

CHAPTER 8. DISCUSSION AND CONCLUSIONS

8.1 Introduction

The purpose of this dissertation has been to propose and evaluate the Event Structure Processing hypothesis, defined in (149).

(149) *Event Structure Processing (ESP) hypothesis: During comprehension, event structure information, accessed in a verb's lexical entry, affects parsing decisions. If the verb is inherently telic, the verb will be parsed as having an underlying direct object.*

As explained in chapters 1 and 2, event structure refers to the type of event denoted by a predicate and the relationship between the subparts of the event and the syntactic components of the predicate, such as the verb and its arguments. For example, if Morpheus is building a ship and then he quits in the middle of doing it, it is not true that he has built a ship. In contrast, if Trinity is running and then she quits in the middle of running, it is still true that she has run. These two events, *building a ship* and *running*, differ in their event structure. A crucial difference in their event structure is the existence or non-existence of an endpoint. The reason that Morpheus has not built a ship if he quits in the middle is that the meaning of *built a ship* includes an endpoint. The reason Trinity has run even if she quits early is because the meaning of *run* does not include the notion of an endpoint. There are other aspects of a verb's event structure such as the existence of a causer and the existence of a process portion, though there is much controversy over the existence of both of these possible event subparts. This dissertation focuses on one

particular component of verb event structure, the endpoint or *terminus*. An event that includes an endpoint in its meaning is known in the event structure literature is known as a *telic* event. An event that does not include an endpoint in its meaning is considered *atelic*.

The primary goal of this dissertation has been to investigate the question of whether verb telicity information is used in language comprehension. The reduced relative clause versus main clause ambiguity was used to investigate this question, because it shows which factors affect the comprehender's degree of commitment to the main clause analysis of a sentence. For example, a number of factors (e.g., animacy, telicity, plausibility) may influence the comprehender to be more committed to the main clause analysis during the ambiguous region in (150)a than in (151)a.

(150)

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

- b. *The horse that was raced past the barn fell.*

(151)

- a. *The chocolate melted by the chef smelled delicious.*

- b. *The chocolate that was melted by the chef smelled delicious.*

In studies on the reduced relative versus main clause ambiguity, the unreduced relative clause version (such as the sentences in (150)b and (151)b) is considered to provide a baseline measure of processing difficulty (but see McKoon and Ratcliff, 2003, for an alternative view). The dependent variable in most experiments on reduced relatives is known as the reduced relative effect (RRE). The RRE is calculated as the reaction times or error rates for the reduced condition minus the reaction times or error rates for the unreduced condition.

One influence that led to the development of the experiments in this dissertation was Sanz's (2000) prediction that verb event structure would affect the RRE (though not strictly in terms of telic versus atelic events). A previous result that suggested that telicity might affect the RRE is MacDonald's (1994) finding that obligatory transitivity affects the RRE. The rationale behind the prediction that telicity would reduce the RRE was that telic verbs, like obligatorily transitive verbs, require an internal argument. The difference between these two types of verb information is that obligatorily transitive verbs require a surface direct object while telic verbs require an underlying direct object. This difference turns out to have a big effect on how these two types of verb information affect the RRE, as shown in the results of the word maze experiment, which will be discussed further in this chapter. Because the difference between telic and atelic verbs is a difference in event structure, the finding that telicity reduces the RRE supports the ESP hypothesis.

8.2 Findings Supporting the ESP Hypothesis

The major empirical question that has been investigated in this dissertation is whether inherent verb telicity information reduces the RRE. This question was investigated in three experimental paradigms: 1) self-paced reading, 2) speaker change monitoring (an auditory task), and 3) the word maze. Results from all three paradigms support the conclusion that the RRE is smaller when the embedded verb is telic than when it is atelic.

The most compelling results are from the word maze experiment. In the word maze paradigm, participants are presented with a series of word pairs, and their task is to choose which of the two words is a grammatical continuation at each step. In order to complete a trial successfully, the participant must process the sentence incrementally, parsing the previously presented words in order to make each choice. The participant's reaction time for integrating each word into the sentence is measured. The word maze experiment showed that reduced relative effects (RREs) measured as reaction times for choosing the word *by* were significantly smaller when the embedded verb was telic as in (152)a than when it was atelic as in (152)b.

(152)

- a. *The actress awakened by the writer left in a hurry.*

- b. *The actress sketched by the writer left in a hurry.*

A significant effect of telicity on the RRE was also found in the speaker change monitoring paradigm. In this paradigm, participants listen to spoken sentences and indicate whether they hear a word or part of a word spoken by a different speaker than the rest of the sentence. Speaker change monitoring experiment 4 in chapter 5 showed that RREs measured as speaker change detection errors were significantly smaller when the embedded verb was telic than when it was atelic. The sentences were almost identical to the ones used in the word maze paradigm (such as the sentences in (152)a and (152)b). The speaker change always occurred on the first syllable of the noun in the *by*-phrase, e.g. *writer* in (152)a and (152)b.

