are two classes of verbs that we identify in this article as meeting the requirement for insertion into the reduced relative construction and two that we identify as failing to meet the requirement. The verbs' interactions with the reduced relative construction are outlined in this section and then are later fleshed out with supporting data.

The first column of Table 2 shows templates for verbs that meet the reduced relative requirement. The entity that would be positioned as the head of a reduced relative is designated *Y*. Verbs like *break*, externally caused change of state verbs, denote events in which the change of state for *Y* is externally caused, so they meet the reduced relative's requirement of external causality and insertion into a reduced relative is permitted. In both *John broke the window* and *the window broken by John*, *Y* is *the window* and $\alpha$ is represented by *John*. Verbs like *examine* are activity verbs. For a verb of this class, the entity *Y* is not represented in the event template; instead, it comes from the other parts of the verb's meaning (and so we have designated it with a subscript in Table 2). When an entity *x* engages in the verbal activity, *x* involves *Y* in that activity. For *Y*, the event is under external control and therefore insertion into the reduced relative is permitted. Thus, for both these

Table 2
*Verb Classes and Reduced Relative Clauses*

| Reduced relative clauses allowed | Reduced relative clauses not allowed |
|---|---|
| *Y* changes state—External cause | *Y* engages in activity—Internal cause |
| $\alpha$ CAUSE (*Y*(BECOME IN STATE)) | *Y*(ACT) |
| $\alpha$-John broke *Y*-the window. | John raced *Y*-the horse |
| *Y*-The window broken by $\alpha$-John . . . | **Y*-The horse raced by John . . . |
| *Y* participates in event—External cause | *Y* changes state—Internal cause |
| $x(\mathrm{ACT}_Y)$ | $Y$ (BECOME IN $\mathrm{STATE}_x$) |
| *x*-John examined *Y*-the window | *x*-The water rusted *Y*-the pipes. |
| *Y*-The window examined by *x*-John . . . | **Y*-The pipes rusted by *x*-the water . . . |

*Note.* Asterisks indicate unacceptable constructions.

classes of verbs, their event templates provide a structure in which the event denoted in the reduced relative is controlled by some entity other than the entity positioned as the head of the reduced relative.

The second column of Table 2 shows how verbs' event templates can block insertion into the reduced relative construction. The two classes of verbs shown are manner of motion activity verbs (e.g., *race*, *swagger*, *jump*) and internally caused change of state verbs (e.g., *bloom*, *rust*, *wilt*).

For manner of motion verbs, their template is the simple activity template, with causality marked as internal to the entity engaging in the activity. The MTS claim is that this is the single, basic template for these verbs. Although some verbs of this class can be used in transitive constructions such as *John raced the horse*, this is considered to be an extension beyond their basic sense (Levin & Rappaport Hovav, 1995; Stevenson & Merlo, 1997). If *Y* is the entity engaging in and controlling the activity, then a reduced relative with *Y* as its head is blocked because the external causality requirement of the reduced relative construction is not met. For the verb *race*, for example, its basic template is *Y*(RACE). Racing is internally controlled by *Y*, the entity engaging in the activity, not externally controlled, and so *race* cannot be used in a reduced relative clause with *Y* as its head.

It should be stressed that it is a verb's template that defines the causality relevant to reduced relatives. Through its template, a verb places a particular construal on an event. *Race* places a construal of internal causality on the events it denotes, whatever the entities doing the racing, whether they are horses on racetracks or cars on highways or sticks racing over a waterfall. Whatever the entity engaging in a racing event, a reduced relative with that entity as head is prohibited.

The same considerations apply to internally caused change of state verbs, although their template structures events differently. Internally caused change of state verbs have a change of state template for which control of the change of state is internal to the entity undergoing the change. These verbs can be used in transitive constructions, but differently from verbs of other classes (McKoon & Macfarland, 2000b). For internally caused change of state verbs, the subject of the verb in a transitive construction is not an external controller but instead a participant in the change of state event. This participant is an argument that comes from the parts of the verb's meaning not represented in the template (and so it is designated with a subscript in Table 2). To the reduced relative clause, the meaning that matters is the meaning in the event template and according to the template, the change of state for *Y* is internally controlled; therefore, use in a reduced relative with *Y* as its head is blocked.

The event templates in Table 2 and their interactions with the reduced relative construction summarize the MTS hypotheses about reduced relative clauses. We turn now to supporting data.

## DATA: LINGUISTIC KNOWLEDGE, SENTENCE PRODUCTION, AND SENTENCE COMPREHENSION

For the MTS view of reduced relatives, there are three main pieces: The meanings of verbs as represented in event templates, the meaning of the reduced relative construction, and the interaction of the two kinds of meaning. In the sections that follow, we present empirical evidence to support all three.

The first data to be presented support the proposed event templates. Of the four classes of verbs shown in Table 2, for only two of them, manner of motion and internally caused change of state, is their template crucial to our thesis because only verbs of these classes invoke prohibition of reduced relatives. For manner of motion verbs, we present one corpus study and three psycholinguistic experiments, and for internally caused change of state verbs, we review one corpus study and four experiments from earlier work. Second, we present evidence supporting the MTS definition of the reduced relative clause, and third, we present evidence of the prohibition of reduced relative clauses for verbs of internal causality.

### Event Templates: Manner of Motion Verbs, Corpus Study 1

The hypothesis is that a reduced relative cannot have as its head an entity engaging in a manner of motion event because a manner of motion event is internally controlled by the entity engaging in it, and such internal control is prohibited for the reduced relative. The claim is that the activity template is the single basic template in the lexicon for these verbs. They do not have a second basic template; in particular, they do not have a basic template for external control of the activity.

The activity template (*x*(ACT)) covers all the events denoted by manner of motion verbs for which control is unequivocally internal to the entity engaging in the activity. This includes instances in which the verbs are used intransitively, such as the examples below for the manner of motion verbs *scamper*, *frolic*, *travel*, *wander*, *slither*, and *limp*:

> *. . . children scamper while their elders frolic.*
>
> *I wanna quit, I wanna get out, I wanna travel, dream, wander!*
>
> *Mr. Guillemette slithered and preened in the company of a glowing neon tube . . .*
>
> *. . . the summer solstice arrived just before Isaiah Thomas, both of them limping.*

The claim that needs to be supported is that the activity template is the only basic lexical template, that there is no second basic template for manner of motion verbs denoting an event for which control is construed to be taken over by an external force. For example, in the sentence . . . *the great Stonewall galloped Little Sorrel up and down this very road*, it appears that Stonewall has taken over control of the galloping event from Little Sorrel. If there were an external causality template, as a basic lexical template, for instances like these, then MTS could not explain the difficulty presented by manner of motion verbs occurring in reduced relatives. The template would incorporate the notion of external control that is hypothesized to be required by the reduced relative construction.

We argue that external control instances of manner of motion verbs do not reflect a basic lexical template. Instead, we argue that these instances are extensions from the basic activity template (Levin & Rappaport Hovav, 1995; Stevenson & Merlo, 1997).

The first line of evidence, Corpus Study 1,[3] comes from examination of the kinds of sentences people produce with manner of motion verbs. We examined all the tokens in the corpus for 69 verbs, all the verbs classified as manner of motion verbs by Levin (1993) that appear in the corpus mainly as verbs (e.g., *backpack* appears mainly as a noun and so it was excluded) and do not have a second meaning that is difficult to distinguish from the manner of motion meaning (e.g., *bounce*). For the 31 verbs *amble*, *clamber*, *climb*, *flit*, *frolic*, *hike*, *hop*, *limp*, *lope*, *meander*, *prowl*, *roam*, *romp*, *saunter*, *scamper*, *scram*, *scramble*, *shamble*, *skulk*, *slink*, *stray*, *swagger*, *tiptoe*, *traipse*, *tramp*, *travel*, *trek*, *trudge*, *waddle*, *wade*, and *wander*, an examination of 56,460 corpus sentences found no instances at all of external control. For the 20 verbs *canter*, *carom*, *cavort*, *coast*, *crawl*, *creep*, *leap*, *lurch*, *plod*, *prance*, *ramble*, *scurry*, *sidle*, *slither*, *stroll*, *stumble*, *swim*, *toddle*, *totter*, and *troop*, there are 5 or fewer instances, a total of 33 instances out of a total of 25,748 sentences. Examples include the following:

> *. . . she tottered him home after one of his binges . . .*
>
> *Now, shut up and crawl your white-ass back over here . . .*
>
> *The black and white transcript shows a President who . . . is rambling his answers . . .*
>
> *. . . she had some rather unusual things to say about CNN, and we . . . very quickly scurried her off.*
>
> *On Saturday Fred . . . zoomed the leaves through his hedge onto Murphy's property . . .*
>
> *Jackson distracts attention . . . by slithering her hands over her torso in erotic confusion.*

For *drift* (3,700 sentences), there are 17 instances; for *glide* (1,339 sentences), 10 instances; for *jump* (16,482 sentences), 24 instances; and for *zoom* (740 sentences), 9 instances.

Only for the 14 verbs *dart*, *float*, *gallop*, *hurry*, *hustle*, *march*, *parade*, *race*, *sneak*, *speed*, *strut*, *trot*, *vault*, and *walk* are there many external cause uses. For both these verbs and those listed above that have less frequent external cause uses, there is a tight restriction on the entities that exercise external control: They are people, groups of people, actions by people, or natural forces (Cruse, 1972; Levin & Rappaport Hovav, 1995; Stevenson & Merlo, 1997). We assume this restriction comes about because internal causality can be overridden only by a force superior to that of the entity engaging in the verbal activity. For example, when "she" "totters him home," "she" is taken to be a superior force to someone who is drunk. For the 14 verbs for which external cause uses are frequent, we checked the restriction by examining a random sample of 12 corpus instances of external control for each verb and found none for which the external cause was not a person, group of people, action by people, or natural force.

Overall, data from Corpus Study 1 provide two lines of evidence that external cause uses of manner of motion verbs do not reflect a basic, external causality template: One is the rarity of external cause sentences for the class of manner of motion verbs, and the second is the restriction that the external controllers that do occur must represent forces superior to the entity engaging in the verbal activity. These data are noteworthy because the sentences in the corpus are sentences that are naturally produced by speakers and writers. The sentences show the knowledge that the speakers and writers have about their language, specifically that manner of motion verbs convey internal control. Speakers and writers follow the constraints imposed by internal control, overriding internal control rarely and only in well-circumscribed situations.

Before moving from sentence production data to comprehension data, we mention one further line of evidence that external cause uses of manner of motion verbs are derived extensions from a basic, internal cause template. If the external cause uses are derived, then it might be expected that they would be morphologically marked. Data for this claim come from the morphology of causatives in languages other than English, reviewed by Levin and Rappaport Hovav (1995). In a number of languages, the morpheme that is used to mark causative uses of manner of motion verbs is different from the morpheme that is used to mark standard causatives, that is, those of the externally caused change of state class. The suggestion is that the two kinds of causatives arise from the lexicon in different ways.

## *Comprehension of Manner of Motion Verbs: Experiments 1, 2, and 3*

The MTS proposal, supported by Corpus Study 1, is that the basic representation of syntactically relevant meaning for manner of motion verbs in the lexicon is the simple activity event template. If this is correct, and if the language system processes simpler

[3] For this corpus study (and only this one), a 280-million-word corpus was used instead of the 180-million-word corpus used for all the other reported studies. The 280-million-word corpus is a superset of the other corpus, and the additional material comes from the same types of sources.

structures faster than more complex ones, then manner of motion verbs should be processed faster than verbs that have external causality templates. We measured speed of processing with sentence comprehension and with lexical decision. We compared the manner of motion verbs against externally caused change of state verbs (e.g., *break*, *fade*) because these verbs have been shown to have $\alpha$ CAUSE($x$(BECOME IN STATE)) as their basic template (McKoon & Macfarland, 2000b; Rappaport Hovav & Levin, 2000).

In Experiment 1, manner of motion verbs and externally caused change of state verbs were tested in three-word intransitive sentences such as *The child crawled* (manner of motion) and *The signal faded* (externally caused change of state). Subjects were asked to read the sentences and the dependent measure was reading time. Because of the greater complexity of the event templates for the external cause verbs, we predicted that their sentences would have longer reading times than sentences with the simpler manner of motion verbs.

In Experiment 2, manner of motion and external cause verbs were tested in anomalous transitive sentences such as *Some captains creep sighs* (manner of motion) and *Some captains split sighs* (externally caused change of state). Subjects were asked to judge whether each sentence was acceptable or unacceptable. Again, it was predicted that the more complex external cause verbs would have longer response times than the simpler manner of motion verbs.

The methodologies of Experiments 1 and 2 complement each other. With acceptable sentences (Experiment 1), there could be a problem in that the sentences for the external cause verbs might differ on some unknown dimension from the sentences for the manner of motion verbs, a dimension other than the verbs' lexical semantic event templates (even though the sets of sentences were closely matched on as many dimensions as possible). Experiment 2 counters this possibility because exactly the same anomalous sentence frames were used for the manner of motion verbs as for the externally caused change of state verbs.

Experiment 3 measured lexical decision response times for the same verbs as were used in the sentences of Experiments 1 and 2. If event templates are an important part of verbs' lexical entries, then the more complex external cause verbs should have longer lexical decision response times than the simpler manner of motion verbs. The two sets of verbs were matched on other factors known to affect lexical decision response times so that response time differences could be attributed to the greater complexity of the external cause verbs' event templates.

## Method

*Materials.* Sixteen manner of motion verbs and 16 external cause verbs (listed by Levin, 1993) were paired such that the two verbs of a pair were matched on several dimensions: Kučera–Francis frequency (Francis & Kučera, 1982; means 34 and 30, respectively), number of syllables, subjective frequency (means 3.3 and 3.2, on a 1 to 5 scale, with 1 being most frequent; Forster, 2000), and imageability (means 2.1 and 2.0, on a 1 to 5 scale, with 1 being most imageable). The two sets of verbs were also matched on the numbers of senses listed for each verb in *WordNet* (Fellbaum, 1998), with the verb *break* deleted (*break* is listed as having 63 senses). The mean number of senses for the manner of motion verbs was 6.6 and the mean number of senses for the external cause verbs was 8.0; the means were not significantly different, $F_2(1, 30) < 1.0$, $p < .05$, throughout this article. The two sets of verbs were also matched on the sizes of their neighborhoods of orthographically similar other words, as measured by the number of words in Francis and Kučera (1982) that differ from them by exactly one letter. The mean numbers of such words were 2.7 and 3.6 for the manner of motion and external cause verbs, respectively—not significantly different, $F_2(1, 30) < 1.0$. The pairs of verbs were (first, the manner of motion verb, then the externally caused change of state verb): *amble*, *shrivel*; *crawl*, *crack*; *creep*, *split*; *float*, *fade*; *gallop*, *vibrate*; *glide*, *fray*; *limp*, *thaw*; *toddle*, *redden*; *stagger*, *stiffen*; *climb*, *freeze*; *scamper*, *mellow*; *trot*, *smash*; *walk*, *break*; *sneak*, *rip*; *wade*, *scorch*; and *rush*, *snap*.

For Experiment 1, one 3-word intransitive sentence was written for each verb. Some examples are as follows: *The cows ambled*, *The crops shriveled*, *The dancers glided*, *The wires frayed*, *The players limped*, and *The chicken thawed.* For all but five of the sentences, the first word was *the*, and for all the sentences, the second word was a noun and the third word was the verb, always in past tense. The noun chosen as the subject for each verb was a noun that is used as the verb's subject in an intransitive sentence in the corpus. Also, the noun that was chosen for the subject of each verb's sentence is an entity of the kind that is most frequently used as a subject of the verb in intransitive sentences in the corpus. For example, the most common entities used as subjects of intransitive sentences with the verb *ambled* are animate entities and so we used an animate entity as the subject of the sentence in the experiment. The sentences for the two verbs of a pair were matched in terms of the Kučera–Francis frequency of their nouns. The two sets of sentences were also matched in terms of plausibility. An independent group of 12 subjects rated the sentences for plausibility on a 1 to 5 scale (with 5 being highly plausible). The mean rating for the sentences with manner of motion verbs was 4.6, and the mean for the sentences with external cause verbs was 4.5. The two sets of sentences were also matched on how semantically related the noun and verb are. An independent group of 15 subjects rated each pair (e.g., *cows–amble*) on a scale of 1 to 5, with 5 being most highly related. The mean rating for the pairs with manner of motion verbs was 2.9 and the mean for the pairs with external cause verbs was 2.8. There were also filler sentences for the experiment, some intransitive and some transitive.

For Experiment 2, 16 pairs of sentence frames were created, each frame made up of four words, all of the form *Some nouns verb nouns*, for example, *Some captains verb sighs* and *Some editors verb moans*. The nouns of the frames were chosen to make the sentences clearly anomalous. For half of the subjects, the first frame of a pair of frames was used with the manner of motion verb of a pair of verbs and the second frame was used with the external cause verb, and for the other half of the subjects, this assignment was reversed. The nouns in the two frames of a pair were closely matched in Kučera–Francis frequency. The verb was always in present tense. There were also filler sentences for the experiment, some transitive and some intransitive, some acceptable and some unacceptable. The unacceptable sentences were all as anomalous as possible in order to keep response times fast and reading rates at a near-normal level.[4]

For Experiment 3, the 16 verbs of each type were tested in lexical decision, with filler words and word-like nonwords.

*Procedure and subjects.* For all the experiments reported in this article, the stimuli were presented on the screen of a personal computer (PC) and responses were collected on the computer's keyboard.

For each subject in Experiment 1 (measuring the reading times of intransitive sentences), there were eight lists of nine randomly ordered sentences each (preceded by one practice list), each list containing two of the sentences with manner of motion verbs and two of the sentences with external cause verbs, plus filler sentences. Subjects began each list by pressing the space bar on the keyboard. For each sentence, subjects were

[4] The materials used in all the experiments reported in this article are available from the authors.

instructed to read it and press the space bar when finished; then the screen was cleared and, after a 500-ms pause, the next sentence was presented. After the nine sentences of a list, there were three true–false test sentences to test comprehension of the sentences that had been read. Subjects averaged 87% correct on the true–false tests, indicating that they read the sentences reasonably carefully. There were 20 subjects. For all of the experiments reported in this article, the subjects were Northwestern undergraduates participating for course credit.

For each subject in Experiment 2 (measuring acceptability judgments of anomalous transitive sentences), there were eight lists of 16 randomly ordered sentences each (preceded by one practice list), 6 unacceptable and 10 acceptable sentences in each list and each list containing 2 of the sentences with manner of motion verbs and 2 of the sentences with external cause verbs, plus filler sentences. Subjects began each list by pressing the space bar, and then the sentences were presented one at a time. For each sentence, subjects were asked to indicate their best judgment of whether the sentence was acceptable or unacceptable, pressing the *?/* key to indicate "acceptable" and the *Z* key to indicate "unacceptable." Subjects were asked to make their judgments as quickly and accurately as possible. There were 32 subjects.

For Experiment 3, the 16 manner of motion verbs and 16 external cause verbs were mixed with filler words and nonwords in 11 blocks of 16 items (preceded by 1 practice block), with 8 words and 8 nonwords in each block. Subjects were instructed to respond as quickly and accurately as possible to each item. A correct response was followed by a 50-ms blank interval and then the next test item. An error was followed by the message *ERROR* for 900 ms, then the 50 ms blank interval, and then the next test item. Order of the items was random except that each of the manner of motion and external cause verbs was immediately preceded by a nonword. There were 35 subjects.