The other paradigm which showed an effect of telicity on the RRE was self-paced reading. On each trial in the self-paced reading paradigm, participants are presented with a sentence, one region at a time, and they successively press a key to make each region of the sentence appear on the screen. After many of the trials, the participant is asked a question about the meaning of the sentence in order to encourage comprehension. A post-hoc analysis of two self-paced reading experiments showed that RREs (measured as self-paced reading times for the “embedded verb plus *by*” region) were significantly smaller when the embedded verb was telic than when it was atelic. The sentences in the two self-paced reading experiments that showed the significant telicity effect were similar to the sentences in (152)a and (152)b.

8.3 Less Clear Results

There was one other self-paced reading experiment for which there was not a significant effect of telicity on the RRE alone or in combination with the other two studies. This study, first reported in Clarke, Townsend, and Bever (2000), differed from the other two studies (McRae, Spivey-Knowlton, & Tanenhaus, 1998; Tabossi, Spivey-Knowlton, McRae, & Tanenhaus, 1994) in that half of the experimental trials used non-*by* prepositional phrases rather than *by*-phrases. The non-*by* prepositional phrase results showed a larger telicity effect in the predicted direction but only on the main verb as opposed to on the earlier region including the embedded verb plus preposition. The *by*-phrase results from Clarke et al.'s study were in the same direction as the results from McRae et al.'s study and Tabossi et al.'s study. RREs for sentences with telic embedded verbs were approximately 100 milliseconds smaller than RREs for sentences with atelic embedded verbs. However, there was a great deal of variance, because there was only one item in one of the four verb types and only two items in another of the four verb types. This imbalance in the number of items is most likely the reason that the results are not statistically significant.


In addition to the speaker change monitoring experiment that showed the significant effect of telicity on the RRE, post-hoc analyses of speaker change detection errors from two other speaker change monitoring experiments showed that the RREs were smaller for sentences with telic verbs than for those with atelic verbs, as predicted. Of the four total speaker change monitoring experiments, only one of them did not show smaller detection error RREs for sentences with telic verbs. In this particular experiment,

the RREs were extremely small or negative for all four verb types, and the difference between the RREs in the telic and atelic conditions was extremely small (0.7%). The reaction times and identification errors in the speaker change monitoring paradigm generally do not show the predicted effect of smaller RREs in sentences with embedded telic verbs. The reason for this is unknown, since these experiments are the first time that reaction times and identification errors have been collected in this paradigm. The identification errors did not show a significant RRE in any of the experiments, which suggests that they are not a good measure of on-line processing difficulty. A look at the word maze results on the position where the speaker change occurred is relevant to the interpretation of the speaker change monitoring results.

The word maze results in Figure 18 show that on the noun in the *by*-phrase, the telicity effect is opposite of the predicted direction and opposite of the direction of the effect that occurs earlier in the sentence on the preposition *by*.³⁶

³⁶ In Figure 18, one star indicates that the factor (telicity or transitivity) had a significant effect ($p < .05$) on the RRE in that sentence region. Two stars indicates that the factor had a highly significant effect ($p < .01$) on the RRE.

Figure 18. Telicity and transitivity effects on the RRE in the word maze experiment



The noun in the *by*-phrase is the where the speaker change occurred in the speaker change monitoring experiment. The word maze results suggest that the location of the speaker change is too late in the sentence for the predicted telicity effect to show up. Given these results, it is surprising that the predicted telicity effect is found in the speaker change detection errors.

8.4 Towards a Unified Account of the Telicity and Transitivity Effects

The comparison of the telicity effect and the transitivity effect in the word maze results shows that these two types of verb information affect the processing of reduced relatives

very differently. In the graph in Figure 18, if the difference scores were computed differently, the telicity and transitivity effects would look very similar. Why is the obligatorily transitive condition subtracted from the optionally transitive condition, instead of the reverse? The reason that the difference scores are computed as they are in Figure 18 is because the telic verb and the obligatorily transitive verb have something important in common: both of them require a direct object. Thus, the transitivity effect as shown in Figure 18 reflects the effects of requiring a surface direct object. The telicity effect as shown in Figure 18 reflects the effects of requiring an underlying direct object. The transitivity effect on *by* seems intuitive enough; in the obligatorily transitive condition, a surface direct object was expected following the verb. The occurrence of *by* shows that the surface direct object is absent; this fact appears to cause processing difficulty. The telicity effect shows that requiring an underlying direct object has the opposite effect of requiring a surface direct object; it seems to facilitate the integration of *by*. These effects are large and significant, and they present a very interesting puzzle. Why does the comprehension system respond differently to the requirement of an object depending on whether the object is required on the surface or underlyingly?

In the following subsections, I will discuss possible accounts of these results in the frameworks of Townsend and Bever's (2001) Late Assignment of Syntax Theory (LAST) and constraint-based models such as McRae, Spivey-Knowlton, and Tanenhaus 1998. Following those accounts, I will turn to the issue of teasing apart the effects of unaccusativity and telicity. The experimental data in this dissertation shed some light on

this issue, but the data also raise further questions that can only be answered through future research.

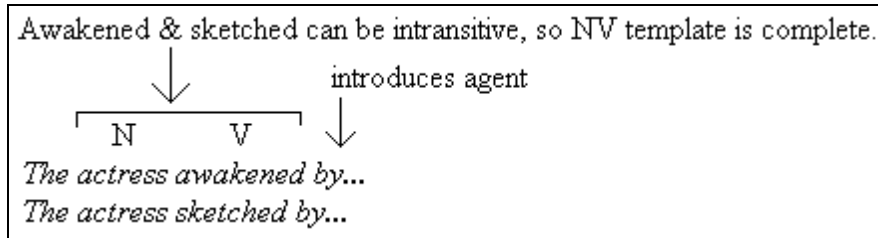
8.4.1 An Account of the Telicity and Transitivity Effects Within LAST

In chapter 7, section 7.2.3, I reviewed the basics of Townsend and Bever's (2001) LAST model, discussed their account of the reduced relative effect (RRE), and discussed the ways that telicity could be represented in the model. Here I will present an account of both the telicity and transitivity effects on the RRE within the LAST framework.

In the early stages of comprehension in the LAST model, pseudo-syntactic strategies are applied to the utterance. An important pseudo-syntactic strategy in the model is the identification of canonical sentence templates such as NV(N). The application of such templates will be critical in accounting for telicity and transitivity effects on the RRE.

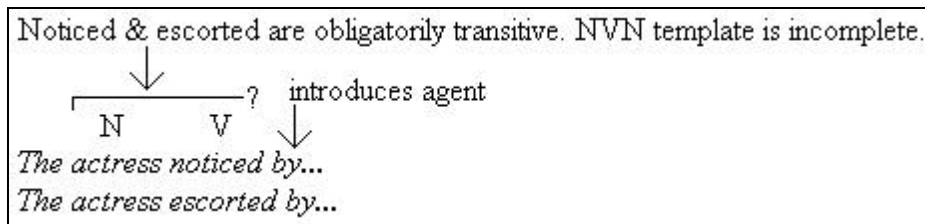
First, I will explain how the application of the NVN and NV(N) templates accounts for the effects of obligatory transitivity throughout the sentence. Since an obligatorily transitive verb requires a direct object, identification of a full predicate will require matching the input with the NVN template (rather than the template involving NV plus an optional N, which can be used with intransitive or optionally transitive verbs) (Townsend & Bever 2001: 253). Figure 19 shows how the NV template is applied in the optionally transitive condition.

Figure 19. Optionally transitive embedded verbs in canonical sentence templates



Once the initial noun and verb is matched up with the NV template, this material is assigned a preliminary meaning and passed off to the second stage of the model, which involves a full syntactic analysis of the material. In the obligatorily transitive condition, the NVN strategy fails, as shown in Figure 20.

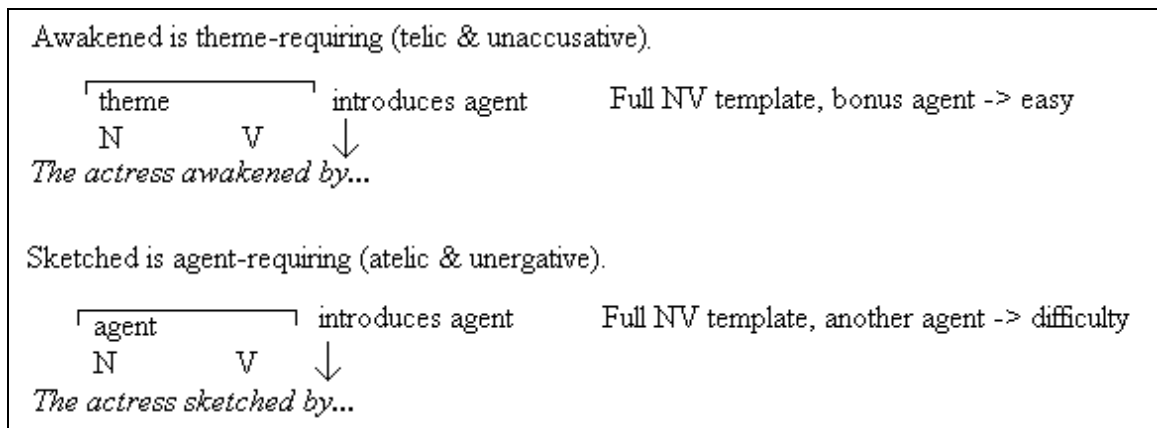
Figure 20. Obligatorily transitive embedded verbs in canonical sentence templates



Following the verb, the NVN template requires a noun, but *by* occurs instead. As a result, the NVN template must be abandoned, resulting in comprehension difficulty on the word *by*. Later on in the sentence, the information that the verb is obligatorily transitive will facilitate the reduced relative analysis, but at this early point in the sentence, it only causes difficulty.