### *Results*

In the first 2 experiments, responses for sentences with externally caused change of state verbs were slower than responses for sentences with manner of motion verbs, as predicted from the complexity difference in the lexical semantic event templates of the two classes of verbs. In Experiment 1, with acceptable intransitive sentences, mean reading time for the sentences with external cause verbs was 1,693 ms and mean reading time for the sentences with manner of motion verbs was 1,546 ms; the difference was significant, $F_1(1, 19) = 7.8$, $F_2(1, 15) = 4.7$, $SE = 38$ ms. In Experiment 2, with unacceptable transitive sentences, the mean time to judge unacceptable the sentences with external cause verbs was 1,744 ms (96% "unacceptable" judgments), compared with 1,656 ms (96% "unacceptable" judgments) for sentences with manner of motion verbs; this difference was also significant, $F_1(1, 30) = 6.3$, $F_2(1, 15) = 5.0$, $SE = 25$ ms.

In Experiment 3, lexical decision responses for the external cause verbs were significantly slower than responses for the manner of motion verbs, $F_1(1, 34) = 12.9$, $F_2(1, 15) = 6.1$, $SE = 12$ ms. The mean response times were 779 ms (5% errors) and 716 ms (3% errors).

We take the data from these experiments as evidence that the single basic template for manner of motion verbs is a simple event template whereas the single basic template for externally caused change of state verbs is more complex.

There are several alternative accounts of the data that might be advanced, but they can be ruled out. First, if manner of motion verbs had two basic templates, an activity template and a more complex, external cause template, then in Experiment 2 the response times for sentences with the manner of motion verbs should not have been shorter than the response times for sentences with externally caused change of state verbs. Or if both classes of verbs had both a simple and a more complex template, there should have been no response time differences between them in any of the experiments.

Second, it might be proposed that only some manner of motion verbs have two basic templates, specifically those verbs (like *gallop*) that are used in external control, transitive sentences, as well as intransitive sentences. According to corpus data, there were eight such verbs in the experiments. The data for these verbs and their externally caused change of state pair mates show the same pattern as for all the verbs: shorter response times for the manner of motion verbs than the externally caused change of state verbs (collapsing across all three experiments), $F_1(1, 86) = 6.7$, $F_2(1, 21) = 4.9$ (the means for the three experiments, manner of motion verbs versus external cause verbs, are, respectively: 1,586 ms vs. 1,690 ms; 1,690 ms vs. 1,811 ms; and 666 ms vs. 722 ms.)

Another possible alternative hypothesis, which might be proposed by the constraint-based models of sentence processing that are discussed in the latter half of this article, is that response times in the experiments depended on the probabilities with which the verbs are used in various syntactic constructions, not on the underlying event templates. To test this hypothesis, we examined response times for manner of motion and externally caused change of state verbs for which the probability of appearing in transitive versus intransitive constructions was approximately equated. Overall, as measured from a random sample of 100 tokens of each verb in the corpus, the manner of motion verbs are less likely to appear in the corpus as transitives than the external cause verbs (the respective probabilities are .13 and .41), but for eight of the verbs of each class, the probabilities are about the same (.19 and .21, respectively). For these verbs, the mean response times show the same pattern as for all the verbs, with responses for manner of motion verbs shorter than responses for external cause verbs, $F_1(1, 86) = 11.0$, and $F_2(2, 42) = 3.4$, $p = .075$. In the three experiments, the mean response times for these eight manner of motion and eight external cause verbs were, respectively: 1,566 ms vs. 1,640 ms; 1,656 ms vs. 1,737 ms; and 728 ms vs. 769 ms. These response times are not noticeably different from the response times for all the verbs together, suggesting that the probabilities with which the verbs appear in surface structure syntactic configurations have little effect on processing time, a point taken up in the latter half of this article.

### *Summary and Discussion*

With the assumption that manner of motion verbs have a single basic lexical template, the activity template *x*(ACT), marked for internal causality, a wide range of data can be explained. Corpus Study 1 investigated the kinds of sentences speakers and writers produce. All the corpus tokens for 69 manner of motion verbs were examined, and the data provide convincing evidence for a single, basic, internal causality, activity template. External control uses of the verbs are rare, and when the verbs do denote events under external control, there is a tight restriction, namely that the external, controlling entity must be a force superior to that of the entity engaging in the verbal activity. The assumption that there is a single basic template for the manner of motion verbs also accounts for the shorter processing times for both transitive and intransitive sentences with manner of motion verbs versus externally caused change of state verbs, and it accounts for shorter lexical decision response times for manner of motion verbs.

The corpus and experimental data jointly rule out a number of other possible accounts of the lexical representation of manner of motion verbs. First, any proposal by which the verbs are assumed to have two basic lexical representations, one for intransitive sentences and one for transitive sentences, has difficulty in explaining the data. Decompositional proposals like this have been made by Pustejovsky (1991; but see Pustejovsky, 1995), Stevenson and Merlo (1997), and van Valin and LaPolla (1997). Also, many constraint-based models assume that the lexical representation of a verb has multiple role lists, one for each syntactic construction in which the verb occurs. If manner of motion verbs and externally caused change of state verbs both had one representation underlying transitive sentences and another representation underlying intransitive sentences, whether the representations were decompositional or role lists, the differences in response times in Experiments 1, 2, and 3 would not be predicted.

The differences in event template complexity proposed by MTS explain the processing time differences between manner of motion and externally caused change of state verbs. In addition, the notion of internal causality (Levin & Rappaport Hovav, 1995; McKoon & Macfarland, 2000b; Smith, 1970) is needed to explain the results of the corpus study. Because manner of motion verbs denote events internally controlled by the entities engaging in them, sentences expressing control taken over by an external entity occur rarely and only under tightly defined circumstances, specifically when the external controller is a superior force. Also, both internal causality and the restriction to superior forces as external controllers are semantic notions. Therefore, it would be difficult to capture them in any system in which the lexical representation of syntactically relevant information for verbs was not formulated in semantic terms.

The internal causality feature of the event template for manner of motion verbs is central to the MTS explanation for the prohibition of manner of motion verbs in reduced relative clauses. The entity engaging in the activity denoted by a manner of motion verb does so under internal, not external, control, and so the entity cannot be the head of a reduced relative. We take up evidence about reduced relative clauses after discussing a second class of verbs marked for internal causality.

## Event Templates: Internally Caused Change of State Verbs

Internally caused change of state verbs are the second class of verbs for which we claim that reduced relatives are blocked. This class of verbs has been extensively studied by McKoon and Macfarland (2000b, 2000c) and their findings are reviewed here.

For internally caused change of state verbs, there must be at least a basic, change of state template to underlie intransitive sentences like *The flower bloomed.* Clearly, the flower changes state and it does so because of some internal process.

The claim is that this is the only basic template for these verbs. Many of the verbs of this class appear frequently in transitive sentences, but the hypothesis is that both transitive and intransitive sentences are projected from the same underlying change of state template. The transitive sentences differ from the intransitives only in that the transitives make explicit an optional argument that arises from the part of a verb's meaning that is not represented in the event template. The optional argument has no expression of its own in the event template. Instead, it is a part of the change of state process (which is why we indicate it with a subscript on the change of state in Table 2). The sentences below show examples of transitive subjects that are natural entities participating in an internally caused change of state, denoted by an internally caused change of state verb. *Water*, for example, participates in the process of *everything rotting*. In each case, we claim, the subject of the transitive verb is not an external controller; instead it is a participant in the internally controlled event. The sentences are as follows:

*The water rots everything . . .*

*. . . the kind of high-noon sun that wilts cactus.*

*But the record cold that chilled the city in January withered many of the azaleas . . .*

*The water table in New York City has risen . . . eroding some subway tunnels so badly . . .*

*. . . the sea air deteriorated the steel work . . .*

McKoon and Macfarland (2000b) supported this analysis of transitive uses of internally caused change of state verbs with corpus data. They compared transitive sentences with internally caused change of state verbs (as in the sentences above) to transitive sentences with externally caused change of state verbs (like *break*). As predicted, the verbs of the two classes differ in the restrictions they place on the subjects of transitives. For externally caused change of state transitives, the subjects can be any of a wide variety of types of entities, including animate agents, natural forces, manmade artifacts, and abstract entities. But for the internally caused change of state transitives, the subjects are restricted to entities that are inherent to the natural course of the change of state described by the verb, as water is inherent to the subway tunnels eroding. Other kinds of entities like agents or instruments cannot be inherent participants of internally caused change of state events, and so they are prohibited from appearing as subjects of transitives. One cannot say, for example, that *John eroded the tunnels* or that *A hammer eroded the tunnels*. (Of course, in the world, there is no reason that a person or a hammer could not accomplish something like the erosion of a tunnel; what cannot happen is *erode* being used to describe such events because *erode* conveys internal control.[5])

McKoon and Macfarland (2000a, 2000b, 2000c) also provided psycholinguistic data to support the change of state template being the single basic template for internally caused change of state verbs. The experiments were analogous to those just described for manner of motion verbs, and illustrative data are shown in Table 3. The pattern of the data is that internally caused change of state verbs lead to faster processing than externally caused change of state verbs. This is true for acceptable intransitive sentences,

[5] McKoon and Macfarland's (2000b) finding covers only transitive sentences with concrete objects like *crops*, not sentences with abstract objects like *spirit*, because there is no good criterion for deciding what counts as a natural force for the change of state of an abstract entity.

Table 3
*Results (in Milliseconds) From Experiments by McKoon and Macfarland (2000b, 2000c)*

| Experiment | Externally caused change of state verbs: $\alpha$ CAUSE ($x$(BECOME IN STATE)) | Internally caused change of state verbs: $x$(BECOME IN STATE) |
|---|---|---|
| Reading times: Acceptable intransitive sentences | 1,664 | 1,557 |
| Reading times: Acceptable transitive sentences | 2,561 | 2,405 |
| Judgment times: Unacceptable transitive sentences | 1,885 | 1,752 |
| Lexical decision response times | 778 | 719 |

acceptable transitive sentences, anomalous transitive sentences, and for lexical decision.

Just as with the manner of motion verbs, a counterargument might be advanced that not all internally caused change of state verbs have only the basic change of state template, that some have a second, external control, basic template. These would be the verbs that have a high probability of occurrence in transitive sentences. McKoon and Macfarland (2000c) provided evidence against this possibility. For subsets of the two classes of verbs, the probabilities that the verbs appeared in the corpus as transitives were about equal. For these verbs, there was still a large difference in response times favoring the internally caused change of state verbs.

We conclude that internally caused change of state verbs have only one basic event template, the simple change of state template. Given this one template, marked for internal causality, all the data can be explained: The restriction on subjects of transitive sentences in the corpus sentences, the shorter processing times for sentences with internally as opposed to externally caused change of state verbs, and the shorter lexical decision times for internally than externally caused change of state verbs.

As with the data for manner of motion verbs, the data for internally caused change of state verbs are difficult to explain by other accounts of verbs' lexical semantic representations. Any proposal by which verbs of both the internally caused and externally caused change of state classes are said to have two different lexical representations, one for transitive and one for intransitive uses of the verbs, would have difficulty predicting the consistently shorter processing times for the internal cause verbs. The restriction on the subjects of transitive sentences for the internal cause verbs (that they be inherent participants in the change of state) is also difficult to explain by other proposals, either because the representations of syntactically relevant information are not semantic or because internal versus external causality is not included as a factor in their representations. Finally, without the notion of internal causality, the prohibition of verbs of internal causality in reduced relatives cannot be explained.

## The Meaning of the Reduced Relative Construction

At this point, we turn from evidence that supports the MTS analysis of the lexical semantic event templates of manner of motion and internally caused change of state verbs to evidence that supports the analysis of the meaning of the reduced relative clause construction. We first provide a test of the critical pretheoretic assumption that different constructions have different meanings and then provide two tests of our analysis of the reduced relative's meaning.

### *Construction-Specific Meaning: Experiment 4*

As was pointed out earlier, it is generally assumed in the psycholinguistic sentence-processing literature that reduced and nonreduced relative clauses have the same meaning and that they differ only in the difficulty with which this meaning can be interpreted from surface form. Nonreduced relatives are "better" because their surface form is easier to process. The meaning that is encoded into memory for the two constructions is assumed to be the same.

The MTS claim is different. The MTS claim is that reduced and nonreduced constructions have different meanings and so nonreduced relative clauses are not, in any universal sense, better than reduced relative clauses. A reduced relative clause will be better when its meaning better matches the intended meaning of its utterance and a nonreduced relative will be better when its meaning is the better match.

Experiment 4 tested this claim by asking subjects to rate corpus sentences with reduced and nonreduced relatives according to how well they were written. The sentences were presented to subjects either in their original form, as they appeared in the corpus, or in a new form such that reduced relatives were changed to nonreduced relatives and vice versa. The MTS prediction was that sentences would be given better ratings when their relative clauses were in the original form because the original form would better match the meaning of the utterance. This was expected to be true no matter whether the original form was reduced or nonreduced.

#### *Method*

Thirty sentences with reduced relative clauses and 30 sentences with nonreduced relative clauses were chosen randomly from the corpus (except that nonreduced relatives with modifiers on the verb such as *might have* were excluded because the modifiers could not easily be retained for the reduced version). There were two versions of each sentence, the original version and a new version with the relative clause changed from reduced to nonreduced or vice versa. To change a sentence from reduced to nonreduced, we added the appropriate pronoun (*who* or *that*) and we added the appropriate auxiliaries to the verb in the clause. Examples are given below. Twenty of the sentences that originally contained reduced relatives and 20 of the sentences that originally contained nonreduced relatives were mixed with 10 filler sentences and presented to 38 subjects for rating on a scale of 1 to 7, with 7 being *very well written* and 1 being *very poorly written*. Half the sentences were presented in their original version, half in the new version, counterbalanced across subjects. The other 10 sentences that were originally reduced and the other 10 that were originally nonreduced were presented in the same way to 48 subjects.

ORIGINAL REDUCED RELATIVES

*"There have been many preposterous reasons advanced to support a capital-gains tax cut," Sen. Mitchell said during his television appearance.*

*The growing controversy comes as many practices historically accepted as normal are coming under close ethical scrutiny.*

*I was particularly interested in the relics of the whaling fleets, that is, the objects fashioned by the sailors from whalebone.*

*The subject most frequently debated by the Fergie followers on Fleet Street is the royal figure.*

ORIGINAL NONREDUCED RELATIVES

*The Swedish company that was edged out of the competition was identified by the investigators as Kockums.*

*The circuitous journey that is embarked upon in John Hughes's "Planes, Trains and Automobiles" is supposed to range from New York to Chicago, but its final destination is surprising.*

*The IAFP recently assembled industry groups to discuss common standards that are applied to planners.*

*The eight designs were originally submitted to bring some life to a skyline that has been deadened by some of the dreariest modern boxes in the world.*

### *Results*

Following the MTS prediction, sentences were rated as better written in their original versions. Sentences that originally contained nonreduced relatives were rated better in their original version than in their new version (mean ratings 4.7 vs. 4.2), $F_1(1, 85) = 12.7$, and $F_2\ (1, 29) = 7.4$.

The more interesting result is for sentences that originally contained reduced relatives. Given the usual assumption that reduced relatives are generally more difficult to process, a switch from reduced to nonreduced should have led to better ratings. It did not. Sentences that originally contained reduced relatives were rated 4.7 in their original version and 4.4 when switched to nonreduced relatives, a significant difference, $F_1(1, 85) = 8.1$, and $F_2(1, 29) = 9.0$.

For 10 of the sentences, the relative clause was in a position such that there was a clear ambiguity between the relative clause reading and a main clause reading. For example, in the *Fergie* sentence, *The subject most frequently debated* could have continued as a main clause. For these 10 sentences, the switch from reduced to nonreduced might have been expected to be especially beneficial, but it was not. Just as for all the sentences, the switch to a nonreduced relative led to worse ratings: For these 10 sentences, the mean ratings were 4.6 and 4.3 for the reduced and nonreduced versions, respectively.

### *Discussion*

The data from Experiment 4 are important not only because they support the MTS pretheoretical assumption of different meanings for different constructions but also because they bring into question a methodology that is often used in experiments designed to investigate the processing of reduced relative clauses. What is typically done is to use nonreduced relatives to provide a control condition against which to measure the processing difficulties engendered by reduced relatives. The assumption underlying this comparison is that a reduced version of a relative clause is always at least as difficult to process as a nonreduced version of the same clause (see Frazier, 1995, for further discussion of this point). The data from Experiment 4 suggest that this is not correct. Subjects rated relative clauses that were originally written as reduced relatives as better sentences than their nonreduced counterparts and this suggests that the reduced relatives were easier to process.

A question raised by these data is why a reduced relative clause is preferred for some sentences and a nonreduced clause for other sentences. Why is it better to say *The subject most frequently debated by the Fergie followers on Fleet Street is the royal figure* than to say *The subject that is most frequently debated by the Fergie followers on Fleet Street is the royal figure*? Given our definitions of the relative clause constructions, it must be because for some sentences, like the *Fergie* sentence, it is better to introduce an entity f($x$) (*The subject most frequently debated by the Fergie followers on Fleet Street*) into the discourse and for other sentences it is better to separately introduce an entity $x$ (*The subject*) and a fact about that entity, g($x$) (*most frequently debated by the Fergie followers on Fleet Street*). For naturally occurring reduced relatives, we assume that the single entity f($x$) is better than the two separate entities $x$ and g($x$) at fitting the meaning of the rest of the sentence plus whatever of our world knowledge is relevant to the sentence.

## The Definition of a Discourse Entity: Corpus Study 2

One element of the MTS definition of the reduced relative clause construction is that the construction serves to identify or define an entity in the sense that it singles out from a larger set of possibilities just those entities for which it is the case that they have participated in the event denoted in the clause. For example, for *Agencies affected by the measure* . . . , the set of agencies being denoted is a subset of all agencies, specifically those that are affected by the measure.

Generally speaking, this element of the MTS view of reduced relatives is consistent with the standard grammarian's prescription that restrictive relative clauses define or limit (e.g., Strunk & White, 1979). It is also consistent with Fox and Thompson's (1990) definition of the function of relative clauses, that is, that they characterize, define, or identify an entity. Fox and Thompson go on to describe the various discourse purposes relative clauses can serve, such as relating the entity denoted by the reduced relative to other information in the discourse, but such functions are outside the scope of what we consider here.

With reduced relatives picking out a subset of entities from a larger set of possible entities, MTS follows Altmann and Steedman's proposals (1988; Crain & Steedman, 1985; Ni, Crain, & Shankweiler, 1996). They emphasize that the comprehension of individual sentences can only be understood in relation to a broader context of knowledge. The set of possible entities presupposed by a reduced relative might be found in the immediate discourse context or it might be found in general world knowledge or it might have to be implicitly constructed on line. In any case, the reader or hearer must accommodate the reference by introducing the specified entity into the discourse.

If it is correct that a reduced relative identifies or defines a particular subset of entities out of a larger set, then often these entities will be new to the discourse and so the head of a reduced relative should often be indefinite. For example, the reduced relatives *agencies affected by the measure* . . . , *all Salvadorans arrested on Long Island* . . . , and *a policeman assigned to the embassy* . . . all have indefinite heads. Of course, the heads of reduced relatives are not required to be indefinite. For example, the definite article *the* with the head of a reduced relative can be used to reference a well-known entity that has not yet been brought into the discourse (e.g., *the worries and fears associated with growing old*) or it can be used to distinguish one entity already in the discourse from another. Overall, however, compared with noun phrases in other syntactic positions, there should be a greater tendency for the heads of reduced relatives to be indefinite.

We tested this prediction by comparing noun phrases that occurred as the heads of reduced relatives with noun phrases that occurred as the subjects of main clauses. For completeness, we also included noun phrases that occurred as the heads of nonreduced relatives. We chose 24 verbs that each appeared in a nonreduced relative clause at least five times in the corpus. For these 24 verbs, we examined all sentences from the corpus with the verb in a reduced relative, in a main clause, and in a nonreduced relative, up to a limit of 15 sentences per verb in each case. The 24 verbs were *accepted*, *accused*, *affected*, *assigned*, *associated*, *attacked*, *convicted*, *cured*, *evaluated*, *expected*, *hired*, *hypnotized*, *identified*, *interrogated*, *kicked*, *overlooked*, *punished*, *reviewed*, *terrorized*, *tortured*, *towed*, *transported*, *washed*, and *worshiped*. These are all verbs that can denote events that fit the reduced relative requirement of external causality. These 24 verbs are a subset of the larger set of verbs examined in the research described in the latter half of this article.

Table 4 shows the total numbers of sentences examined and the proportions of noun phrases of each kind. Indefinite noun phrases are divided into those for which the noun is "bare," meaning there is no article or descriptor of any kind (as in *agencies affected by the measure* . . .) and those for which there is an article or other definite descriptor. The data are consistent with the prediction. The entities that are the heads of reduced relative clauses are much more often indefinite than the entities that are the subjects of main clauses, a significant difference $F_2(4, 207) = 5.9$. The nonreduced relative clauses show the same pattern as reduced relatives in these data, but nonreduced relatives can be distinguished from reduced relatives by other patterns of data, as is discussed later.

Table 4
*Proportions of Naturally Occurring Noun Phrases*

| Construction | Indefinite: Bare | Indefinite: Article or other indefinite descriptor | Definite |
|---|---|---|---|
| Subjects of main clauses (164 sentences) | .26 | .25 | .49 |
| Heads of reduced relatives (312 sentences) | .39 | .38 | .24 |
| Heads of nonreduced relatives (174 sentences) | .35 | .39 | .27 |

## *The Packaging of a Discourse Entity: Sentence Memory, Experiments 5 and 6*

When the reduced relative construction introduces an entity into a discourse, the entity is defined and denoted by the reduced relative construction as a whole. In the sentence *Miss Comfort suddenly stretched to reveal fingers lengthened by spikes of artificial nails*, the single entity *fingers lengthened by spikes of artificial nails* is introduced by the reduced relative. Of course, there are other entities mentioned in the clause (e.g., the artificial nails) and these might, for some discourse purposes, be treated as discourse entities in their own right. But the MTS hypothesis is that the main function of the reduced relative is to provide a package ($f(x)$) that references a single discourse entity. In contrast, a nonreduced relative introduces both an entity and information about the entity ($x$ and $g(x)$). A nonreduced version of the *Miss Comfort* sentence would introduce both *fingers* and *fingers were lengthened by artificial nails.*

The prediction that follows from these hypotheses is that the elements of a reduced relative should be packaged more tightly together than the elements of a nonreduced relative. We tested this prediction with cued recall. Subjects studied lists of sentences and then were asked to recall the sentences in response to cues. The measure was the probability of recalling the contents of a relative clause given that the head was recalled. We expected that this probability would be greater for reduced relatives than for nonreduced relatives. In other words, for the *Miss Comfort* sentence, given recall of the head *fingers*, subjects should be more likely to recall *lengthened by spikes of artificial nails* if the information had been presented in a reduced relative clause than in a nonreduced relative clause.

### *Method*

*Materials.* For Experiment 5, there were two sets of sentences. In the first set were 20 sentences chosen from the *WSJ* part of the corpus, and in the second set were 40 sentences chosen from the *NYT* part of the corpus. In each set, half the sentences contained a reduced relative clause and half a nonreduced relative clause. For each sentence, a question was written to which the answer was the relative clause information. The question contained most of the words of the sentence except the relative clause and its head, as shown by the examples below.

Sentence: *A stock split accompanied by a dividend increase can pack a powerful punch.*

Question: *WHAT can pack a powerful punch?*

Sentence: *Still other projects that have been announced could put expenditures above three billion dollars over the next few years.*

Question: *WHAT could put expenditures above three billion dollars over the next few years?*

Sentence: *The experience gained by residents from their participation on community boards is a practical preparation which can be applied to all issues and decisions.*

Question: *The WHAT is a practical preparation which can be applied to all issues and decisions?*

For the two sets of sentences, reduced and nonreduced, the mean numbers of words in the sentences were, respectively, 13.6 and 17.7

(*WSJ*), 18 and 19 (*NYT*); the mean numbers of words in the heads of the relative clauses were, respectively, 2.2 and 2.4 (*WSJ*), 2.7 and 2.6 (*NYT*); the mean numbers of words in the clauses were, respectively, 4.5 and 3.6 (*WSJ*), 5.0 and 4.9 (*NYT*) (not counting the relative pronoun and the auxiliary verb in the nonreduced relatives); the mean frequencies (Francis & Kučera, 1982) of the content words of the heads of the relative clauses were, respectively, 129 and 192 (*WSJ*), 122 and 96 (*NYT*); and the mean frequencies of the content words in the clauses were respectively, 107 and 210 (*WSJ*), 53 and 66 (*NYT*).

In Experiment 5, each sentence was used in the experiment just the way it had appeared in the corpus, as either a reduced relative or a nonreduced relative. For Experiment 6, each sentence as it appeared in the corpus contained a reduced relative clause and each was studied in the experiment either in its original reduced relative version or a new nonreduced version (counterbalanced across subjects). There were 36 sentences randomly chosen from the corpus.

*Procedure and subjects.* For both experiments, the sentences were presented on a PC screen in blocks of four, each block containing two fillers and two experimental sentences. Each block began with an instruction to *press the space bar* on the PC keyboard to initiate the block. The four sentences were presented one at a time in random order, each for an amount of time determined by summing 200 ms times the number of words in the sentence plus 20 ms times the total number of letters in the words. There was a 1,500-ms blank interval following each sentence. After the fourth sentence, the word *questions* was displayed for 1,000 ms, and then the questions were displayed one at a time. Subjects were instructed to answer each question in writing in their own words on an answer sheet that was provided. They pressed the space bar on the keyboard to advance from each question to the next. For Experiment 5, 36 subjects were tested with the first set of sentences and 28 with the second set of sentences, and for Experiment 6, 12 subjects were tested with 16 of the sentences and 32 with the other 20 sentences.

### Results

For each experimental sentence, the probability of recall of any content word of the relative clause, conditionalized on recall of any content word of the head, was calculated. The MTS prediction was that this probability would be greater for reduced than nonreduced relative clauses because the information in the clause is hypothesized to be more tightly bound to the head for reduced relatives. This prediction was confirmed by the data. For Experiment 5, the probabilities were .64 and .38 for reduced and nonreduced clauses, respectively, for the first set of sentences (*WSJ*) and .64 and .48, respectively, for the second set of sentences (*NYT*). These differences were significant by analyses of variance with sentences as the random variable, $F_2(1, 16) = 4.4$, and $F_2(1, 38) = 5.6$. For Experiment 6 (from which four sentences were deleted because no subject recalled any words from them), the probabilities were .45 and .35, respectively, $F_2(1, 31) = 4.0$. Analyses of variance with subjects as the random variable were not conducted for Experiments 5 and 6 because many subjects recalled nothing from the heads of many of the experimental sentences. (Subjects hated these experiments.)

The probabilities of recalling any content word of the heads of the reduced relatives were also calculated in order to check that the heads themselves did not differ significantly in difficulty between the reduced and nonreduced clauses. The probabilities for reduced and nonreduced, respectively, were as follows: Experiment 5 (*WSJ*), .52 and .38; Experiment 5 (*NYT*), .61 and .67; Experiment 6, .31 and .30 (all $F_2$s $< 1.7$).

### Discussion

We have interpreted the results of Experiments 5 and 6 as supporting the MTS hypothesis that the information in the head of a relative clause and the clause's contents are more tightly bound together if the clause is reduced than if it is nonreduced. An alternative hypothesis might be suggested, namely that reduced relatives are restrictive clauses and nonreduced relatives are nonrestrictive. According to Strunk and White (1979), restrictive clauses "limit or define" (p. 4) the entities that are their heads, whereas nonrestrictive clauses simply add information. For the experimental sentences in Experiments 4, 5, and 6, all of the relative clauses were written to limit or define the entities that are their heads and so both the reduced and nonreduced relatives were restrictive. However, the question remains as to whether the reduced and nonreduced relative clauses from our corpus studies are restrictive. This question is taken up in the next section.

## Restrictive and Nonrestrictive Relative Clauses: Corpus Study 3

At the outset, we restricted the MTS hypotheses about reduced and nonreduced relative clauses to restrictive clauses. Thus, it is necessary to ensure that the corpus data collected to support MTS are, at least in large proportion, restrictive clauses. Strict rules of grammar (e.g., Strunk & White, 1979; Wilson, 1993) limit nonrestrictive clauses to those for which the relative pronoun is *which* and those for which the clause is set off by punctuation, usually commas. No such clauses were included in any of the tabulations of data from the corpus that are reported here. The only clauses counted as nonreduced relatives were those with the pronouns *who* and *that*, and no clauses were counted as reduced or nonreduced if they were set off by punctuation.

However, users of English may not follow the strict rules; some relative clauses may be nonrestrictive even though they are not set off by punctuation and do not use *which*. To check whether this is the case, we examined a large number of relative clauses from the corpus. The clauses came from corpus studies described later in this article in which 200 randomly chosen sentences were examined for each of 85 verbs. The verbs are all those used in relative clauses in Experiments 8–11 in this article and in studies by MacDonald, Pearlmutter, and Seidenberg (1994), McRae, Spivey-Knowlton, and Tanenhaus (1998), and Trueswell, Tanenhaus, and Garney (1994). For the total of about 17,000 sentences examined, the verbs were used in reduced relative clauses in 6% of the sentences and in nonreduced relative clauses in 2% of the sentences, as is discussed in more detail later. The total number of relative clauses was 1,236.

Each of these 1,236 clauses was categorized as restrictive or nonrestrictive. A clause was counted as restrictive only if it limits or defines the entity appearing as head of the clause. For the great majority of the clauses, the head is some generic category of entities and the relative clause is restrictive in that it defines or identifies a subset of those entities, as in the two sentences below where *chips* and *styles* are two generic categories and the clauses define subsets of them:

*Those royalty payments were based on chips sold in the United States.*

*In the old days, district merchandise specialists would make frequent buying trips to New York to pick out the styles that would be sold in Penney stores in their territories.*

Other clauses are restrictive in that they define the entity in head position by giving an identifying characteristic. In the first sentence below, the relative clause defines what kind of man Mr. Wehner is, and in the second sentence, the relative clause defines a critical property of the ex-friend:

*Mr. Wehner, a man once feared and admired, loved and despised, sits puffing a large briar pipe, his sandy hair gone white, his broad brow . . .*

*. . . if Rose gets drawn, as expected, into the trial of Thomas Gioiosa, an ex-friend who is accused of not paying taxes on a winning Pik-Six ticket worth $47,646 . . .*

The clauses that were counted as nonrestrictive did not limit or define the entity in head position in any obvious way. Some examples are given below. The entities in the head position—outraged peasants; the response by the Reagan Administration; the damp, badly heated apartments; and the Wallis pictures—are all at least reasonably well-defined without the information in their respective relative clauses. Sometimes, the judgment about whether a relative clause is restrictive or nonrestrictive is difficult; we categorized clauses about which we were unsure as nonrestrictive.

*Outraged highland peasants armed with sickles and homemade shotguns captured 13 Shining Path guerrillas . . .*

*. . . prompted a response by the Reagan Administration that was carried out in secrecy.*

*. . . having to live in damp, badly heated apartments concealed by the freshly painted facades.*

*. . . four other Wallis pictures to be auctioned later that are expected to bring $30,000 to $150,000 each.*

Even with the inclusion of “unsure” cases, there were few nonrestrictive clauses, only 45 out of the total of 1,236 examined. Only 3% of the reduced relative clauses and 5% of the nonreduced relative clauses were nonrestrictive—a difference that was not significant, $F_2(1, 84) < 1.0$. With so few nonrestrictive clauses (and with many of these “unsure” cases that might possibly be restrictive) and with the clauses categorized as nonrestrictive about evenly distributed between reduced and nonreduced, we can be reasonably sure that the empirical differences we find in corpus studies between reduced and nonreduced relative clauses are not due to a confound with whether the clauses are restrictive versus nonrestrictive.

### Summary

At this point, we have provided strong support for the MTS definition of the reduced relative clause. First, subjects’ ratings show that contrary to the usual belief in the sentence processing literature, reduced relatives can be preferred over nonreduced relatives. According to the MTS view, this is because reduced and nonreduced relative clauses have different meanings, and one meaning may be preferred in one context, the other meaning in another context. Second, according to both the MTS view and Altmann and Steedmann’s (1988) view, reduced relative clauses often serve to enter new entities into a discourse and so should be likely to be indefinite. Corpus sentences provide evidence that this is the case. Third, the MTS view is that reduced relatives package information more tightly than nonreduced relatives do. As a consequence, given recall of the head of the clause, subjects should be more likely to recall the contents of the clause for a reduced relative than a nonreduced relative. Experiments 5 and 6 confirm this.

Not only do these findings support the MTS view, they make the more general point that reduced and nonreduced relative clauses cannot be treated as interchangeable in all ways except surface form. If they actually were interchangeable, then subjects would not be expected to disprefer sentences switched from reduced relatives to nonreduced relatives—the switch to a nonreduced relative should only make processing easier, which is not a reason to disprefer it. Furthermore, if they actually were interchangeable, then subjects would be just as likely to recall the contents of a nonreduced relative clause as the contents of a reduced relative clause. In sum, it appears that reduced and nonreduced relative clauses do have different meanings, consistent with MTS and with the claims made by linguists about the meanings of syntactic constructions (e.g., Goldberg, 1995).

## REDUCED RELATIVE CLAUSES ARE BLOCKED FOR INTERNAL CAUSE VERBS

Previously, we summarized the MTS thesis for reduced relative clauses by saying that *The horse raced past the barn fell* is unacceptable because the lexical, event template meaning of the verb *race* denotes causality for the racing event as being internal to the horse, and such internal causality is incompatible with the construction meaning of the reduced relative clause. So far, we have presented corpus and psycholinguistic data to support our hypotheses about the event template meanings of verbs like *race* and about the meaning of the reduced relative construction. It remains to demonstrate the consequences of the interaction of the verb template meanings with the construction meaning, that is, to provide evidence that *The horse raced past the barn fell* is indeed unacceptable.

Given the claims that internal causality is incompatible with the meaning of the reduced relative and that verbs like *race* denote internal causality, then MTS makes a strong prediction: Reduced relatives of the type *The horse raced past the barn* should never occur in naturally produced sentences. Of course, given errors by speakers and writers, the probability of such sentences may not be exactly zero, but it should be close to zero. The corpus data presented in the next section confirm this prediction.

### Reduced Relatives and Corpus Data: Corpus Studies 4 and 5

For Corpus Study 4, we examined 39,159 sentences that contain manner of motion verbs. These were all the tokens in the corpus for the *-ed* ending form (or the irregular form for verbs like *creep*) of 69 manner of motion verbs listed by Levin (1993). We examined tokens for all the manner of motion verbs listed by Levin except those for which the majority of uses were for a meaning of the verb different than the manner of motion meaning (e.g., *bolt*, *skip*). The 69 verbs are as follows: *amble*, *canter*, *carom*, *cavort*, *clamber*, *climb*, *coast*, *crawl*, *creep*, *dart*, *drift*, *float*, *frolic*, *gallop*, *glide*, *hike*, *hop*, *hurry*, *hustle*, *jump*, *leap*, *limp*, *lope*, *lurch*, *march*, *parade*, *plod*, *prance*, *prowl*, *race*, *ramble*, *roam*, *romp*, *saunter*, *scamper*, *scoot*, *scram*, *scurry*, *shamble*, *sidle*, *skulk*, *slide*, *slink*, *slither*, *sneak*, *speed*, *stagger*, *stray*, *stroll*, *strut*, *stumble*, *swagger*, *swim*, *tiptoe*, *toddle*, *totter*, *traipse*, *tramp*, *trundle*, *trek*, *troop*, *trot*, *trudge*, *vault*, *waddle*, *wade*, *walk*, *wander*, and *zoom.*

Among the 39,159 sentences, there are only 6 sentences in which a manner of motion verb is used in a reduced relative clause with the entity that engages in the motion event appearing as the head of the clause. The verbs in 4 of these reduced relatives are *paraded* and *hurried*. For these 4 sentences, the head of the reduced relative is an inanimate entity that cannot ever engage in the verbal activity under its own internal control. The reduced relatives for these sentences are as follows:

> *The ready-to-wear collections paraded at the Italian fashion center* . . .
>
> . . . *the beefy silhouette best paraded in a bevy of unisex exercise duds* . . .
>
> . . . *banners and posters paraded by millions of demonstrating East Germans* . . .
>
> *Extra editions hurried into print by Japanese newspapers* . . .

The fact that collections, silhouettes, and posters cannot ever have internal control over their own parading, nor editions control their own hurrying, strongly suggests that these tokens reflect different meanings for the verbs than the internally caused manner of motion meanings.

The only two reduced relatives in which manner of motion verbs are clearly used in internal control, manner of motion meanings are as follows:

> *Malham is a village with stone cottages and farmsteads of the 17th and 18th centuries, periods when it was a major assembly point for cattle walked down the Pennine range from Scotland.*
>
> *As he typed he blundered like a young foal strayed from his mother* . . .

Counting all 6 tokens, the reduced relatives are .00015 of the total of 39,159 sentences. This probability is near zero, as MTS predicts. The result means that *The horse raced past the barn fell* can appear as a naturally occurring sentence only with near-zero probability. This is because, according to MTS, the template for the verb *race* denotes an internally caused event, that is, an event under the control of the entity engaging in the event. The reduced relative construction requires that control of the event be external to the entity positioned as head of the reduced relative.

A possible alternative account of these data is that there is a prohibition, not on reduced relative uses of the verbs, but more generally on all passive uses of the verbs. In an object reduced relative, the head of the clause is the underlying object of a passive verb; for example, in *The horse raced past the barn fell,* the head *the horse* is the underlying object of the passive verb *was raced*. However, a prohibition on all passives can be ruled out because passive uses of manner of motion verbs are substantially more frequent in other constructions than they are in reduced relatives. We checked the corpus for passive tokens for the 69 manner of motion verbs used for Corpus Study 1. The probability of passives that were not reduced relatives was .04. Thus passives occur in other constructions much more frequently than they do in reduced relatives, indicating that the near-zero probability of occurrence for reduced relatives is not due to a general restriction on passives.

Manner of motion verbs have as their event structure *x*(ACT), marked for the activity being under the internal control of the entity engaging in it. The other class of verbs examined in this article for which reduced relatives are predicted to be blocked are internally caused change of state verbs (*x*(BECOME IN STATE)). For Corpus Study 5, we examined 4,775 sentences for 17 internally caused change of state verbs (listed by Levin, 1993), all the tokens that appeared in the corpus for the verbs in their *-ed* ending form in the appropriate change of state meaning. The verbs were as follows: *blister*, *bloom*, *blossom*, *decay*, *deteriorate*, *erode*, *ferment*, *flower*, *germinate*, *molt*, *rot*, *rust*, *sprout*, *stagnate*, *swell*, *wilt*, and *wither.*

There were only 21 reduced relatives out of the total of 4,775 sentences, 10 of them for the verb *eroded*, the others for the verbs *decay*, *deteriorate*, *ferment*, *germinate*, *rot*, *swell*, and *wilt*. These were .004 of the total number, .002 if *erode* is discarded because it is possible that it was misclassified as an internally caused change of state verb. Examples are as follows:

> . . . *manor homes now decayed or demolished* . . .
>
> *Frank MacShane has come to him from a book about Raymond Chandler, another good writer deteriorated by drink.*
>
> . . . *topsoil eroded from bald hills* . . .
>
> . . . *carpets rotted by water seepage* . . .
>
> . . . *annuals wilted by the recent freeze* . . .