Accounting for the telicity effects in the LAST model requires associating the NV(N) templates with thematic roles, which is assumed by Townsend and Bever (2001: 267). A telic verb such as *awakened* requires an argument that is assigned the role of theme; Townsend and Bever refer to such verbs as “theme-requiring”. Thus, the NV template will associate the initial noun with the role of theme as shown in the upper half of Figure 21.

Figure 21. Telic and atelic verbs in canonical sentence templates



An atelic verb such as *sketched* does not require a theme argument, so it is considered “agent-requiring”. Thus, the NV template associates the initial noun with the role of agent, as shown in the lower half of Figure 21. Encountering *by* activates the strategy that marks the following noun as an agent. Since there is already an agent, this causes immediate processing difficulty on *by*. An additional reason that telicity should affect comprehension on the word *by* is because both telicity information and the word *by* activate strategies that make reference to the D-structure underlying level, the level where

thematic roles are assigned. In contrast, obligatory transitivity makes reference to the S-structure level, the level in which words are in their pronounced order.

I have just explained why the LAST model using canonical sentence templates predicts that in the experimental material sentences, telicity should facilitate the integration of *by* in the experimental material sentences while obligatory transitivity makes the integration of *by* more difficult. In the next subsection, I will explain how constraint based models can make the same predictions. Then, I will discuss some experimental data that supports one of the theoretical frameworks but not the other.

8.4.2 A Constraint-based Account of the Telicity and Transitivity Effects

All of the results in this dissertation can also be accounted for in a constraint-based model, such as the ones proposed by MacDonald, Pearlmutter, and Seidenberg (1994), Trueswell and Tanenhaus (1994), and McRae, Spivey-Knowlton, and Tanenhaus (1998). In such models, telicity information in the lexical entry of a verb can be part of a constraint that increases the activation of the reduced relative structure and decreases the activation of the main clause structure. The transitivity effects can be explained by a combination of constraints. Obligatory transitive verbs increase the expectation of an immediately following noun phrase, and this may explain the large RRE for obligatorily transitive verbs on the early *by* region. An obligatorily transitive verb plus an entire *by* phrase would seem to increase the activation of the reduced relative structure much more than an optionally transitive verb plus a *by*-phrase. This can explain the small RRE for obligatorily transitive verbs on the main verb region.

8.4.3 Accounting for the Reactivation of Noun Phrases Following Trace Positions

The original experimental results presented in this dissertation cannot distinguish between the two accounts discussed above: the LAST account and the constraint-based account. However, the two accounts differ in their ability to explain the results of two recent studies on the reactivation of antecedents following unaccusative verbs and their NP trace positions. Friedmann, Taranto, Shapiro, and Swinney (2003) reported cross-modal priming evidence of an NP trace in English following all non-alternating unaccusative verbs (pure intransitives) and some alternating unaccusative verbs (optionally transitive verbs). Sanz (2000) provides evidence for an NP trace that occurs following unaccusative verbs in Spanish but not in English. My own knowledge of Spanish is not good enough to determine whether her unaccusative verbs are alternating (optionally transitive) or non-alternating (purely intransitive). Both Sanz's and Friedmann et al.'s results suggest that unaccusativity has implications for the processing of syntactic structure in comprehension. A more detailed description of these two experiments is given in chapter 7.

Townsend and Bever's (2001) LAST model predicts the existence of NP traces in passive constructions and following at least some unaccusative verbs. They also predict that the evidence of the NP trace should show up a few words downstream from the verb, which is exactly what Friedmann et al. and Sanz found. In the LAST model, the NP trace arises from the synthesis of a full syntactic structure in the second stage of the model and not from the use of heuristic strategies in the first stage of the model. Thus, evidence of

the NP trace should show up a few words downstream as a result of the basic architecture of the model.

Constraint-based models really have no basis for predicting the reactivation of antecedents at trace positions. Constraint-based models assume that syntactic structures are chosen through competition driven by constraints that differentially activate syntactic alternatives. Thus, nothing like a syntactic derivation occurs during a constraint-based model, and there is no reason for such models to predict reactivation of an antecedent at a trace position.

In conclusion, both the LAST model and constraint-based models account for the effects of telicity and transitivity on the RRE. However, only the LAST model can also account for the NP trace results reported by Sanz (2000) and Friedmann et al. (2003).