As with manner of motion verbs, the probability of reduced relative uses of the verbs is much lower than is the probability of passive uses of the verbs in other constructions. For the 17 internally caused change of state verbs, the probability of passive uses other than the reduced relative is .07 (.04 without *erode*).[6]

Averaging over the manner of motion verbs and the internally caused change of state verbs, the probability of a reduced relative in which the clausal event is internally controlled by the entity that is the head of the reduced relative is .0006 (.0003 if *erode* is excluded and only the two manner of motion tokens with animate heads are included). This probability is 100 times less than the probability of a reduced relative in which the clausal event is controlled by some entity external to the head. This latter probability is about .06 (as we show later in the article).

The near-zero probabilities of reduced relatives for the manner of motion and internal cause verbs and the large difference in reduced relative probabilities between externally caused change of state verbs and internally caused change of state and manner of motion verbs both conform to MTS predictions and so give strong support to the MTS view.

[6] With internally caused change of state verbs, it is sometimes not possible to know, either for reduced relatives or for other constructions, whether a sentence represents a passive use of a verb or an adjectival use. For example, . . . *the manor homes were decayed* . . . could be either. In counting corpus tokens, we did not attempt to make a distinction; instead, all tokens that could possibly be passives were counted as passives. This applied for both reduced relatives and other constructions.

## Experiment 7

From the MTS point of view, reduced relatives with manner of motion verbs or internally caused change of state verbs that denote internal causality for the clause head should be judged unacceptable. However, researchers with a different point of view might find such reduced relatives quite acceptable. A reduced relative like *cattle walked from Scotland* might even seem as acceptable as a reduced relative with a verb that provides external causality.

We addressed this issue experimentally by asking subjects to judge the acceptability of sentences with reduced relatives like those we found in the corpus with manner of motion and internally caused change of state verbs. Twenty-two sentences were written, some based directly on the reduced relative sentences we found in the corpus, such as *Cattle walked down the Pennine range from Scotland were assembled at Malham Village*. Others were based on manner of motion and internally caused change of state verbs that are used transitively (e.g., *The boys marched toward the reporters were frightened and refused to answer any questions*). The issue was whether these sentences would be judged as acceptable as sentences with reduced relatives with external cause verbs.

### Method

The 22 sentences with manner of motion and internally caused change of state verbs in reduced relative clauses were matched to sentences with external cause verbs in reduced relative clauses. Just as for the sentences with manner of motion and internally caused change of state verbs, the sentences with externally caused change of state verbs were based on corpus sentences. Examples of the sentences are shown below. For both sets of sentences, the reduced relative clause was always in the subject position for the main verb of the sentence, as in *Cattle walked down the Pennine range from Scotland were assembled at Malham Village*. The mean numbers of words in the reduced relative clauses of the two sets of sentences were equated. The two sets of sentences were also matched in that when their reduced relative clauses were taken out of their sentences and presented as simple, short sentences (e.g., *Cattle walked down the Pennine range from Scotland*), they did not differ significantly in acceptability, as rated by an independent group of subjects (the mean ratings for the sentences with reduced relatives with external vs. internal causal verbs were 4.1 and 4.1 on a scale of 1–5, with 1 being unacceptable). Also, the two sets of sentences with the reduced relative clauses deleted (e.g., *Cattle were assembled at Malham Village*) did not differ significantly in acceptability (again, as rated by an independent group of subjects; mean ratings were 4.3 and 4.2).

REDUCED RELATIVE CLAUSES WITH INTERNAL CAUSE VERBS

*The horse raced yesterday at Del Mar will be sent to stud at the end of the year.*

*The vice-presidential possibilities trooped to Minnesota talked to Mondale, who was taking weeks to make a decision.*

*The taxi drivers trundled off to English classes had come under severe criticism from the tourist bureau.*

*Spindly saplings germinated from seeds are selected for their productive capacity.*

*Manor homes decayed into ruin must be demolished for safety reasons.*

*The shrieking children zoomed to the sky on the roller coaster had made the mistake of filling themselves with candy.*

REDUCED RELATIVE CLAUSES WITH EXTERNAL CAUSE VERBS

*A voice cracked with age told listeners of a long-ago time and a young and pretty wife dead for many years.*

*Adults awakened from deep sleep will often report being in the middle of a dream.*

*Glassy debris shattered into fragments was found under an eight-inch layer of ashes.*

*A ferry capsized in the Burhi Ganga River held 200 drowned passengers, the Times reported today.*

*A submarine submerged thirty miles off the coast successfully test-fired a Poseidon missile today in the Atlantic.*

*The reserves drained from the banking system led the Treasury to compensate by selling bonds.*

Each subject was asked to judge the acceptability of only 5 of the reduced relative sentences, either 2 or 3 with internal cause verbs and either 2 or 3 with external cause verbs. The 5 sentences were mixed with 48 filler sentences, none of which contained relative clauses. The subjects were asked to rate each sentence according to how easy it was to read on a scale of 1 to 5, with 1 being most difficult. Each of the reduced relative sentences was rated by 4 of the total of 41 subjects.

### Results

Despite our efforts to make the reduced relatives with internal cause verbs as acceptable as possible by basing them on naturally occurring tokens, they were still judged less acceptable than the sentences with reduced relatives with external cause verbs. The mean ratings were 2.68 and 3.40 for the sentences with internal and external cause verbs, respectively, with the difference between them significant, $F_2(1, 42) = 10.6$.

## Nonreduced Relatives: Corpus Study 6

Previously, we suggested that one way reduced and nonreduced relative clauses differ is that reduced relatives package the elements of the clause more tightly than nonreduced relatives do. Experiments 5 and 6 supported this idea because, given recall of the head, the contents of a relative clause were more likely to be recalled if the clause was reduced than if it was nonreduced.

The idea that nonreduced relative clauses do not package their elements tightly suggests that the external causality restriction that applies to reduced relatives might not apply to nonreduced relatives, or at least that it might be weakened. To investigate this possibility, we examined the probabilities of reduced versus nonreduced relative clauses for sentences from the corpus.

In general, nonreduced relative clauses are rarer than reduced relative clauses. For the 85 externally caused change of state verbs that are examined in depth later in this article (a total of about 17,000 sentences), the probability of a nonreduced relative is .02 and the probability of a reduced relative is .06, a significant difference, $F_2(1, 84) = 37.8$. If reduced relatives are blocked for internal cause verbs because of the reduced relative's external causality restriction and if the restriction is loosened for nonreduced relative clauses, then the usual trend of nonreduced relatives occurring more rarely than reduced relatives should be lessened or even reversed for internal cause verbs. Corpus data show that this is the case.

We collected all the sentences in the corpus with nonreduced relatives for the same 69 manner of motion verbs and 17 internally caused change of state verbs that were examined for reduced relatives. For the manner of motion verbs, there were a total of 23 nonreduced relatives out of the 39,159 sentences, an occurrence probability of .0006. This is larger than the probability of a reduced relative (which is .00015), and significantly different, $F_2(1, 68) = 5.9$. For the internally caused change of state verbs, there were a total of 15 nonreduced relatives out of the 4,775 sentences, an occurrence probability of .003, not significantly different from the occurrence probability of reduced relatives (.004), $F_2(1, 16) = 2.0$.

In sum, for both the manner of motion and the internally caused change of state verbs, the pattern is different from that for the 85 externally caused change of state verbs. The trend observed for the external cause verbs toward fewer nonreduced than reduced relatives is eliminated for the internally caused change of state verbs and reversed for the manner of motion verbs.

The difference between the manner of motion verbs and the external cause verbs is particularly striking. For the manner of motion verbs, the probability of a reduced relative, .00015, is four times smaller than the probability of a nonreduced relative, .0006. For the external cause verbs, the probability of a reduced relative, .06, is three times larger than the probability of a nonreduced relative, .02. The MTS claim is that the restriction against manner of motion verbs for the reduced relative is loosened for nonreduced relatives and this is exactly what the data show. However, it should also be noted that nonreduced relatives for the manner of motion verbs are much rarer than nonreduced relatives for the external cause verbs—apparently the restriction against manner of motion verbs is merely loosened for nonreduced relatives, not entirely removed. There are two possible ways to think about this loosening. One is that the restriction should apply equally to reduced and nonreduced relatives, but speakers and writers sometimes err in producing nonreduced relatives because of their discourse function: A part of the function is to separate the entity in head position from the information in the clause; that is, in the terms introduced earlier, $x$ is separated from $g(x)$. The separation might lead speakers and writers to err by failing to abide by the prohibition against $x$ being an internal cause of the event denoted by $g(x)$. The second way to think about the prohibition being loosened for nonreduced relatives is that it is caused by some characteristic of the discourse function of nonreduced relatives that we have not yet discovered. As said earlier, a full investigation of nonreduced relatives has not yet been carried out, and so a full explication of their infrequency with manner of motion verbs must be left to future research. Nevertheless, the interaction between the reduced–nonreduced probabilities and type of verb supports the MTS view.

For the internally caused change of state verbs, there is no significant difference between the probabilities of reduced and nonreduced relatives. The three-times-larger difference in probability for a reduced versus nonreduced relative that appears for the external cause verbs is eliminated. This is not so striking an interaction as for the manner of motion verbs, but again it conforms to the MTS prediction of a loosening of the internal causality restriction for nonreduced relatives. Some nonreduced relatives with internally caused change of state verbs seem acceptable, as in these sentences with *fermented* and *wilted*:

> *Mort Subite—An extremely tart beer, the result of the cherries added during fermentation, and the only one of the 10 that is naturally fermented by organisms in the air.*
>
> *At the Quilted Giraffe, Second Avenue, between 50th and 51st Streets, a slice of duck liver is seared and served with lettuce that is wilted in rendered duck's fat, then moistened with a sauce of reduced cassis berries.*

However, despite the apparent acceptability of these tokens, the probability of nonreduced relatives is much smaller for the internal cause verbs than the external cause verbs, .003 versus .02, and further research into nonreduced relatives is needed to explain this.

Overall, the data provide strong support for MTS. They show an interaction between the type of verb in a relative clause and the probability that the clause will be reduced versus nonreduced, an interaction predicted by MTS. For verbs with templates that provide the appropriate external causality for the reduced relative construction, reduced relatives are more frequent than nonreduced relatives. For verbs that do not provide the appropriate causality, reduced relatives are less frequent or equally as frequent as nonreduced relatives. This is an interaction not predicted by other current views of sentence processing.

## SUMMARY

The data from the corpus studies of reduced and nonreduced relative clauses reflect the convergence of the MTS hypotheses about verb meanings and the MTS hypothesis about construction meaning. Manner of motion verbs and internally caused change of state verbs are found only extremely rarely in reduced relative clauses like *The horse raced past the barn fell* because their template meanings include internal control by the entity engaging in the clausal activity, something that the meaning of the reduced relative prohibits. Thus, the hypotheses about verb meanings and construction meaning jointly explain the corpus data for reduced relatives. They do the same for nonreduced relatives. A nonreduced relative is hypothesized to connect the elements of the clause less tightly than a reduced relative, and so the restriction against internal control is weakened. As a result, the probability of a nonreduced relative does not decrease relative to the probability of a reduced relative, as it does with the externally caused change of state verbs that are allowed by the reduced relative construction; instead it increases or is about the same. Thus, for both reduced and nonreduced relative clauses, the data conform to MTS predictions. To our knowledge, no other current theory can make predictions consistent with these patterns of data.

## HOW DOES MTS RELATE TO OTHER VIEWS OF SENTENCE PROCESSING?

Currently, most researchers who study sentence processing study "the part of comprehension in which a syntactic representation is recovered from the input" (MacDonald, 1997, p. 122). Recovering the syntactic representation of a sentence is assumed to be a crucial step performed by the language processing system on the way to understanding meaning (but see Bever, Sanz, & Townsend, 1998). This assumption often leads researchers to focus on syntactic processes to the exclusion of semantic processes and to design experiments without consideration of the meanings of the words, constructions, and sentences used in the experiments. In contrast, the goals of the MTS view are to bring meaning to the forefront in sentence processing research and to discover whether there are semantic explanations for phenomena most often considered syntactic.

One class of current theories rests on the assumption that syntactic processing takes place in an encapsulated module separate from other cognitive systems (Frazier, 1987; Frazier & Clifton, 1996; Frazier & Rayner, 1982; Rayner, Carlson, & Frazier, 1983). This hypothesis guided early sentence processing research in general and the study of reduced relative clauses in particular. Frazier (1987) proposed that the syntactic processing system is designed to minimize cognitive effort, preferring those syntactic structures that place the least burden on resources. For sentences that are structurally ambiguous between reduced relative clause and main clause readings, the main clause choice has the simplest structure and so is always the first choice. If it is the wrong choice, then processing must backtrack and the input must be reanalyzed.

The goal of MTS is to foster the development of explanations of sentence processing phenomena that emphasize meaning rather than syntax. Nevertheless, a syntactic theory of the class just described could be consistent with the MTS view and the data presented in this article. It could be that there is a syntax module that parses sentences into syntactic structures that are then forwarded to processes that make use of the semantic information postulated by MTS. Under this scenario, the structural ambiguity between a reduced relative and a main clause would be resolved for all reduced relatives by the syntactic system, leaving other processes to pursue the elements of meaning defined by MTS. The reduced relative clause structure might be a bit more difficult for the syntactic system than a main clause structure, but the difference in difficulty might be so small as to be unobservable empirically (see Frazier, 1995, who makes the same point). The extreme difficulty of comprehending reduced relatives of the type *The horse raced past the barn* would be attributed not to the syntactic module but instead to the processes that make use of semantic information of the type delineated by MTS. A scenario like this, with a separate fast modular syntax underlying semantic processing, would not detract from MTS as a potential theory of how meaning is understood or as a source of empirical hypotheses.

The theory most like the MTS view is Stevenson's (Stevenson & Merlo, 1997) competitive attachment parser. In this system, syntactically relevant lexical information for verbs is represented in structural syntactic configurations. Like MTS event templates, the configurations are decompositions, but instead of laying out semantic event structures, they lay out syntactic structures. Verbs differ in the number and kind of their syntactic configurations. Verbs of the class that we call "externally caused change of state" have two lexical configurations, one for intransitives like *the vase broke* and one for transitives like *John broke the vase*. Both structures place *the vase* in a theme syntactic position and the transitive structure places *John* in an agent syntactic position. Manner of motion verbs have only a single lexical configuration, with *the horse* in *the horse raced* in an agent syntactic position. When manner of motion verbs are used in apparently causal transitive sentences like *John raced the horse*, the causal sentence is formed by extension from *race*'s single lexical configuration and *the horse* still fills an agent syntactic position. In the competitive attachment parser, the reason reduced relatives like *The horse raced past the barn fell* are so difficult is that the reduced relative requires the verb in the clause to provide a theme syntactic position for the entity that is the head of the clause. The lexical structure of manner of motion verbs cannot provide this for entities like *the horse* because *the horse* holds an agent position. Thus, the parser fails to successfully parse reduced relatives like *The horse raced past the barn*.

The competitive attachment parser shares several key features with the MTS view. Syntactically relevant information about verbs is represented lexically in decomposition structures. The lexical structures are different for verbs of different classes, reflecting and determining the syntactic environments in which they appear. For reduced relative clauses, the structure required by the clause clashes with the structure provided by manner of motion verbs, thus rendering reduced relatives like *The horse raced past the barn* unparsable.

However, there are also key differences between the two systems. One is that MTS postulates that there is only a single basic event template for a verb (unless the verb is truly ambiguous). It is this assumption that explains why verbs with simpler templates are processed more quickly than verbs with more complex templates in both transitive and intransitive sentences and in lexical decision. With multiple representations, one for each of a verb's possible syntactic constructions, the competitive attachment parser cannot easily account for the speed-of-processing data.

A second difference is that the reduced relative difficulty for manner of motion verbs is caused by their lexical syntactic structure in the competitive attachment parser. In MTS, it is caused by their event template plus the template's semantic, internal causality feature. In MTS, internally caused change of state verbs share the internal causality feature with manner of motion verbs and so both are predicted to be difficult in reduced relative clauses. But in the competitive attachment parser, internally caused change of state verbs do not have the syntactic structure that gives the reduced relative difficulty. Thus, the competitive attachment parser could not predict the low probability of occurrence of reduced relatives for the internally caused change of state verbs.[7]

An important class of sentence-processing theories are the constraint-based theories (e.g., MacDonald et al., 1994; McRae et al., 1998; Tabor, Juliano, & Tanenhaus, 1997; Trueswell et al., 1994). The variables that predict sentence processing difficulty for many of these theories are different from those postulated by MTS, and they are different in empirically testable ways, as we show below. Constraint-based theories are connectionist theories in which multiple kinds of information ("constraints") all interact simultaneously during sentence processing. All the kinds of information are stored together in the lexicon. With each word is stored its own syntactic, semantic, discourse, and general knowledge information. For an input string of words, the various sources of

[7] In the competitive attachment parser, the syntactic structure for manner of motion verbs is an unergative structure. Although Stevenson and Merlo (1997) do not discuss internally caused change of state verbs, it can be inferred that they would not classify them as unergative because they do not show the two critical characteristics of unergative verbs that Stevenson and Merlo list: Internally caused change of state verbs almost never appear in the *X's way* construction and they almost never appear with reflexive objects.

information for all the words interact to produce the most likely syntactic representation of the input. The syntactic possibilities are ranked by their probabilities of occurrence in the language, and some of the sources of information that contribute to deciding which possibility is the most likely are more strongly weighted than others.

Many of the constraint-based models of sentence processing that have appeared in the literature are in conflict with MTS in three ways. First, in these models, processing is heavily influenced by information encoded in the linguistic processing system about the frequencies of occurrence of linguistic structures. In MTS, frequency information (beyond the frequencies of individual words) plays no part in the processing that forms meaningful representations from input sentences. Second, in many constraint-based theories, syntactic information is encoded in the lexicon separately and independently from semantic information, for example, in subcategorization frames for syntactic information and thematic role lists for syntactically relevant semantic information. In MTS, there are no separate representations of the syntactic structures for words. Instead, there are only semantic event templates to encode syntactically relevant information (with general linking rules to project event template information to sentence structures). In addition, MTS has only a single lexical event template for a verb (unless it is truly ambiguous). In many constraint-based models, for each verb, there are multiple subcategorization frames and thematic role lists to represent all the possible sentence environments in which the verb could occur. Third, constraint-based models do not include construction meanings as a source of parsing information.

Because MTS is a new and therefore relatively untested view and because the constraint-based theories with which it is in conflict are older, more established, and more numerous, we conducted a series of studies to attempt to distinguish between them. In corpus studies, we calculated the frequencies of occurrence of structures that have been claimed by constraint-based models to be important determinants of sentence processing difficulty in the reduced relative domain. Then we examined whether these frequencies are significant predictor variables for the reading times observed in experiments. To anticipate, we found that they are not. In the next sections, we first outline a generic constraint-based model and then proceed to describe the corpus studies and the experiments.

## CONSTRAINT-BASED MODELS

Many constraint-based theories of sentence processing make roughly the same assumptions about the representation and processing of linguistic information (e.g., MacDonald et al., 1994; McRae et al., 1998; Trueswell et al., 1994). They assume that there are constraints based on lexical, syntactic, argument role, general knowledge, and discourse information. For the most part, for reduced relative clauses, the theories have concentrated on constraints relevant to verbs because the structural ambiguity between a main clause and a reduced relative clause arises from the fact that the past tense form of the verb in a main clause is, for most verbs, the same word as the participle form of the verb in a reduced relative clause. For example, the string *The speaker proposed* is ambiguous between a main clause and a reduced relative because the past and participle forms of the verb are both the word *proposed*.