8.5 The Possibility that Unaccusativity is the Operative Factor

In this section, I discuss the possibility that unaccusativity rather than telicity could be the factor that affects the comprehension of reduced relative clause sentences. The results support the Event Structure Processing hypothesis regardless of whether the operative factor is unaccusativity or telicity. This is because unaccusative verbs and unergative verbs differ in their event structures. Strictly speaking, only intransitive verbs can be unaccusative or unergative. However, there are optionally transitive verbs that are unaccusative in their intransitive usage. This is the case with the verbs in the optionally transitive telic condition in the speaker change monitoring experiments in chapter 5 and those in the word maze in chapter 6. The verbs in the optionally transitive atelic condition

are unergative in their intransitive usage. Thus, if unaccusativity is the operative factor, telicity will appear to have a significant effect on comprehension for optionally transitive verbs but not for obligatorily transitive verbs. If this turns out to be the case, a purely structural account of the data is possible. Such an account would explain why it is very easy to reanalyze a simple main clause structure with an unaccusative verb as having a reduced relative structure. Stevenson and Merlo (1997) have proposed one such account.

I will begin by discussing differences in the thematic roles that the four verb types assign to their arguments, since this highlights a difference between the optionally transitive telic unaccusative verbs and the other three verb types. Then I will discuss the experimental data with a focus on whether telicity has a significant effect in both the obligatory transitive and optionally transitive conditions.

8.5.1 Thematic Role Assignment in the Four Verb Types

In a reduced relative clause sentence, the initial noun phrase always ends up receiving the theme role. Comprehension difficulty in the disambiguating region of the reduced relative may be a result of whether or not the initial noun phrase was originally assigned a theme role or an agent/causer role. For this reason, I discuss differences in the thematic roles assigned by the four verb types. First, I consider the case of an unambiguous unaccusative verb, which is always intransitive and thus cannot occur in a reduced relative structure. The initial noun phrase preceding a purely unaccusative verb such as *arrived* in (153) must always receive a theme role.

(153) *The woman arrived.*

In order for this fact to be used in comprehension, the fact that the initial noun phrase must be a theme must be encoded in the lexical entry of the verb *arrived*. The verb lexical entry for *arrived* could either encode the fact that its initial noun phrase is always a theme, or it could encode the fact that it is an intransitive, non-alternating unaccusative. The initial noun phrase immediately preceding the verb *awakened*, as in the sentence in (154), is often a theme. It could also be an agent/causer, as in *The woman awakened Neo*.

(154) *The woman awakened.*

The verb lexical entry for *awakened* could encode the fact that the initial noun phrase is either a theme or a causer, or it could encode the fact that *awakened* is an unaccusative verb that alternates with a transitive structure. This is the case with the verbs in the optionally transitive telic condition of the experiments in this dissertation. Importantly, the initial noun phrase of the optionally transitive atelic verb *sketched*, as in (155)a, can never be a theme unless passive morphology is present as in (155)b.

(155)

a. *The woman sketched.*

b. *The woman was sketched.*

Thus, the verb lexical entry for *sketched* may encode the fact that a noun phrase that immediately precedes it is an agent/causer. Of course, the reduced relative clause

structure is an exception to this generalization. The experimental results suggest that it is an exception that the comprehension system does not take into account. That would explain why a large garden path effect occurs on *by* following an optionally transitive atelic verb in a reduced relative. For both atelic obligatory transitive verbs as in (156)a, and telic obligatory transitive verbs as in (157)a, the initial noun phrase also cannot be a theme unless passive morphology is present, as shown in (156)b and (157)b.

(156)

- a. *The woman chased Neo.*
- b. *The woman was chased by Neo.*

(157)

- a. *The woman spotted Neo.*
- b. *The woman was spotted by Neo.*

Why should there be a difference between the lexical entries for the obligatorily transitive atelic verb *chase* and the obligatory transitive telic verb *spot*, with regard to their argument structure? An initial noun phrase immediately preceding either of them must be an agent/causer in any simple main clause sentence. In the next subsection, I consider

whether the processor is affected by the difference between the obligatory transitive telic and optionally transitive telic conditions.

8.5.2 A Detailed Review of the Findings for Obligatorily Transitive Verbs

A verb that immediately follows an initial noun phrase with a theme role is generally an unaccusative (either alternating or non-alternating) in a simple main clause sentence. Obligatory transitive verbs are never unaccusatives. This raises the question of whether the reduced relative effect (RRE) is really smaller for obligatorily transitive telic verbs as it is for optionally transitive telic verbs (the real unaccusatives). In the post-hoc analyses of self-paced reading experiments in chapter 4, the mean RRE in the obligatorily transitive telic condition was smaller than the mean RREs for both the optionally transitive atelic condition and the obligatorily transitive atelic condition. This result supports the idea that telicity reduces the RRE even in the obligatorily transitive condition. In the word maze results, the RRE on *by* was also smaller in the obligatorily transitive telic condition than in the optionally transitive atelic and obligatorily transitive atelic conditions. Thus, the maze results also support the idea that telicity reduces the RRE even for obligatorily transitive verbs. In the speaker change monitoring results, the detection error RREs in experiments 1 and 4 were smaller in the obligatorily transitive telic condition than in either the optionally transitive atelic or obligatorily transitive atelic conditions. In experiment 2, the detection error RREs were one percentage point smaller in the obligatorily transitive atelic condition than the obligatorily transitive telic condition. Both conditions had smaller RREs than the optionally transitive atelic

condition. The results of speaker change monitoring experiment 3 had extremely small or negative RREs in all four verb type conditions, and the results were not at all as predicted (see chapter 5, section 5.3 for more information). Thus, in two out of the three speaker change monitoring experiments that showed smaller RREs for telic verbs generally, the RRE was smaller with obligatorily transitive telic verbs than for either of the atelic conditions. These results support the notion that it is telicity rather than unaccusativity that reduces the RRE on the preposition *by*.

8.5.3 Evidence that the Optionally Transitive Telic Condition is Special

In the last subsection, I reviewed results supporting the claim that telicity rather than unaccusativity is the operative factor in predicting comprehension difficulty associated with the reduced relative effect. In this subsection, I show that a distinctive pattern of results was found in the unaccusative (optionally transitive telic) condition in the word maze experiment. Future research will be necessary to tease apart the effects of telicity and unaccusativity.

Figure 22 shows RREs for each of the four verb types in the word maze paradigm. The absolute reaction times for the reduced and unreduced conditions are given in Table 13 in chapter 6, section 6.3.

Figure 22. Reduced relative effects for the four verb types in the word maze paradigm

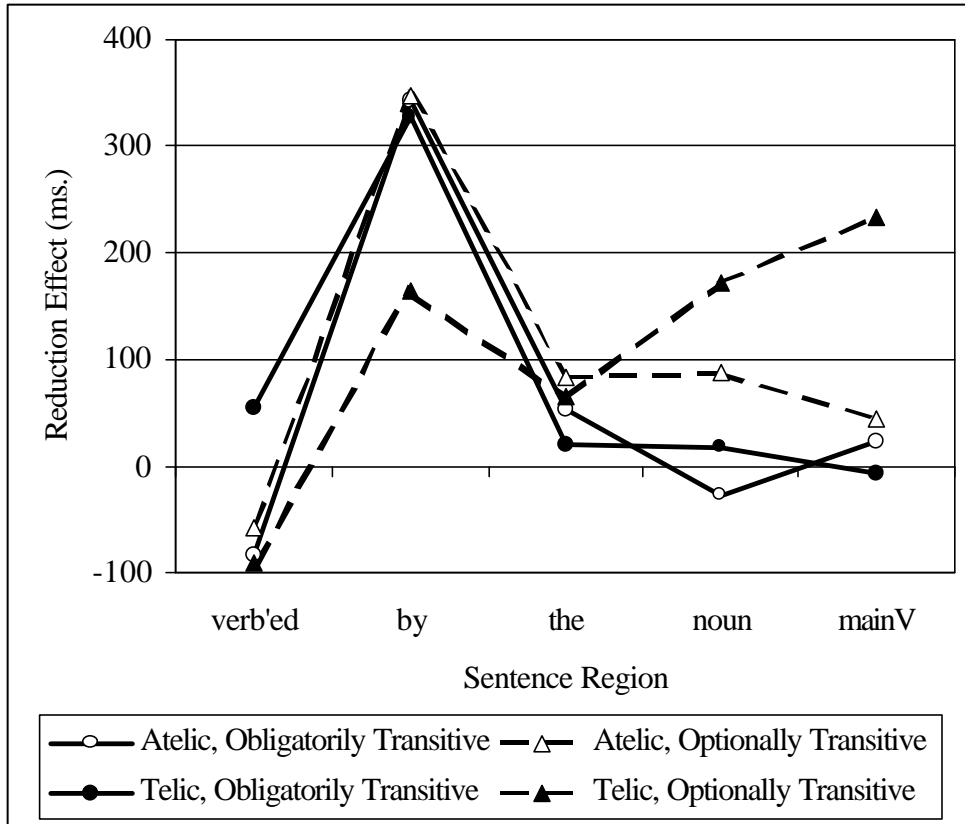


Figure 22 clearly shows that the pattern of results in the telic optionally transitive condition is quite different than the other three verb type conditions. The optionally transitive telic embedded verb appears to facilitate the integration of *by*. Later on the noun in the *by*-phrase and on the main verb, the telic optionally transitive condition shows the largest RREs.