Table 5 shows five verb-relevant constraints plus one clause constraint, constraints said by constraint-based models to contribute to a decision about whether an input string of words is the beginning of a main clause or a reduced relative clause. The table shows examples of the constraints with the verb *cooked*. For each of the verb constraints, it is assumed that the processing system knows the probability of the verb being used in the possible structures. For example, the first constraint in the table is argument roles, which is the form in which lexical semantic information relevant to verb structures is often represented. The processing system is assumed to know the probabilities of *cooked* appearing with both an agent and a patient, an agent alone, and a patient alone. All three could be used in a main clause, but only the agent–patient and patient structures are possible in a reduced relative. (Constraint-based models use argument roles on an informal "as needed" basis; they do not define an exact list of all the possible roles for the verbs of a language nor do they give explicit definitions of each of the individual roles.)

The other three kinds of lexical information the system is assumed to know about a verb are the probabilities of it being used in transitive form, in active form, and in its past versus participle forms. The system is assumed to know the probabilities of each of these different forms occurring in the language. In a main clause like *John cooked the dinner*, the verb *cooked* is transitive, active, and past. In a reduced relative like *the dinner cooked by John*, the verb *cooked* is transitive, passive, and a participle.

General knowledge, real world information usually enters constraint-based models as information about the typicalities of combinations of words. There are many combinations of words that could be considered in constraint-based theories to be relevant to the main-clause/reduced-relative ambiguity, but never have been mentioned. One is the combination of a noun and its modifiers. The corpus study described earlier showed a difference between main clause subjects and reduced relative heads in whether they are definite. Another possibility is the combination of a verb with a *by* phrase that introduces an agent. Some passive verbs frequently appear with such *by* phrases (e.g. *captured*) and some almost never do (e.g. *armed*).

The combinations of words that have been singled out for attention by constraint-based theories are the combinations of verbs and the nouns that fill the verbs' agent and patient argument roles. This constraint is often labeled the "thematic fit" between a verb and a noun in a particular argument role, or the "typicality" of the noun in the argument role (the constraint is labeled "typicality" in Table 5). This constraint has been treated in two ways. Sometimes it has been treated as a general constraint across verbs (e.g., Trueswell et al., 1994): An agent is prototypically an animate entity and so, for all verbs, a noun better fills an agent role if it denotes an animate entity, and a patient is prototypically inanimate and so an inanimate entity better fills the patient role. The other way the constraint has been treated is as a constraint that is specific for each verb–noun pair. For each noun in the lexicon paired with each verb in the lexicon, the system is assumed to know the frequency with which the noun occurs as an agent of the verb and the frequency with which the noun occurs as a patient of the verb.

Of all the kinds of information relevant to the main-clause/reduced-relative ambiguity that a constraint based system is as-

Table 5
*Constraints in Constraint-Based Models*

| Constraints | Main clause | Reduced relative |
|---|---|---|
| Argument roles | | |
| Agent–patient | John (had) cooked the dinner. | The dinner cooked by John . . . |
| | The dinner was cooked by John. | |
| Agent | John (had) cooked last night. | |
| Patient | The dinner (had) cooked slowly. | The dinner cooked last night . . . |
| | The dinner was cooked slowly. | |
| Transitivity | | |
| Transitive | John (had) cooked the dinner. | The dinner cooked by John . . . |
| | The dinner was cooked by John. | The dinner cooked last night . . . |
| | The dinner was cooked slowly. | |
| Intransitive | John (had) cooked last night. | |
| | The dinner (had) cooked slowly. | |
| Voice | | |
| Active | John (had) cooked the dinner. | |
| | John (had) cooked last night. | |
| | The dinner (had) cooked slowly. | |
| Passive | The dinner was cooked by John. | The dinner cooked by John . . . |
| | The dinner was cooked slowly. | The dinner cooked last night . . . |
| Tense | | |
| Past | John cooked the dinner. | |
| | John cooked last night. | |
| | The dinner cooked slowly. | |
| Participle | John had cooked the dinner. | The dinner cooked by John . . . |
| | The dinner was cooked by John. | The dinner cooked last night . . . |
| | The dinner was cooked slowly. | |
| | John had cooked last night. | |
| | The dinner had cooked slowly. | |
| Typicality | | |
| Agent = animate | John cooked the dinner. | The dinner cooked by John . . . |
| Agent = inanimate | The charcoal cooked the dinner. | The dinner cooked by charcoal . . . |
| Syntactic structure | | |
| Main clause | John (had) cooked the dinner. | |
| | The dinner was cooked by John. | |
| | John (had) cooked last night. | |
| | The dinner was cooked slowly. | |
| Reduced relative | | The dinner cooked by John . . . |

sumed to know, the only one not tied to the verb of the clause is the last constraint shown in Table 5, the frequencies with which syntactic clause structures appear in the language. Generally, only the frequencies of main clauses and reduced relative clauses are considered.

The constraint-based models that have been tested empirically are connectionist models that are designed to choose which of several candidate syntactic structures is correct for a sentence (e.g., McRae et al., 1998; MacDonald et al., 1994, suggest how syntactic structures might actually be built, but Frazier, 1995, points out the difficulties of such an enterprise). To model results from experiments that investigate the reduced relative–main clause ambiguity, a constraint-based system is designed to decide whether an input string of words is a reduced relative clause or the subject and main verb of a sentence. Given the words in the input string and the probabilities of occurrence of each word in its various possible structures and forms (past, active, intransitive, main clause, etc.), activation reverberates among the constraints, according to their frequencies of occurrence in the language and the weights on the connections among them, to yield a decision about which is the most likely structure. For example, for the input string *the dinner cooked*, if *cooked* is more frequently transitive than intransitive, active than passive, past than participle, and main clause than reduced relative, then the main clause structure will be most highly activated and therefore chosen as the most likely syntactic structure for the string. However, if the input string continues *the dinner cooked by John*, and if the word *by* followed by an animate entity signals an agent with high probability, then activation at this point in the input string will begin to disfavor the main clause structure, and given the appropriate relative probabilities and weights for the various structures and forms, perhaps begin to favor the reduced relative structure.

## Tests of Constraint-Based Models: Corpus Study 7

The central proposal embodied in the constraint-based models under discussion here is that the difficulty with which a sentence is parsed is determined by statistical information stored in the language processing system. The stored information includes the likelihoods of occurrence of all the various structures listed in Table 5. In the MTS view, the main determinants of processing difficulty are semantic, not statistical. To test these views against each other, we carried out several investigations of corpus sentences.

To investigate how the statistics of actual language usage fit with constraint-based models, we examined sentences from the

180-million-word corpus. We reasoned that the probabilities of occurrence assumed to be known by a constraint-based processing system must be derived from sentences that are input to the system, sentences like those in the corpus. Of course, the corpus is not a perfect sample of all the sentences that would be input to the human language processing system, but it is a reasonable approximation.

For each of 85 verbs, we examined 200 randomly chosen tokens of the verb in its *-ed* ending form. Of the 85 verbs, 73 came from studies of reduced relative clauses by MacDonald (1994), McRae et al. (1998), and Trueswell et al. (1994), and the other 12 were used in Experiments 8 and 9 reported below. All 85 verbs were verbs with templates that provide the appropriate external causality and so, according to MTS, the verbs can appear in reduced relative clauses. Table 6 shows the mean probabilities of occurrence of the verbs in each of the structures from Table 5. For example, the 85 verbs occurred mainly in transitive constructions (probability .94). The verbs were as follows: *accepted*, *accused*, *acquired*, *admired*, *affected*, *approved*, *armed*, *arrested*, *assigned*, *associated*, *attacked*, *banned*, *captured*, *carried*, *certified*, *chased*, *concealed*, *considered*, *convicted*, *cured*, *devoured*, *dismissed*, *entertained*, *evaluated*, *examined*, *expected*, *frightened*, *graded*, *guarded*, *harmed*, *haunted*, *hired*, *hit*, *honored*, *housed*, *hurt*, *hypnotized*, *identified*, *instructed*, *interrogated*, *interviewed*, *investigated*, *invited*, *kicked*, *lectured*, *lifted*, *moved*, *observed*, *overlooked*, *paid*, *painted*, *praised*, *proposed*, *punished*, *pushed*, *questioned*, *read*, *recognized*, *recorded*, *replaced*, *requested*, *rescued*, *reviewed*, *searched*, *selected*, *sent*, *sentenced*, *served*, *sketched*, *slaughtered*, *sold*, *studied*, *supervised*, *surrendered*, *taught*, *terrorized*, *tortured*, *towed*, *transported*, *unloaded*, *visited*, *washed*, *watched*, *witnessed*, and *worshiped.*

Table 6
*Constraints in Constraint-Based Models*

| Constraints | Frequencies of occurrence |
|---|---|
| Argument roles | |
| Agent–patient | .50 |
| Agent | .06 |
| Patient | .44 |
| Transitivity | |
| Transitive | .94 |
| Intransitive | .06 |
| Voice | |
| Active | .45 |
| Passive | .55 |
| Tense | |
| Past | .35 |
| Participle | .65 |
| Typicality | |
| Agent = animate | .69 |
| Patient = animate | .46 |
| Syntactic structure | |
| Main clause | .92 |
| Reduced relative | .06 |
| Nonreduced relative | .02 |

*Note.* The frequencies of occurrence are based on 200 tokens for each of 85 verbs. For agent typicality frequencies, only transitive, active, main clause tokens were used. For patient typicality frequencies, only transitive, passive, main clause tokens were used.

## *Two Observations*

Before using the data summarized in Table 6 to test constraint-based models, we make two observations. First, constraint-based models are designed to explain why reduced relative clauses might often be difficult to process. The data in Table 6 point to where the difficulty must lie. Of all the probabilities for all the constraints, the one that most strongly disfavors the reduced relative is the clause structure constraint: The likelihood of a reduced relative is only .06, compared with a .92 likelihood of a main clause. If the head of the reduced relative denotes an animate entity, the typicality constraint also disfavors the reduced relative: The head of a reduced relative is the patient of the verb in the clause, and the frequency of an animate patient is only .46. All of the other constraints favor the reduced relative: For the argument role, transitivity, tense, and voice constraints, the form of the verb that is required by a reduced relative clause (agent–patient or patient argument roles, transitive, passive, participle) is the most frequently occurring form of the verb.

The second observation concerns the question of exactly how to count corpus tokens to determine the probabilities of the various constraints. For the data in Table 6, the probabilities for one constraint are assumed to be independent of the probabilities for all the other constraints. For example, we did not count the probability of an agent–patient role structure for active sentences separately from the probability of an agent–patient role structure for passive sentences. Constraint-based models have not explicitly addressed this issue, but it is clearly important. An example of its importance is provided by the typicality constraint. In applications of constraint-based models, MacDonald (1994), McRae et al. (1998), and Trueswell et al. (1994) attributed differences in the difficulty of processing for two sets of reduced relative clauses to differences in the typicalities of their heads as patients. The reduced relatives with heads that were typically patients were easier to process than the reduced relatives with heads that were typically agents. But MacDonald (1994), McRae et al. (1998), and Trueswell et al. (1994) measured the typicality of nouns as patients with passive sentences, not reduced relatives. In the corpus, the probabilities of animate entities appearing as patients are different in passive sentences than reduced relatives. In passive main clauses, the probability of the patient being animate is .46 but in reduced relatives it is only .38, a significant difference, $F_2(1, 84) = 30.8$. From the MTS point of view, this difference points to a difference in construction meaning between the passive construction and the reduced relative construction (although exactly what that meaning difference is we have not explored). For consistency with previous research, in Table 6 the agent and patient frequencies were calculated in the same way as has been done for constraint-based models: The patient frequency was counted from past tense, passive main clauses and the agent frequency was counted from past, active, transitive main clauses.

## *Are Reduced Relative Clauses Less Frequent Than Nonreduced Relative Clauses?*

Generally speaking, it is reasonable to suppose that sentences that are difficult to comprehend are less likely to be produced. This is not absolutely necessary, but it is at least a reasonable supposition. Given this assumption, then constraint-based models can be tested against sentences that are naturally produced by speakers and writers.

According to constraint-based models, reduced relative clauses are difficult to process because of their ambiguity with main clauses. Main clauses are much more frequent and so a constraint-based system is biased against reduced relatives. Nonreduced relatives, on the other hand, are not ambiguous and so, on average, they should be easier to comprehend. Easier comprehension is the only thing that distinguishes nonreduced from reduced relative clauses in constraint-based models because the two constructions are assumed to have the same meaning.

If speakers and writers tend to avoid constructions that are difficult to comprehend, and if reduced relatives are more difficult to comprehend than nonreduced relatives, then reduced relatives should appear in the corpus less frequently than nonreduced relatives. However, the opposite is true (see Table 6): Reduced relatives are more frequent. The probability of occurrence for reduced relatives is .06 (913 sentences) versus .02 (323 sentences) for nonreduced relatives; as was mentioned above, this is a significant difference, $F_2(1, 84) = 31.6$.

### *Are Reduced Relatives With Animate Heads Less Frequent Than Nonreduced Relatives With Animate Heads?*

Perhaps, it might be argued, the reduced relative clauses in the corpus, reduced relatives produced naturally by speakers of the language, are not difficult to process because the clauses fit constraints in such a way that difficulty is reduced. MacDonald (1994) has pointed out how naturally produced reduced relatives can be so easy to process as not to be noticed even by psycholinguistic researchers. Reduced relatives with inanimate entities as their heads are supposed to be easier to process than reduced relatives with animate heads, so perhaps all, or almost all, of the corpus reduced relatives have inanimate heads. But this is not the case: 38% of the reduced relatives from the corpus sample have animate heads. Moreover, there are not significantly fewer animate heads for reduced relatives than for nonreduced relatives; the nonreduced relatives have slightly fewer (37%).

Perhaps, however, the difficulty of a reduced relative clause with an animate head varies substantially across verbs. If so, it should be the case that if a verb with an animate head is especially difficult in a reduced relative, then a nonreduced relative clause should be used instead. For example, if *captured* is difficult to process in a reduced relative with an animate head, then when *captured* appears in a relative clause, it should be a nonreduced relative. That is, there should be a negative correlation between the probability of a verb appearing with an animate head (vs. an inanimate head) in a reduced relative and the probability of the verb appearing with an animate head (vs. an inanimate head) in a nonreduced relative. However, the correlation is not negative, it is highly positive ($r = .73$).

### *Are Reduced Relatives Less Likely to Appear in Subject Positions of Sentences Than Nonreduced Relatives?*

To calculate the statistics just described, all instances of reduced and nonreduced relative clauses were counted. However, reduced relative clauses are ambiguous with main clauses only when the clause is in a subject position in its sentence. For example, consider the sentence *The subject most frequently debated by the Fergie followers on Fleet Street is the royal figure*. The subject of the sentence is *the subject most frequently debated by the Fergie followers on Fleet Street*; it is the subject of *is the royal figure*. It might be argued that reduced relatives should be difficult only in this situation, not when they occur in object positions where there is no ambiguity (e.g., *The royal figure is the subject most frequently debated by the Fergie followers on Fleet Street*). (It should be noted that it is not obvious that a processing system could always know for sure whether a noun phrase that was potentially the head of a reduced relative clause was in a subject position or an object position, but we leave this problem aside.)

We examined the subset of relative clauses that appear in subject positions in the corpus sentences. Counter what would be predicted from constraint-based models, the probability of appearing in subject position (vs. other sentence positions) is higher for reduced than nonreduced relative clauses, .27 versus .15, respectively. Also, for the clauses that appear in subject positions, we calculated the probabilities of animate heads: For reduced relatives, the probability was .56 ; for nonreduced relatives, the probability was .60. The probability of an animate head for a reduced relative was higher when the reduced relative appeared in subject position (.56) than it was for all reduced relatives combined (.38, see above). (We did not perform statistical tests with the probability values for relative clauses in subject positions because the values are zero for many verbs.)

In sum, just as when all reduced and nonreduced relatives were counted, the statistics for relative clauses in subject position are not what would be expected from a constraint-based point of view. The supposedly difficult reduced relatives are not less frequent than the supposedly easier nonreduced relatives, and they are not less likely to have animate heads.

### *Reading Times for Sentences With Reduced Relative Clauses: Experiments 8, 9, 10, and 11*

To further test predictions of constraint-based models, reading times for sentences with reduced relative clauses were collected in Experiments 8–11. For constraint-based models of the type under discussion here, reading times for sentences with reduced relative clauses are predicted from the likelihoods of occurrence in the language of the various linguistic structures in the reduced relative. If, for example, the verb of a reduced relative rarely occurs in participle form, then reading times for the reduced relative should be lengthened.

Three sets of sentences for the experiments were taken from previous studies of reduced relatives (MacDonald, 1994, Experiment 2; McRae et al., 1998; Trueswell et al., 1994) and two other sets were new sentences that we took from the corpus. Examples of the sentences are shown in Table 7.

In MacDonald's (1994) and Trueswell et al.'s (1994) tests of constraint-based models, experiments manipulated whether the head of a reduced relative was more likely to be an agent of the reduced relative's verb or a patient. For the sentence that begins *The speaker proposed by the group*, the noun *speaker* is an animate entity. In the other version of the sentence, *The solution proposed by the group*, the noun *solution* is inanimate. The assumption made by Trueswell et al. (1994) and MacDonald (1994) was that an

Table 7
*Sentences Used in Experiments 8–11*

| Experiment and sentence source | Head noun | Examples of sentences |
|---|---|---|
| Experiments 8, 9: Trueswell et al. (1994) | Animate | The speaker proposed by the group would work perfectly for the program. |
| | Inanimate | The solution proposed by the group would work perfectly for the program. |
| Experiment 8: The corpus | Animate | The special prosecutor requested by Edwin Meese will probably not be provided by the Justice Department until early next week. |
| | Inanimate | The special ruling requested by Edwin Meese will probably not be provided by the Justice Department until early next week. |
| Experiment 9: The corpus | Animate | When the fairies sent by the good witch flew in to break the evil spell, the children's singing and dancing began. |
| | Inanimate | When the snowflakes sent by the good witch flew in to break the evil spell, the children's singing and dancing began. |
| Experiment 10: MacDonald (1994) | Animate | The spy concealed in the passageway was discovered by the maid. |
| | Inanimate | The microfilm concealed in the passageway was discovered by the maid. |
| Experiment 11: McRae et al. (1998) | Animate, typical agent | The postman carried by the paramedics was having trouble breathing. |
| | Animate, typical patient | The newborn carried by the paramedics was having trouble breathing. |

animate entity is more likely to be an agent of a verb and an inanimate entity is more likely to be a patient. If this assumption is correct, then, according to the constraint-based models, reading times for the reduced relative with *speaker* as head should be longer than reading times for the reduced relative with *solution* as head. This is because the head of the reduced relative must be a patient of the reduced relative's verb, and the inanimate *solution* is more likely to fill a patient role than the animate *speaker*.

Whereas Trueswell et al. (1994) and MacDonald (1994) manipulated the typicality of nouns as agents and patients by using animate and inanimate entities, McRae et al. (1998) manipulated typicality with only animate entities, varying whether they are more typically agents or more typically patients (see Table 7).

For all three of these sets of sentences, the previous studies found that reading times for the sentences were consistent with constraint-based model predictions: When the head of the reduced relative was an entity typically an agent, reading times were longer than when the head was typically a patient. Experiments 8, 9, and 10 replicated this result for these three sets of sentences.

From the point of view of MTS, there is no reason to expect this result. All else being equal, there is no reason that an entity that is typically an agent should be more difficult to comprehend as the head of a reduced relative than an entity that is typically a patient.

To further explore typicality effects on reading times for reduced relatives, the two new sets of sentences were constructed. They were based closely on naturally produced sentences from the corpus. Surprisingly, from the context of the results for the other sets of sentences and constraint-based theories, we found that for these new sets of sentences whether the head of a reduced relative was animate versus inanimate had no significant effect on reading times.

The main portion of the analyses of the results of Experiments 8–11 is devoted to the question of why the agent versus patient typicality of the head of a reduced relative sometimes does affect reading times and sometimes does not. But first, before those analyses, the experiments and their data are presented.

### Experimental Design

The five sets of sentences encompass a wide range of different kinds of sentences. The sentences used by previous researchers are relatively short, and they were constructed by the experimenters. The new sentences from the corpus are longer, and they are closely based on sentences naturally produced by speakers and writers.

Experiments designed to examine the processing of reduced relative clauses often have two design features that our experiments did not have. First, reading times are usually measured for each word of a sentence; instead, we measured whole sentence reading times. We did this because our concern was not at exactly what point in a sentence difficulty occurred but only in the overall difference in difficulty between reduced relatives with one kind of head versus another. In other words, we were concerned with which variables can predict difficulty rather than at what point in a sentence difficulty occurs.

Second, the difficulty of a sentence with a reduced relative clause is usually measured relative to the difficulty of the same sentence with the clause turned into a nonreduced relative. This method of measuring difficulty depends on the assumption that nonreduced and reduced relative clauses differ only in the reduced relative being more difficult to parse syntactically, not in any other aspect of comprehension.

We think this assumption is wrong, that instead reduced and nonreduced relative clauses differ in their meanings, and we gave evidence earlier in this article of the differences. First, we showed that a naturally produced sentence with a reduced relative can be preferred by subjects to the same sentence with the reduced relative changed to a nonreduced relative, and vice versa. This finding indicates that whether reduced or nonreduced is preferred depends on how the construction interacts with the meaning of its sentential and discourse context. Second, we showed that subjects remember the contents of a relative clause better for reduced than nonreduced relatives (given recall of the head of the clause). Again, this indicates a difference in the encoded meaning of the two constructions.

If reduced and nonreduced relatives have different meanings, then nonreduced relatives cannot be used as a control condition for examination of the relative difficulties of parsing reduced relative clauses. However, as it turns out, resolution of this issue is less critical than it might be, because of the patterns of results from previous studies. In the standard experimental design used by MacDonald (1994), McRae et al. (1998), Trueswell et al. (1994), and others, each sentence comes in four versions: one with a reduced relative clause with a head that is an entity more typical as agent (e.g. *speaker*), one with a nonreduced relative clause with the same entity as head (e.g. *speaker*), one with a reduced relative clause with a head that is an entity more typical as a patient (e.g. *solution*), and one with a nonreduced relative clause with the same entity as head (e.g. *solution*). The usual finding is the following: The only significant differences among the reading times of the sentences are that reading times are longer for reduced relative sentences with typically animate entities as heads than for the other three types of sentences. What this means is that reading times for the reduced relative sentences with the two types of heads (*speaker* and *solution*) can be directly compared, and this is the design we adopted for Experiments 8–11.

For all four experiments, example sentences are shown in Table 7. For each sentence, there are two versions, one with an animate entity or a typically agent entity as head of the reduced relative and one with an inanimate entity or a typically patient entity as head. As in the example sentences, all of the reduced relative clauses were in subject positions in their sentences.

In each experiment, subjects read the sentences from a PC monitor, one sentence at a time, pressing the space bar key on the PC keyboard to advance to each next item. Comprehension tests were interspersed among the sentences to be sure subjects read to some reasonable level of comprehension.

## Method

*Materials.* The first set of 16 sentences, exemplified in Table 7 and used in Experiments 8 and 9, was from Trueswell et al.'s (1994) Experiment 2. The sentences averaged 9.8 words in length. The mean word frequency of the animate heads of the reduced relative clauses was 354, and the mean frequency for the inanimate heads was 132 (Francis & Kučera, 1982).

The second and third sets of items, used in Experiments 8 and 9, were two new sets of 16 pairs of sentences, with all sentences containing a reduced relative clause. For one of these sets of sentences, the first sentence of a pair was based on a corpus sentence for which the head of the reduced relative clause was a person or people, and the second sentence of the pair was identical except that the original head was replaced by something that was not animate. The sentences of this set averaged 18.0 words in length, and the mean frequencies of the animate and inanimate heads were 174 and 55, respectively (Francis & Kučera, 1982). For the other set of new sentences, the first sentence of a pair was also based on a sentence from the corpus, but the head was inanimate, and it was replaced by a person or people for the second sentence of the pair. The sentences of this set averaged 15.7 words in length, and the mean frequencies of the animate and inanimate heads were 57 and 52, respectively (Francis & Kučera, 1982). Examples of the sentences are listed below:

*A Welsh family/estate haunted by an ominous heritage is the focus around which the author builds the long saga of three generations.*

*Relief to farmers/hotels hurt by low prices and bad weather cost an average of 200 million dollars a year.*

*Cars/Guards assigned to nearly 200 officials have telephones to ensure instant communication.*

*The Palestinian headquarters/terrorists hit in the two raids caused serious rethinking of the peace process.*

*Documents/Scientists associated with the organization were reported Sunday to have arrived in the Soviet Union.*

The sentences of the fourth set of items, used in Experiment 10, were versions of sentences from MacDonald (1994; Experiment 2), shortened so that all the words of each sentence could appear on a single line of the PC screen used for stimulus presentation. There were 32 pairs of sentences. In MacDonald's sentences, the reduced relative clause could contain either a time adverbial or a location phrase; we used the versions with the location phrase. The shortened sentences averaged 9.8 words in length, and the mean frequencies of the animate and inanimate heads were 58 and 42, respectively (Francis & Kučera, 1982).

The sentences of the fifth set were used by McRae et al. (1998). The heads of the reduced relatives were always animate entities, but for one version of a sentence, this was an entity that is typically an agent and for the other version, it was an entity that is typically a patient (typicalities were determined by McRae et al. from subjects' ratings, as discussed below). These sentences averaged 10 words in length, and the mean frequencies of the typically agent and typically patient heads were 27 and 29, respectively (Francis & Kučera, 1982).

*Procedure.* For all four experiments, experimental and filler sentences to read and true–false test sentences were mixed together in blocks of 11 sentences to read and 6 true–false tests. There were 16 blocks, plus 1 practice block at the beginning of the experiment. Within a block, the sentences and true–false tests were presented in random order except that in Experiment 8, a true–false test was either 1, 2, 3, or 4 positions after the sentence for which it was the test item; in Experiments 9 and 11, a true–false test immediately followed the sentence for which it was a test; and in Experiment 10, a true–false test was 2, 3, or 4 positions after the sentence for which it was a test item. Subjects read faster in Experiment 9 than the other experiments, but the placement of the test phrases did not otherwise affect any results of interest.

Subjects controlled reading time for the sentences by pressing the space bar on the PC keyboard when they were finished with a sentence and ready for the next item in the block. Test phrases were marked with seven leading and seven following asterisks. Subjects were instructed to press the *?/* key if a test phrase had been in one of the sentences they had read and to press the *Z* key if the phrase had not been in one of the sentences.

Subjects were encouraged to read the sentences as quickly as possible while keeping their accuracy on the test phrases high. In Experiments 8, 9, and 10, subjects were given a point system: They were told that they started the experiment with 500 points. For each block, points were added or subtracted according to a scheme intended to weight accuracy over speed: If mean reading time was under 200 ms per word, 100 points were added; for each correct response to a test phrase, 100 points were added; and for each incorrect response, 100 points were subtracted.

*Subjects and design.* In Experiment 8, the corpus sentences were those that originally appeared in the corpus with people as the heads of their reduced relative clauses. Reading times for Trueswell et al.'s (1994) sentences with animate versus inanimate reduced relative heads were compared with reading times for the corpus sentences with animate versus inanimate heads. In Experiment 9, the corpus sentences were those that originally appeared in the corpus with inanimate entities as the heads of their reduced relative clauses. Reading times for Trueswell et al.'s sentences with animate versus inanimate heads were compared with reading times for the corpus sentences with animate versus inanimate heads. In Experiment 10, only the MacDonald (1994) sentences were tested; in Experiment 11, only the McRae et al. (1998) sentences were tested. In each experiment, the assignment of reduced relative head was counterbalanced across subjects. There were 16 subjects in Experiment 8, 22 subjects in Experiment 9, 42 subjects in Experiment 10, and 40 subjects in Experiment 11.

## Results: Reading Times

Subjects read at rates that approximated normal reading times for college students. Over all sentences (including fillers), reading times averaged 264 ms per word in Experiment 8, 212 ms per word in Experiment 9, 243 ms per word in Experiment 10, and 284 ms per word in Experiment 11. The mean probability of a correct response on a true–false test was .84 for Experiment 8, .90 for Experiment 9, .94 for Experiment 10, and .92 for Experiment 11.

Mean reading times for the sentences of interest, the sentences with reduced relative clauses, are summarized in Table 8. For the Trueswell et al. (1994) sentences (Experiments 8 and 9), the MacDonald (1994) sentences (Experiment 10), and the McRae et al. (1998) sentences (Experiment 11), reading times were significantly longer when the heads of the reduced relative clauses were typically agentive entities than when they were typically patient entities, replicating Trueswell et al.'s, MacDonald's, and McRae et al.'s findings. However, for the corpus sentences (Experiments 8 and 9), whether the head of the reduced relative was an animate or an inanimate entity did not significantly affect reading times.

In Experiment 8, the main effects of sentence set (Trueswell et al.'s, 1994, sentences vs. the corpus sentences)—$F_1(1, 15) = 283.7$, and $F_2(1, 30) = 134.2$—and animate versus inanimate entities as heads of the relative clauses—$F_1(1, 15) = 10.8$, and $F_2(1, 30) = 5.4$—were both significant. The interaction between the two factors was also significant, $F_1(1, 15) = 5.1$, and $F_2(1, 30) = 4.6$, $SE = 104$ ms. In Experiment 9, the main effect of sentence set was significant, $F_1(1, 21) = 182.3$, and $F_2(1, 28) = 94.9$, and the main effect of animate versus inanimate entities was not, $F_1(1, 21) = 3.4$, and $F_2(1, 28) = 2.1$. The interaction of the two factors was significant, $F_1(1, 21) = 5.3$, and $F_2(1, 28) = 5.9$, $SE = 70$ ms. In Experiment 10, the main effect of animate versus inanimate heads was significant, $F_1(1, 41) = 24.1$, and $F_2(1, 31) = 12.6$, $SE = 28$ ms. In Experiment 11, the main effect of typically agent versus typically patient head was significant, $F_1(1, 39) = 31.5$, and $F_2(1, 39) = 18.0$, $SE = 53$ ms.

The finding that an animate versus an inanimate head made no significant difference in reading times for the corpus sentences is surprising from a constraint-based modeling point of view. In response to the finding, it might be argued that the animate–inanimate difference was somehow lost in the long reading times for the corpus sentences. (The long reading times came about because the corpus sentences were several words longer on average than the sentences of the other sets.) To check this possibility, we ran an additional experiment that included both sets of corpus sentences and the set of sentences from Trueswell et al. (1994). In this experiment, 15 subjects were asked to press the space bar after each line of an experimental sentence (rather than at the end of each sentence) so that reading times could be recorded for the first lines of the corpus sentences. For each of the corpus sentences, the first line included all of the reduced relative clause plus the next two words of the sentence. The next two words included the main verb of the sentence. The pattern of results was the same for this experiment as for Experiments 8 and 9. The mean reading times for corpus sentences with animate and inanimate heads were 2,391 ms and 2,427 ms, respectively, compared with 2,645 ms and 2,333 ms, respectively, for the Trueswell et al. sentences. The interaction of sentence sets with animate–inanimate heads was significant, $F_1(1, 14) = 6.4$, and $F_2(2, 45) = 5.3$, $SE = 67$ ms. Thus, the failure to find a significant animate–inanimate difference for the corpus sentences was not due to the long reading times caused by their length.

For the Trueswell et al. (1994) and MacDonald (1994) sets of sentences, the sizes of the differences observed in Experiments 8–10 are about the same as were observed in the previous studies with the sentences. For the Trueswell et al. sentences, the difference they found in mean reading times between reduced relatives with animate versus inanimate heads was about 400 ms (using an eyetracker, they reported both first pass and second pass reading time differences; 400 ms is about the sum of the two). The difference in our experiment was 481 ms. MacDonald (using self-paced word-by-word reading) reported the mean reading time per word difference between reduced relatives with animate heads versus inanimate heads: The mean was about 10 ms. In our experiments, the mean difference per word was 15 ms (the slightly shorter mean reading times in MacDonald's study probably are due to the sentences in her study being longer; we shortened the sentences so that the whole sentence would appear as a single line on a PC screen).

For the McRae et al. (1998) sentences, mean reading times for reduced relatives with typical agents as heads and reduced relatives with typical patients as heads were reported for only some, not all, of the words of the sentences (McRae et al., 1998). Reading times for the words they did report averaged about 500 ms per word, much longer than the reading times we found. This is likely due to their method, self-paced word-by-word reading, which often leads to slow reading.

Table 8
*Sentence Reading Times (in Milliseconds)*

| Experiment | Sentence set | Head of reduced relative: Animate | Head of reduced relative: Inanimate |
|---|---|---|---|
| 8 | Trueswell et al. (1994) | 3,121 | 2,640 |
| | Corpus, Set 1 | 5,154 | 5,133 |
| 9 | Trueswell et al. (1994) | 2,407 | 2,084 |
| | Corpus, Set 2 | 4,091 | 4,095 |
| 10 | MacDonald (1994) | 2,844 | 2,653 |
| | | Typical agent | Typical patient |
| 11 | McRae et al. (1998) | 3,056 | 2,639 |

## Reading Times, Corpus Statistics, and Typicality Ratings

The question immediately raised by the results of the four experiments is why the typicality of the head of the reduced relative affects reading times for three of the sets of sentences but not the other two. According to constraint-based models of the type under discussion here, we should find the answer in the differences among the sentences in the likelihoods of occurrence of the elements of their reduced relative clauses. For example, the finding is that it takes longer to read a sentence that begins *The speaker proposed by the group . . .* than to read a sentence that begins *The solution proposed by the group . . .*, but it does not take longer to read *The fairies sent by the good witch . . .* than *The snowflakes sent by the good witch . . . .* According to constraint-based models, this should be because the difference in typicality as an agent for *speaker* versus *solution* is significantly larger than the difference in typicality as an agent for *fairies* versus *snowflakes*. In this section, we explore these likelihoods, but in the end, we find that the constraint-based predictions are not supported by the data. The constraints hypothesized by the models do not explain why there is a significant reading time difference due to typically agent versus typically patient heads for three of the sets of sentences but not for the other two sets. We conclude that the reasons for the observed differences in reading times must lie outside the domain of the constraint-based models.

### *Corpus Statistics*

In introducing constraint-based models, we laid out a number of constraints that have been claimed by the models to be relevant to the processing of reduced relative clauses. The likelihoods of occurrences, as calculated from the corpus, for the various linguistic constructs and argument role typicalities were shown in Table 6. The 85 verbs for which the statistics were calculated were all the verbs used in reduced relative clauses for the sentences in Experiments 8–11.

Table 9 shows statistics for the same constraints, this time broken down into the statistics for the verbs of each set of sentences individually. For each constraint, the mean probability of occurrence is shown and (in parentheses) the interquartile range of the probability values is shown. The ranges of probability values are reasonably wide for the argument role, voice, tense, and typicality constraints, indicating considerable variation in these characteristics across verbs. For the transitivity and syntactic structure constraints, the verbs are much more uniform, almost always being used in transitive, main clause sentences. Overall, looking at both the mean probability values and the ranges, there are few large differences between the statistics for the sets of sentences for which the typicality of the reduced relative head affected reading times (the sentences from MacDonald, 1994; McRae et al., 1998; and Trueswell et al., 1994) and those for which it did not (the sentences based on corpus sentences).

Consider, as an example, the statistics for the first set of corpus sentences and the statistics for the sentences from Trueswell et al. (1994). For the argument role, transitivity, voice, tense, and syntactic structure constraints, the statistics for the verbs of the reduced relatives of the two sets of sentences differ insignificantly, all $F_2(1, 30) < 1.0$.

For the typicality constraint, the statistics do differ significantly: The verbs of the reduced relatives in Trueswell et al.'s (1994) sentences are more likely to occur with animate agents than the verbs of the corpus reduced relatives (.77 vs. .58), $F_2(1, 30) = 5.9$, and they are less likely to occur with animate patients (.17 vs. .35), $F_2(1, 30) = 6.9$.

The differences in the typicality constraint are exactly what should drive the difficulty of reduced relatives, according to constraint-based models. For Trueswell et al.'s (1994) reduced relatives with animate heads, the high probability (.77) of the verbs taking animate agents should bias the parsing system strongly toward a main clause reading, away from the correct, reduced relative reading. That is, a high probability of the verb *proposed* occurring with an animate agent should bias the system toward a main clause reading of *The speaker proposed* . . . . At the same time, the low probability of the verbs occurring with animate patients (.17) should also drive the system toward a main clause reading, away from the correct, reduced relative reading. Thus, animate heads should make the reduced relatives quite difficult compared with inanimate heads. For the corpus sentences, on the other hand, the reduced relative verbs are not so highly biased (.58

Table 9
*Constraints in Constraint-Based Models: Mean Frequencies of Occurrence and Interquartile Ranges*

| | Sentence source | | | | |
|---|---|---|---|---|---|
| Constraints | Trueswell et al. (1994) | Corpus, Set 1 | Corpus, Set 2 | MacDonald (1994) | McRae et al. (1998) |
| Argument roles | | | | | |
| Agent–patient | .53 (.20) | .59 (.20) | .39 (.41) | .54 (.32) | .50 (.22) |
| Agent | .04 (.05) | .04 (.05) | .01 (.01) | .09 (.07) | .06 (.06) |
| Patient | .43 (.31) | .37 (.28) | .60 (.41) | .37 (.32) | .44 (.33) |
| Transitivity | | | | | |
| Transitive | .96 (.05) | .96 (.05) | .99 (.01) | .91 (.07) | .94 (.05) |
| Intransitive | .04 | .04 | .01 | .09 | .06 |
| Voice | | | | | |
| Active | .49 (.32) | .51 (.26) | .31 (.28) | .53 (.33) | .45 (.34) |
| Passive | .51 | .49 | .69 | .47 | .55 |
| Tense | | | | | |
| Past | .38 (.25) | .39 (.27) | .23 (.19) | .40 (.26) | .35 (.31) |
| Participle | .62 | .61 | .77 | .60 | .65 |
| Typicality | | | | | |
| Agent = animate | .77 (.20) | .58 (.45) | .59 (.33) | .69 (.28) | .74 (.24) |
| Patient = animate | .17 (.21) | .35 (.44) | .42 (.26) | .27 (.38) | .66 (.70) |
| Syntactic structure | | | | | |
| Main clause | .92 (.05) | .92 (.04) | .84 (.11) | .94 (.05) | .93 (.05) |
| Reduced relative | .06 (.04) | .05 (.04) | .13 (.08) | .04 (.05) | .05 (.04) |
| Nonreduced relative | .02 (.01) | .01 (.01) | .03 (.04) | .02 (.01) | .02 (.02) |

*Note.* Values in parentheses represent the interquartile ranges of the probability values.