The results suggest that in the optionally transitive telic condition, a simple main clause structure may be maintained until the noun in the *by*-phrase. In other words, the

comprehension system may be interpreting *by* as introducing a locative as in (158) or a time adverbial as in (159).³⁷

(158) *The actress awakened by the side of the swimming pool.*

(159) *The actress awakened by the time the guests arrived.*

A future experiment will be necessary in order to determine whether this is the case. The experiment would show differential effects depending on whether the noun in the *by*-phrase is an animate potential agent or a location.

8.6 Summary of Accounts of the Data

Here, I will briefly summarize the results and then review their implications for the accounts just discussed. The main result is that reduced relative effects are significantly smaller when the embedded verb is telic rather than when it is atelic. This effect is found at the earliest possible disambiguation point, on the word *by*. While the word maze results show that the telicity effect is much larger for optionally transitive verbs, the self-paced reading and speaker change monitoring results show that the telicity effect is also found for obligatorily transitive verbs. The word maze results suggest that unaccusativity may affect processing decisions; this issue requires future investigation. The word maze results strongly suggest that telicity and obligatory transitive have opposite effects on the comprehension of reduced relative clause sentences throughout the disambiguating

³⁷ The time adverbial interpretation of *by* was pointed out to me by John Lapeyre, personal communication.

region. These results are consistent with experiments from other laboratories on the effects of obligatory transitivity.

The LAST model and constraint-based models are both able to explain the opposite effects of telicity and obligatory transitivity. The LAST model is also able to explain the results of recent experiments that show reactivation of the antecedent at the locations where NP traces should occur. In the LAST model, telicity can be represented in a pseudo-syntactic strategy of applying the NV(N) canonical sentence template, where one of the Ns must receive the semantic role of theme. This mechanism predicts that telicity will have immediate processing effects.

8.7 How are the Experimental Results in this Dissertation Similar to and Different from Others' Results?

In this section, I discuss some relevant experimental results from other researchers and compare their methods and results to those in this dissertation. The research that is the most similar is Seegmiller, Ingrassia, and Townsend's (2003) self-paced reading study and their *for an hour* versus *in an hour* test. Seegmiller et al. investigated the effects of telicity and transitivity on the RRE in a word-by-word self-paced reading study that was somewhat similar to the experiments included in the post-hoc analyses in chapter 4. There were two major methodological differences between Seegmiller et al.'s experiment and the experiments in chapter 4. The biggest difference is that Seegmiller et al. used reduced relative clause sentences with non-*by* prepositional phrases as in (160) as opposed to those with *by*-phrases as in (161).

(160) *The woman healed at the clinic rejoined the demonstration.*

(161) *The defendant examined by the lawyer turned out to be unreliable.*

The results of Clarke, Townsend, and Bever (2000), one of the three self-paced reading experiments analyzed in chapter 4, suggest that the use of *by* versus another preposition is a factor that interacts with the effect of obligatory transitivity on the RRE. Specifically, Clarke et al. found a significant RRE only for optionally transitive verbs with non-*by* prepositional phrases. Also, Clarke et al.'s results show that when non-*by* prepositional phrases are used, the RRE does not show up until the main verb. The experiments analyzed in chapter 4 all show RREs earlier on the preposition *by*. The other difference between Seegmiller et al.'s experiment and the experiments analyzed in chapter 4 is the number of words that appeared on the screen together. Seegmiller et al. presented one word at a time, while in the experiments in chapter 4, the sentences were presented two words at a time. The number of words presented on the screen together in self-paced reading has been shown to affect the results. Burgess (1991) found that the animacy of the initial noun had no effect on the RRE using one word at a time presentation, but he found that animacy did have a significant effect when the words were presented two at a time. Seegmiller et al. found that obligatory transitivity had a significant effect on the RRE on the non-*by* preposition and on the noun two words downstream. They found that telicity had no significant effect on the RRE. The contrast between Seegmiller et al.'s results and the results in this dissertation suggest that the presence of the *by*-phrase may be necessary for seeing the telicity effect. One other difference is that Seegmiller et al.

used different verbs in the four verb types. There may be systematic differences in the types of verbs that they included. For example, they may have used manner of motion verbs such as *raced* in the optionally transitive atelic category. Manner of motion verbs were avoided in the experiments in chapters 5 and 6 that included telicity and transitivity as planned factors. The reason that they were avoided is because their structure is controversial in the event structure and argument structure literature.

Seegmiller et al. also report the results of their own *for an hour* versus *in an hour* “forced choice” task. Their experiment is very similar to the *for an hour* versus *in an hour* forced choice portion of the telicity questionnaire that I present in chapter 3. In their experiment, the participants were asked to choose between sentences such as (162)a and (162)b.