Table 10
*Correlations Between Corpus Statistics, Typicality Ratings, and Reading Time Differences*

| | Sentence source | | | | |
|---|---|---|---|---|---|
| Correlations | Trueswell et al. (1994) | Corpus, Set 1 | Corpus, Set 2 | MacDonald (1994) | McRae et al. (1998) |
| Corpus statistics | | | | | |
| A: Probability of animate agent | −.33 | −.41 | .23 | −.28 | |
| B: Probability of animate patient | .24 | .30 | .02 | −.42 | |
| A − B | −.37 | −.39 | .15 | .17 | |
| Typicality ratings | | | | | |
| A: Typicality of animate agent − typicality of animate patient | −.18 | −.53 | −.50 | −.37 | .29 |
| B: Typicality of inanimate agent − typicality of inanimate patient | .24 | .09 | −.57 | −.46 | −.15 |
| A − B | −.31 | −.53 | .19 | −.03 | .27 |

*Note.* "Correlations" refers to reading time differences between reduced relatives with animate heads versus reduced relatives with inanimate heads correlated with the entries in the first column. All the sets of sentences except McRae et al. compared animate heads with inanimate heads. McRae et al. compared typically agent heads with typically patient heads.

and .35), so the difference in difficulty between the reduced relatives with animate and inanimate heads should be reduced.

If these statistical differences are the reason for the differences in the patterns of reading times for the two sets of sentences, then the statistics should correlate with the reading time differences. For both Trueswell et al.'s (1994) sentences and the corpus sentences, the more strongly biased a verb is toward an animate agent and away from an animate patient, the larger the reading time difference should be between sentences with animate heads for their reduced relatives and sentences with inanimate heads. For both sets of sentences, there should be significant correlations between the biases of the verbs and the animate–inanimate reading time differences. However, the correlations were not significant.

Table 10 (under "Corpus statistics") shows the relevant correlations for the Trueswell et al. (1994), MacDonald (1994), and corpus sentences. According to constraint-based models, the size of the reading time difference should increase as the probability of the verb occurring with an animate agent increases (a positive correlation), it should decrease as the probability of an animate patient increases (a negative correlation), and it should increase as the difference between these two probabilities increases (a positive correlation). In only four cases are the correlations in the predicted direction, and only one of these is significant: the negative correlation in Row B for the MacDonald sentences, $t(30) = -2.5$; all other $t$s $< 1.0$. Figure 1 shows scatter plots for these same correlations, with reading time per word plotted on the vertical axes and the corpus probabilities on the horizontal axes. The ranges of the probability values are wide enough that it would be possible to see the predicted correlations if they had been present in the data.[8]

Most of the corpus statistics listed in Table 9 have not been investigated by constraint-based researchers because corpora large enough to gather meaningful statistics have not been available. The exception is verb tense and voice, which were jointly tested by Trueswell (1996) for their abilities to predict the difficulty of reduced relative clauses. He found that reduced relatives with verbs highly likely to occur in their past active forms were more difficult than reduced relatives with verbs less likely to occur in their past active forms. In the corpus statistics in Table 9, there are no apparent differences in the voice and tense statistics between those sets of sentences that produced an animate–inanimate reading time difference and those that did not. Nevertheless, we conducted an additional analysis to examine the voice and tense constraints. For the Trueswell et al. (1994) sentences, the MacDonald (1994) sentences, and both sets of corpus sentences, we computed the mean probabilities with which the verbs occurred in the active voice in the corpus, and then divided the sentences into those for which the probability was higher than the mean and those for which it was lower. A verb more likely to occur in the active voice should be more difficult in a reduced relative (where it occurs in the passive voice). Combining across the four sets of sentences, the mean numbers of words in the sentences were matched for the sentences with the high probability verbs and the sentences with the low probability verbs (12.8 and 12.5 words per sentence, respectively), as were the mean Kučera–Francis frequencies per word (293 and 300, respectively; the McRae et al. (1998) sentences were not used in this analysis because high–low probability sentences did not match in Kučera–Francis frequencies). Contrary to what would be predicted by constraint-based models, there was no significant difference in reading times for the sentences with high versus low probability verbs, $F_2(1, 158) < 1.0$. Verbs more likely to occur in the active voice did not present significantly more difficulty. The same result holds for tense. For the verbs in the four sets of sentences, it turns out that splitting them (at the mean) into those more likely versus those less likely to occur in the active voice is equivalent to splitting them into those more likely versus less likely to occur in their past forms rather than their participle forms. So neither tense nor voice probabilities significantly affected reading times in our experiments.

[8] To check that we had sufficient power to find significant correlations if they were present in the data, for each experiment, we correlated one half of the subjects' data against the other half; the correlations were always .48 or greater.

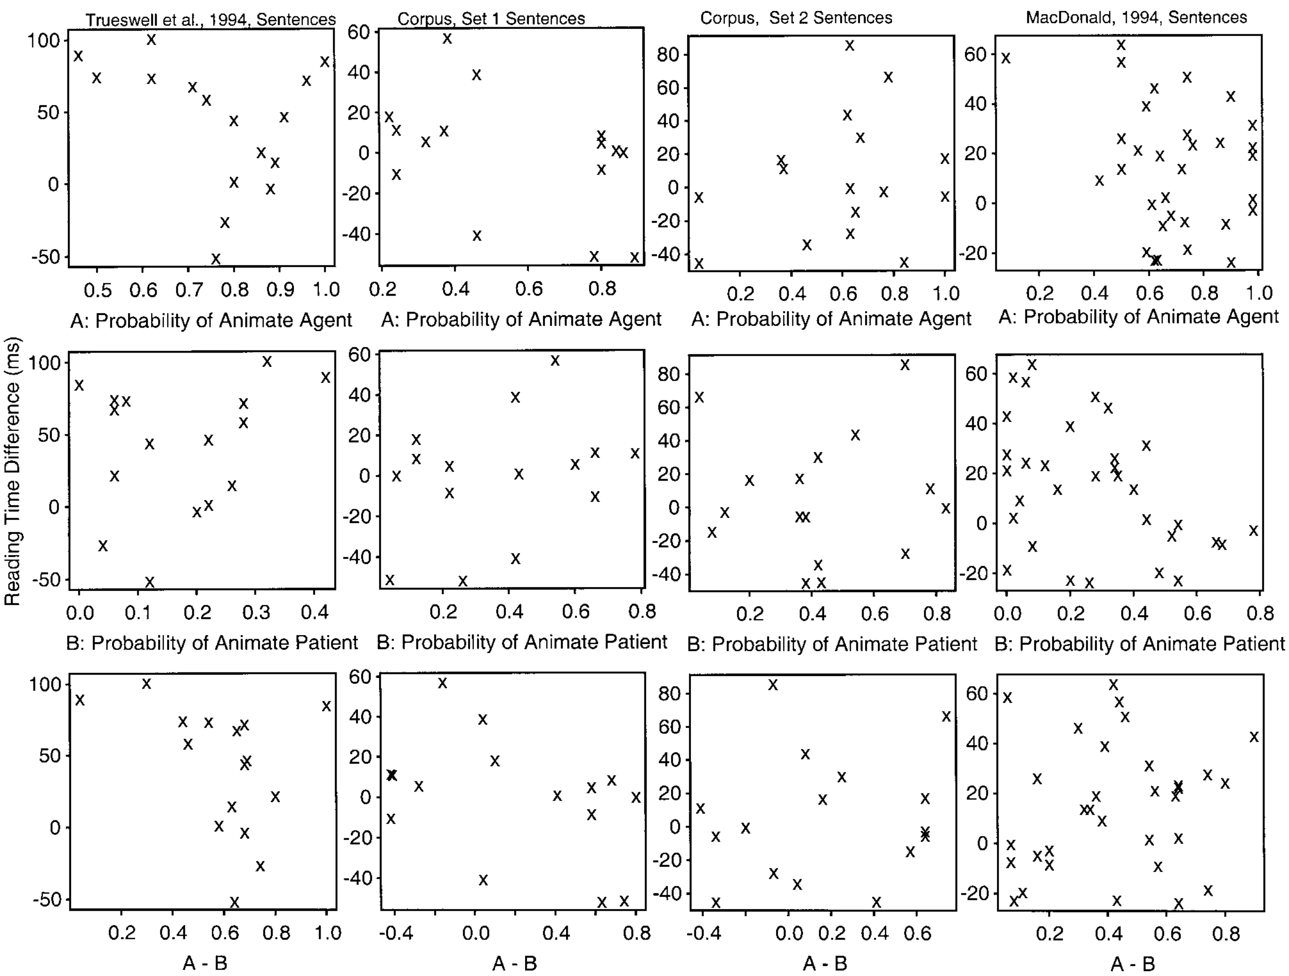


*Figure 1.* Scatter plots for correlations of reading time per word differences with corpus statistics.

This finding is contrary to Trueswell's (1996) result, so we looked for what might be the cause of the discrepancy. Some of the 20 verbs in the reduced relatives in Trueswell's study are among the 85 summarized in Table 9. For his two groups of verbs, one more likely to occur in active voice and the other less likely, we noticed a difference: For the group more likely to occur as active, they are also less likely to occur in reduced relative clauses. To determine whether this is true when all of his verbs are considered, we computed the probabilities of occurrence as reduced relatives for those verbs that had not been in the original 85 (using 200 randomly chosen sentences from the corpus for each verb). For the 10 verbs in Trueswell et al.'s (1994) study that are most likely to occur in active voice, the mean probability with which they occur in reduced relatives is .02; for the other 10 verbs, those with lower probabilities of occurring in the active voice, the mean is .09—a significant difference, $F_2(1, 18) = 12.4$. Thus, there was a large difference between the two groups of verbs not only in likelihood of occurring in active voice, but also in likelihood of occurring in reduced relatives. Because in our own data we found no significant effect of probability of occurrence in active voice, we suggest Trueswell's result may be due to the difference in reduced relative probability. However, we do not believe difficulty of reading in his study depended directly on reduced relative probability (as might be consistent with a constraint-based model). Instead, we suggest that the difficulty for reduced relatives with verbs unlikely to occur in reduced relatives was due to the difficulty of writing, for those verbs, reduced relative sentences that are consistent with other language processing factors such as discourse functions and real world knowledge. We take up this point later in the summary of findings relevant to constraint-based models.

### *Typicality Ratings*

In previous applications of constraint-based models (e.g., McRae et al., 1998; Tabossi, Spivey-Knowlton, McRae, & Tanenhaus, 1994; Trueswell et al., 1994), typicality values for the agents and patients of verbs were not calculated from large corpora, because none were available, but instead were calculated from typicality ratings collected from subjects. Subjects were asked to rate how common it is for some entity to do something to someone or something, and they were asked to rate how common it is for something to be done to someone or something by someone; for example, *How common is it for a speaker to propose someone?* and *How common is it for a speaker to be proposed by someone?*

To further test constraint-based models, we collected typicality ratings for the Trueswell et al. (1994) sentences, the MacDonald (1994) sentences, and the two sets of corpus sentences. (For the McRae et al., 1998, sentences, we used the typicality values published with their sentences.) All the pairs formed by combining each head noun of a reduced relative clause with the verb of the clause were rated using the same "how common is it" questions as have been used by constraint-based modelers. Forty subjects rated each pair on a scale of 1–7, with 7 being most common. The means of the ratings with their interquartile ranges are reported in Table 11. Trueswell et al. reported typicality ratings for the inanimate nouns for the sentences of theirs that we used in our experiments. To ensure that our ratings were consistent with earlier work, we calculated the correlation between our ratings for Trueswell et al.'s noun–verb pairs and their ratings, and the correlation was .97.

As an aside, it should be mentioned that typicality ratings do not produce exactly the same pattern of likelihoods for agents and patients that corpus statistics do (a point made earlier by Merlo, 1994). Table 11 shows both the mean typicality ratings and the corresponding corpus statistics (the same statistics as were reported in Table 9). The main difference between the corpus probabilities and subjects' ratings is that subjects rate the likelihood of an animate entity being a patient much higher than actually occurs in the corpus—the typicality ratings for animate patients (*How common is it for a speaker to be proposed by someone?*) average about 4.7, not much different than the ratings for animate agents and inanimate patients, but the probabilities of animate patients in the corpus average only about .30, considerably lower than the probabilities for animate agents and inanimate patients.

The main question for the typicality ratings is whether they can explain why the sets of sentences used by previous researchers (the sentences from Trueswell et al., 1994; MacDonald, 1994; and McRae et al., 1998) yield longer reading times for reduced relatives with typically agent heads than for reduced relatives with typically patient heads, and the sets of sentences from the corpus do not. Consider the typicality ratings for the Trueswell et al. noun–verb pairs and the ratings for the pairs from the first set of corpus sentences. The ratings do not differ significantly, $F_2(1, 126) = 2.1$. Even though reduced relatives with animate heads had significantly longer reading times than reduced relatives with inanimate heads for Trueswell et al.'s sentences but not for the corpus sentences, the typicality ratings for the two sets of sentences were not significantly different. This suggests that typicality ratings cannot explain the differences in reading time results.

Table 10 (under "Typicality ratings") shows the relevant correlations for all the sets of sentences. According to the constraint-based models under discussion here, reading time for a reduced relative with an animate head should increase as the typicality of an animate agent for the verb increases, and it should increase as the typicality of an animate patient decreases. The difference between the two typicality values should correlate positively with the animate–inanimate reading time difference. As the results in Row A in the table show, the correlation was significantly positive only for the McRae et al. (1998) sentences, $t(38) = 1.8$. For reduced relatives with inanimate heads, reading time should decrease as the typicality of an inanimate agent decreases and reading time should decrease as the typicality of an inanimate patient increases. The difference between these two typicality values should correlate negatively with the animate–inanimate reading time difference. For two of the five sets of sentences (Row B in the table), the predicted correlation was significant, $t(14) = -2.5$ for the second set of corpus sentences, and $t(30) = -2.8$ for the MacDonald (1994) sentences. Finally, the animate–inanimate reading time difference should be positively correlated with the difference between the typicality values for animate heads and the typicality values for inanimate heads (the row labeled "A − B" in the table). A significant positive correlation was found only for the McRae et al. sentences, $t(38) = 1.7$. Figure 2 displays scatter plots of the correlations, showing that the ranges of the typicality values are wide enough that the predicted correlations could have been observed had they been present in the data.

In previous research, there have been few efforts to examine whether constraint-based factors correlate in the predicted ways with reading times. Trueswell et al. (1994) found the predicted negative correlation for typicality of inanimate patients but did not report correlations for any of the other typicality ratings (animate and inanimate agents and animate patients). For our data, the correlation of the reading time difference with the typicality of inanimate patients was also negative (−.19) but not significant, $t(14) < 1.0$. The most complete examination of whether typicality values predict reading times was done by McRae et al. (1998). They built typicality values for typically agent and typically patient heads of reduced relatives into a constraint-based model and

Table 11
*Typicality Ratings (With Interquartile Ranges) and Corpus Probabilities*

| Sentence items | Corpus probability | Typicality rating |
|---|---|---|
| Animate agent: How common is it for a speaker to propose someone? | | |
| Trueswell et al. (1994) | .77 | 4.5 (2.1) |
| MacDonald (1994) | .69 | 4.0 (2.4) |
| Corpus, Set 1 | .58 | 4.6 (1.6) |
| Corpus, Set 2 | .59 | 3.8 (1.7) |
| McRae et al. (1998) | .74 | 4.2 (0.6) |
| Inanimate agent: How common is it for a solution to propose someone? | | |
| Trueswell et al. (1994) | .23 | 1.8 (0.4) |
| MacDonald (1994) | .31 | 1.8 (0.6) |
| Corpus, Set 1 | .41 | 2.1 (0.5) |
| Corpus, Set 2 | .41 | 2.3 (0.8) |
| Animate patient: How common is it for a speaker to be proposed by someone? | | |
| Trueswell et al. (1994) | .17 | 5.3 (1.4) |
| MacDonald (1994) | .27 | 4.7 (2.0) |
| Corpus, Set 1 | .35 | 4.6 (1.4) |
| Corpus, Set 2 | .42 | 4.1 (2.4) |
| McRae et al. (1998) | .66 | 4.2 (1.0) |
| Inanimate patient: How common is it for a solution to be proposed by someone? | | |
| Trueswell et al. (1994) | .83 | 5.9 (0.6) |
| MacDonald (1994) | .73 | 5.4 (1.6) |
| Corpus, Set 1 | .65 | 4.2 (2.1) |
| Corpus, Set 2 | .58 | 4.2 (1.9) |

*Note.* Typicality was rated on a scale of 1 to 7, with 7 being most common.

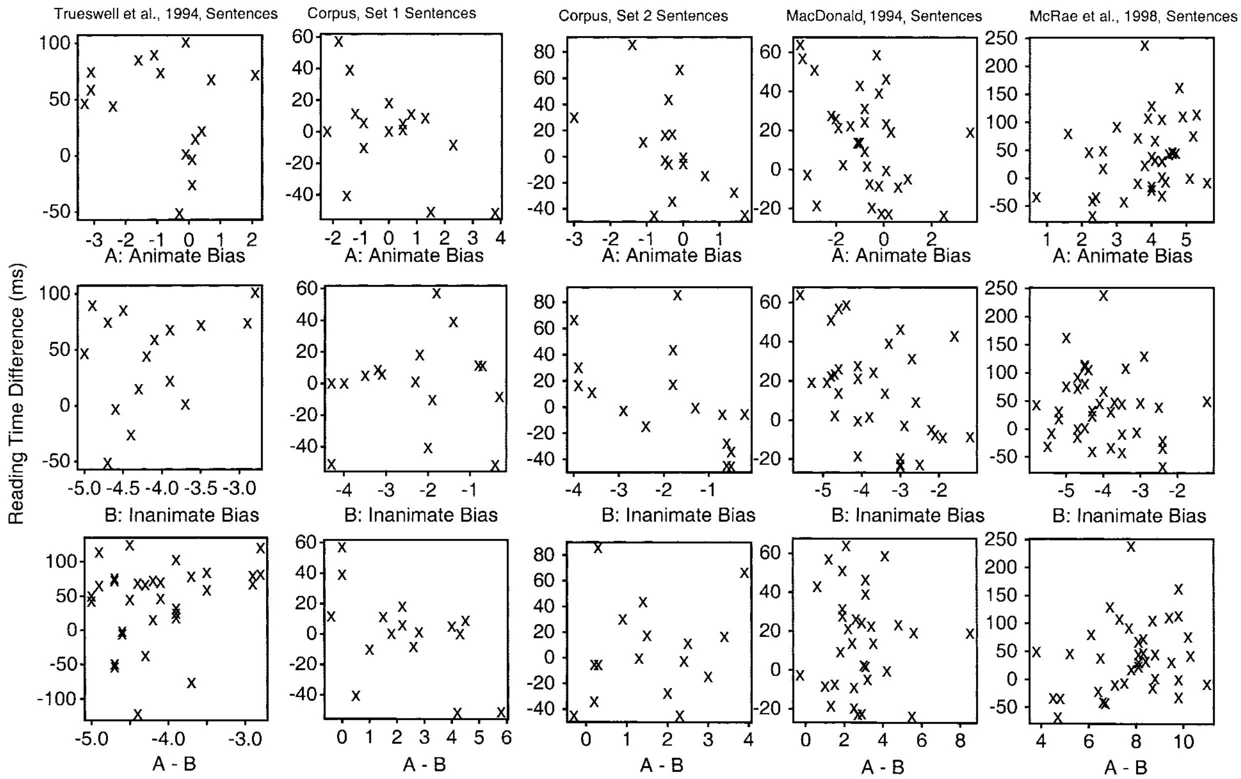


*Figure 2.* Scatter plots for correlations of reading time per word differences with biases determined from typicality ratings.

found that the model did a good job of predicting reading times. The significant correlation we obtained between typicality ratings and reading times for their sentences replicates their finding. On the whole, however, our studies provide a larger and more systematic examination than has been available previously for the empirical effects of constraints defined by constraint-based models and the data fail to show the correlations predicted by the models. In other words, the data are not well accounted for by the models.

## Summary

Constraint-based models are directly in conflict with MTS. As part of the lexical entry for a verb, constraint-based models encode syntactic information separately from semantic information. MTS encodes only semantic information, and syntactic information emerges from the semantic information (through linking rules) in the sense that the entities in an event template fill positions in sentences that can be syntactically labeled (e.g., as the subject of a sentence). For constraint-based models, the main determiner of processing difficulty for an input string of words is the frequency with which various candidate analyses are used in the language. If the correct analysis of an input string is rarely used, processing will be slowed as more frequently occurring candidates win the competition for activation. According to MTS, frequency information should play little if any role in processing (except for the frequencies of words themselves).

Empirically, the critical question, in the context of this article, for evaluating the constraint-based versus the MTS views is which of the two generates the right predictor variables for the likelihood of occurrence of reduced relative clauses and for their processing difficulty. MTS predicts that reduced relatives with verbs that denote causality internal to the entity in head position should never occur in reduced relatives, barring error on the part of the speaker or writer. The finding of an occurrence rate of about 1 in 3,500 for internal control verbs, compared with an occurrence rate of about 1 in 16 for verbs that do provide the appropriate causality, is in accord with this prediction.

For all the constraint-based predictor variables that we examined, none consistently accounted for the data from the corpus studies or the reading time experiments. Constraint-based models should predict that reduced relatives occur less frequently in naturally produced language than nonreduced relatives; they should predict that reduced relatives are less likely to occur in ambiguous, subject positions in sentences; and they should predict that reduced relatives are less likely to occur with animate heads—all of these predictions are wrong according to the corpus data we examined. In Experiments 8–11, reading times were collected for several sets of sentences with reduced relative clauses. For some of the sets of sentences, reading times were longer when the head of the relative was more typically an agent of the verb, not a patient as required by the relative clause. For others of the sets of sentences, there was no significant reading time difference. Most importantly, for none of the sets of sentences was the reading time difference consistently correlated with the predictor variables identified by constraint-based theories. For the nouns that appeared as the heads of reduced relatives, their frequency of occurrence in the corpus as agents versus patients for the reduced relative verbs should correlate with the reading time difference: The more likely a noun is to occur as the patient of a verb, the easier the reduced relative should be; the more likely the noun is to occur as the agent of the verb, the more difficult the reduced relative should be. But the correlations were consistently not significant. Likewise, when the fit of a noun to the argument role positions of a verb was measured with typicality ratings, the correlations with reading time differences were consistently not significant.

Given that the variables that have been defined by constraint-based models are not good predictors of processing difficulty for sentences with reduced relative clauses, then what does account for their sometimes being difficult? Why, for example, is *The speaker proposed by the group would work perfectly for the program* more difficult than *The solution proposed by the group would work perfectly for the program*? The most obvious guess is that when differences in difficulty are observed, they come from factors outside of syntactic processing. We offer two suggestions about what such factors might be. Perhaps, relative to the real world, it is easier to understand how a group could propose a solution than how a group could propose a speaker. Consider these reduced relatives used by MacDonald (1994), McRae et al. (1998), or Trueswell et al. (1994), all with animate, typically agent, heads:

*A slave sold in the crowded bazaar . . .*

*The snake devoured by the tribesman . . .*

*The lion chased by the zookeeper . . .*

*The juror convicted by the judge . . .*

*The hangman executed by the government . . .*

*The babysitter punished by her mother . . .*

In the real world, selling slaves is politically incorrect, eating snakes is disgusting, and chasing lions is outright dangerous. In each case, any difficulty in processing likely does not come from syntactic analysis but rather from an attempt to accommodate the information in the reduced relative with knowledge of the world. Similarly, a judge convicting a juror, a government executing a hangman, and a mother punishing a babysitter all violate standard schemas. Whether this explanation can extend to all the sentences of the sets that show differential difficulty of reduced relatives is not clear, but we believe it is a strong possibility.

Real world factors and schemas are outside the domain of MTS (although we assume that the semantic structures postulated by MTS would be the propositions out of which larger sentence and text representations are constructed). For sentences to be used in experiments with relative clauses, we chose sentences from the corpus, that is, sentences that were naturally produced by speakers and writers. Thus, we have some degree of confidence that these sentences do not stray far from real world acceptability.

The second suggestion about factors that might make the reduced relatives with animate heads in the sentences from MacDonald (1994), McRae et al. (1998), and Trueswell et al. (1994) difficult comes from Fox and Thompson's (1990) analysis of the discourse functions of relative clauses. They suggest that the head of a reduced relative (like any other discourse entity) must be grounded in the discourse. If it is not already grounded, then it will often be the function of the relative clause to provide grounding. Grounding gives a link between a discourse entity and knowledge the reader already has. For the isolated sentences that have been used in many reduced relative experiments, no larger discourse is provided and so grounding must come from the reader's general knowledge of the world. Fox and Thompson distinguish between methods of grounding for people and methods of grounding for inanimate objects. Thus, it is possible that the requisite grounding was not as easily available for the reduced relatives with animate (almost always people) heads as for the reduced relatives with inanimate heads.

In conclusion, constraint-based models did not stand up well to the empirical tests described here. Besides the corpus statistics and the correlations between reading times and possible predictor variables, constraint-based models also have trouble with the finding that switching a relative clause from nonreduced to reduced can make it a better sentence (Experiment 4) and the finding that the information in a reduced relative is more tightly connected in memory to the head of the clause than the information in a nonreduced relative (Experiments 5 and 6). Constraint-based models treat reduced and nonreduced relatives as equal in all respects except their parsing difficulty, so they have no way to predict findings that show differences in the meanings they encode into memory.

## GENERAL DISCUSSION

Modern psycholinguistics began in the 1960s when cognitive psychologists borrowed some of the constructs of Chomsky's transformational grammar and attempted to test their psychological reality (see Fodor, Bever, & Garrett, 1974, for a review of this work). A large number of studies examined how various hypothesized syntactic structures might be revealed by the empirical methods of cognitive psychology. The consensus that emerged by the early 1970s was that the human language processing system does not engage in the transformations postulated by transformational grammar. The response to this conclusion was twofold: Some researchers with close ties to linguistics began to investigate alternative syntactic processing systems (Bever, 1970; Frazier, 1978). But a number of cognitive psychologists rejected syntax as an interesting domain of study and moved instead to the study of meaning (e.g., Bransford, Barclay, & Franks, 1972; Kintsch, 1974; Norman, Rumelhart, & the LNR Research Group, 1975). This split between psycholinguists interested in syntax and psycholinguists interested in meaning has continued until the present time. The two domains of interest rarely interact with each other, with the consequence that questions about how syntax and semantics interact are rarely addressed.

The research described in this article is an initial attempt to begin to fill this gap. Researchers interested in meaning have often supposed that the units of meaning are propositions—small bits of information made up of a relation and its arguments. The favored questions for investigation have been how propositions are combined into larger bodies of information and how inferences are generated from that information. Left aside has been the question of how propositions and relations among propositions are recovered from input strings of words. Our hope is that MTS may go some way toward providing an answer to this question. For verbs, the proposed event templates are representations of meaning and as such are much like propositions. The event template for the verb *break*, for example, is $\alpha$ CAUSE (*x*(BECOME BROKEN)). For the sentence *John broke the window*, the entity *John* would represent the cause of the breaking event and *window* would be the entity that became broken. Propositionally, this would be written (*break*, *John*, *window*).

For constructions like reduced and nonreduced relatives, the meanings of the constructions could guide the formation of propositions. We hypothesize that a reduced relative clause introduces into a discourse an entity placed into some state by an external force (f(*x*)), whereas a nonreduced relative introduces two pieces of information, an entity (*x*) and a fact about *x* (g(*x*)). The reduced relative sentence *The window broken by John was expensive* could be translated into two propositions. One proposition, from the reduced relative, would be (*broken by John*, *window*), which expresses a relation *broken by John* and its argument *window*. The second proposition (*expensive*, (*broken by John*, *window*)) would express a modifier, *expensive*, on the entity introduced by the first proposition. The same sentence but with a nonreduced relative clause (*The window that was broken by John was expensive*) could be translated into two different propositions. One would express the modifier *expensive* of the window (*expensive*, *window*), and the other would express the window being broken by John (*break*, *John*, *window*). The differences between the propositions for the reduced relative sentence and the propositions for the nonreduced relative sentence capture the differences we hypothesize between the meanings of the two constructions.

At the moment, MTS is not a model of how words are processed into propositions. It is still at the stage of generating predictor variables and testing their effects. This is as it should be because the predictor variables are new to psycholinguistics. Investigations of event templates as possible representations of the syntactically relevant parts of verb meanings are only in their earliest stages, yet the corpus studies and experiments presented here and by McKoon and Macfarland (2000a, 2000b, 2000c) show they have large effects on language production and comprehension. Reduced relative clauses have been studied in dozens of previous experiments, but it has never before been shown that they differ from nonreduced relative clauses in the meanings that are encoded as a product of their comprehension and it has never been explicitly demonstrated that reduced relative clauses can sometimes be better than nonreduced relatives.

As well as demonstrating MTS's usefulness in generating important new variables for study, we also showed that MTS has merit over major competitors, some classes of constraint-based models. The predictor variables defined by the constraint-based models did a poor job of correlating with sentence processing times. Our demonstrations do not, of course, rule out connectionist models altogether—no one can ever do that—but they do show that connectionist schemes coupled with the constraints defined by some current models are inadequate.

In testing MTS itself, we showed that reduced relatives of the kind *The horse raced past the barn* occur with a frequency of only about 1 in 3,500 in naturally produced language, compared with about 1 in 16 for other kinds of verbs. The prediction that reduced relatives of the kind *The horse raced* should never occur (except through error) was based on a combination of the meanings hypothesized for verbs (event templates) and the meaning hypothesized for the reduced relative construction. We provided considerable independent empirical support for these hypothesized meanings. The convergence of that support with the observed rarity of reduced relatives of the type *The horse raced past the barn fell* makes MTS a compelling alternative to other current views of sentence processing, especially in its ability to generate new ideas about what kinds of information govern language processing.

## References


Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30,* 191–238.

Anderson, S. R. (1971). On the role of the deep structure in semantic interpretation. *Foundations of Language, 6,* 197–219.

Bates, E., & Goodman, J. (1997). On the inseparability of grammar and the lexicon: Evidence from acquisition, aphasia, and real-time processing. *Language and Cognitive Processes, 12,* 507–584.

Bever, T. G. (1970). The cognitive basis for linguistic structure. In J. R. Hayes (Ed.), *Cognitive development of language* (pp. 279–352). New York: Wiley.

Bever, T. G., Sanz, M., & Townsend, D. J. (1998). The emperor's psycholinguistics. *Journal of Psycholinguistic Research, 27,* 261–284.

Bolinger, D. L. (1968). Entailment and the meaning of structures. *Glossa, 2,* 119–127.

Borkin, A. (1985). *Problems in form and function.* Norwood, NJ: Ablex.

Bransford, J. D., Barclay, J. R., & Franks, J. J. (1972). Sentence memory: A constructive versus interpretive approach. *Cognitive Psychology, 3,* 193–209.

Chomsky, N. (1986). *Knowledge of language: Its nature, origin and use.* New York: Praeger.

Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntax processor. In D. R. Dowty, L. Karttunen, & A. M. Zwicky (Eds.), *Natural language parsing* (pp. 320–358). Cambridge, England: Cambridge University Press.

Croft, W. (1998). Event structure in argument linking. In M. Butt & W. Geuder (Eds.), *The projection of arguments: Lexical and compositional factors* (pp. 21–64). Stanford, CA: CSLI Publications.

Cruse, D. A. (1972). A note on English causatives. *Linguistic Inquiry, 3,* 520–528.

Fellbaum, C. (Ed.). (1998). *WordNet.* Cambridge, MA: MIT Press.

Fillmore, C. J. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory* (pp. 1–88). New York: Holt, Rinehart & Winston.

Fodor, J. A., Bever, T. G., & Garrett, M. F. (1974). *The psychology of language: An introduction to psycholinguistics and generative grammar.* New York: McGraw-Hill.

Forster, K. I. (2000). The potential for experimenter bias effects in word recognition experiments. *Memory & Cognition, 28,* 1109–1115.

Fox, B. A., & Thompson, S. A. (1990). A discourse explanation of the grammar of relative clauses in English conversation. *Language, 66,* 297–316.

Francis, W. N., & Kučera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar.* Boston: Houghton-Mifflin.

Frazier, L. (1978). *On comprehending sentences: Syntactic parsing strategies.* Bloomington: Indiana University Linguistics Club.

Frazier, L. (1987). Theories of syntactic processing. In J. L. Garfield (Ed.), *Modularity in knowledge representation and natural language processing* (pp. 291–307). Cambridge, MA: MIT Press.

Frazier, L. (1995). Constraint satisfaction as a theory of sentence processing. *Journal of Psycholinguistic Research, 24,* 437–468.

Frazier, L., & Clifton, C., Jr. (1996). *Construal.* Cambridge, MA: MIT Press.

Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in analysis of structurally ambiguous sentences. *Cognitive Psychology, 14,* 178–210.

Givon, T. (1985). Function, structure, and language acquisition. In D. I. Slobin (Ed.), *The crosslinguistic study of language acquisition* (pp. 1005–1028). Hillsdale, NJ: Erlbaum.

Goldberg, A. (1995). *A construction grammar approach to argument structure.* Chicago: University of Chicago Press.

Goldberg, A. (1997). The relationships between verbs and constructions. In M. Verspoor, K. D. Lee, & E. Sweetner (Eds.), *Lexical and syntactical constructions and the construction of meaning* (pp. 383–398). Amsterdam: John Benjamins.

Grimshaw, J. (1994). Lexical reconciliation. *Lingua, 92,* 411–430.

Hale, K., & Keyser, S. J. (1993). On argument structure and the lexical expression of syntactic relations. In K. Hale & S. J. Keyser (Eds.), *The view from Building 20: Essays in linguistics in honor of Sylvain Bromberger* (pp. 53–109). Cambridge, MA: MIT Press.

Jackendoff, R. S. (1972). *Semantic interpretation in generative grammar.* Cambridge, MA: MIT Press.

Kintsch, W. (1974). *The representation of meaning in memory.* Hillsdale, NJ: Erlbaum.

Levin, B. (1993). *English verb classes and alternations: A preliminary investigation.* Chicago: University of Chicago Press.

Levin, B. (1999). Objecthood: An event structure perspective. In *Proceedings of CLS 35: Vol. 1: The main session, Chicago Linguistic Society* (pp. 223–247). Chicago: University of Chicago.

Levin, B., & Rappaport Hovav, M. (1995). *Unaccusativity: At the syntax–lexical semantics interface.* Cambridge, MA: MIT Press.

Levin, B., & Rappaport Hovav, M. (1996). Lexical semantics and syntactic structure. In S. Lappin (Ed.), *The handbook of contemporary semantic theory* (pp. 487–508). Oxford, England: Blackwell.

MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes, 9,* 157–201.

MacDonald, M. C. (1997). Lexical representations and sentence processing: An introduction. *Language and Cognitive Processes, 12,* 121–136.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review, 101,* 676–703.

McKoon, G., & Macfarland, T. (2000a). *Complexity in the meaning and structure of verbs.* Manuscript submitted for publication.

McKoon, G., & Macfarland, T. (2000b). Externally and internally caused change of state verbs. *Language, 76,* 833–858.

McKoon, G., & Macfarland, T. (2000c). The interaction of syntax and semantics in the lexical representations of verbs. *Cognitive Psychology, 45,* 1–44.

McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review, 99,* 440–466.

McKoon, G., Ward, G., Ratcliff, R., & Sproat, R. (1993). Morphosyntactic and pragmatic factors affecting the accessibility of discourse entities. *Journal of Memory and Language, 32,* 56–75.

McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language, 38,* 283–312.

Merlo, P. (1994). A corpus-based analysis of verb continuation frequencies for syntactic processing. *Journal of Psycholinguistic Research, 23,* 435–457.

Moens, M., & Steedman, M. (1988). Temporal ontology and temporal reference. *Computational Linguistics, 14,* 15–28.

Ni, W., Crain, S., & Shankweiler, D. (1996). Sidestepping garden paths: Assessing the contributions of syntax, semantics and plausibility in resolving ambiguities. *Language and Cognitive Processes, 11,* 283–334.

Norman, D. A., Rumelhart, D. E., & the LNR Research Group. (1975). *Explorations in cognition.* San Francisco: Freeman.

Pustejovsky, J. (1991). The syntax of event structure. *Cognition, 41,* 47–81.

Pustejovsky, J. (1995). *The generative lexicon: A computational theory of lexical semantics.* Cambridge, MA: Bradford Books/MIT Press.

Rappaport Hovav, M., & Levin, B. (1998). Building verb meanings. In M. Butt & W. Geuder (Eds.), *The projection of arguments: Lexical and compositional factors* (pp. 97–134). Stanford, CA: CSLI Publications.

Rappaport Hovav, M., & Levin, B. (2000). Classifying single argument verbs. In P. Coopmans, M. Everaert, & J. Grimshaw (Eds.), *Lexical specification and insertion* (pp. 269–304). Hillsdale, NJ: Erlbaum.

Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. *Journal of Verbal Learning and Verbal Behavior, 22,* 358–374.

Smith, C. S. (1970). Jespersen's "move and change" class and causative verbs in English. In M. A. Jazayery, E. C. Polome, & W. Winter (Eds.), *Linguistic and literary studies in honor of Archibald A. Hill: Vol. 2. Descriptive linguistics* (pp. 101–109). The Hague, the Netherlands: Mouton.

Stevenson, S., & Merlo, P. (1997). Lexical structure and parsing complexity. *Language and Cognitive Processes, 12,* 349–399.

Strunk, W., & White, E. B. (1979). *The elements of style.* Boston: Allyn and Bacon.

Tabor, W., Juliano, C., & Tanenhaus, M. K. (1997). Parsing in a dynamical system: An attractor-based account of the interaction of lexical and structural constraints in sentence processing. *Language and Cognitive Processes, 12,* 211–271.

Tabossi, P., Spivey-Knowlton, M. J., McRae, K., & Tanenhaus, M. K. (1994). Semantic effects on syntactic ambiguity resolution: Evidence for a constraint-based resolution process. In C. Umilta & M. Moscovitch (Eds.), *Attention & performance XV: Conscious and nonconscious information* (pp. 589–616). Cambridge, MA: MIT Press.

Tenny, C. L. (1994). *Aspectual roles and the syntax–semantics interface.* Dordrecht, the Netherlands: Kluwer Academic.

Trueswell, J. C. (1996). The role of lexical frequency in syntactic ambiguity resolution. *Journal of Memory and Language, 35,* 566–585.

Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language, 33,* 285–318.

van Valin, R. D., & LaPolla, R. J. (1997). *Syntax: Structure, meaning and function.* Cambridge, MA: Cambridge University Press.

van Valin, R. D., Jr., & Wilkins, D. P. (1996). The case for "effector": Case roles, agents, and agency revisited. In M. Shibatani & S. A. Thompson (Eds.), *Grammatical constructions: Their form and meaning* (pp. 289–322). Oxford, England: Oxford University Press.

Wasow, T. (1985). Postscript. In P. Sells (Ed.), *Lectures on contemporary syntactic theories* (pp. 193–205). Stanford, CA: CSLI Publications.

Wierzbicka, A. (1988). *The semantics of grammar.* Amsterdam: John Benjamins.

Wilson, K. (1993). *The Columbia guide to standard American English.* New York: Columbia University Press.

Received July 21, 2000
Revision received November 27, 2001
Accepted July 24, 2002 ■