(162)

- a. *The mother burped the baby for three minutes.*

- b. *The mother burped the baby in three minutes.*

The rationale behind this task is that linguists claim that telic events should be used with the adverbial *in <some time interval>* and atelic events should be used with a durative adverbial such as *in <some time interval>*. In both Seegmiller et al.’s experiment and my own experiment in chapter 3, the events tested were based on the sentences used in the reduced relative clause experiment(s) and on classic examples of telic and atelic events.

In the telicity questionnaire in chapter 3, the task is actually forced choice; the participant had to choose one or the other. In Seegmiller et al.'s experiment, the participants were also allowed to choose "both" or "neither". In spite of this difference, Seegmiller et al.'s experiment yielded similar results. They found that the *for three minutes* variant was chosen significantly more often for atelic verbs than for telic verbs and that the *in three minutes* variant was chosen significantly more often for telic verbs than for atelic verbs. Seegmiller et al.'s conclusion from this result differs from my conclusion. They concluded that telicity is a gradient variable rather than a categorical variable, because *in three minutes* was not chosen 100 percent of the time with telic events. My conclusion is that this experiment cannot reveal whether telicity is gradient or categorical. Rather, I interpret the results to show that coerced interpretations (such as the repeated interpretation that results from a telic event used with *for three minutes* as in (162)a) are chosen significantly less often than interpretations that do not require aspectual coercion (interpretations in which a telic event is used with *in three minutes*, for example, as in (162)b).

Another result that is related to the experiments in this dissertation is reported by Stevenson and Merlo (1997). They conducted an informal survey in which they asked several colleagues to rate the acceptability of reduced relative clause sentences with alternating unaccusative verbs and with unergative manner of motion verbs. The reduced relative clause sentences with unergative manner of motion verbs were rated as much more difficult. These results are not truly comparable with the results of the experiments in this dissertation that included telicity and transitivity as planned variables. Those

experiments did not include manner of motion verbs. However, the results of the experiments in this dissertation are consistent with Stevenson and Merlo's findings in the sense that the smallest RREs were found on *by* for optionally transitive telic (alternating unaccusative) verbs.

McKoon and Ratcliff (2003) also have investigated the degree of processing difficulty associated with different verb types in the reduced relative clause construction. Their experiments are also not truly comparable with the experiments in this dissertation, because like Stevenson and Merlo, they used manner of motion verbs for their unergative verbs. McKoon and Ratcliff argue that the reduced relative clause construction has a meaning. They claim that a crucial part of its meaning is that the event associated with the embedded verb was externally caused. They claim that the lexical semantics of an unergative manner of motion verb makes it incompatible with the reduced relative clause structure. They present experimental data showing that reduced relatives with unergative verbs are indeed very difficult to understand. Their results are entirely compatible with the results in this dissertation. However, the results in this dissertation suggest that there are other aspects of verb event structure which have significant effects on the RRE aside from the factor that McKoon and Ratcliff investigate (internal causation associated with manner of motion verbs). McKoon and Ratcliff's results, Stevenson and Merlo's results, and the results in this dissertation all agree on one thing: alternating unaccusative verbs (the optionally transitive telic condition) in the reduced relative construction cause less processing difficulty than other verb types.

8.8 Significance of the Findings for Linguistics

The finding that inherent verb telicity has immediate effects on the comprehension of reduced relative clause sentences provides psycholinguistic evidence for telicity and for event structure more generally. This finding suggests that information about whether a verb is inherently telic or not is stored in the verb's lexical entry. The research also addresses the issue of whether telicity information and verb transitivity information are both stored in the lexicon or whether one of these types of verb information is derived from the other. The results show that telicity and transitivity have very different, essentially opposite effects on the comprehension of reduced and unreduced relative clause constructions throughout the critical region. Deriving one type of verb information from the other would presumably take processing time, but the effects of telicity and transitivity can be seen immediately during on-line comprehension. Thus, the results indicate that verb telicity and verb transitivity are both represented in the lexicon.

The results support theories of syntax that include telicity as a feature that has syntactic implications and theories that relate telicity to the theta roles that noun phrases bear. The results presented in this dissertation cannot differentiate between two prevailing and somewhat opposing views of lexical semantics. One view is that telicity is a syntactic feature and that aspectual semantics ends up being determined by the syntactic structure that results from syntactic derivation. This is a very simplified description of the view held by many researchers on the syntax-semantics interface, such as Travis (1994), Borer (1998), Ritter and Rosen (1998), Harley and Noyer (2000), Sanz (2000), Harley and Folli (2002), and many others. The other view is that telicity is purely semantic and that

semantics determines the syntactic structure of sentences. This is a very simplified description of the view held by Jackendoff (1990, 1997) and Levin and Rappaport-Hovav (1995).

The NP trace results found by Sanz (2000) and Friedmann et al. (2003) and discussed above support the view that telicity is syntactic. A number of syntacticians have argued that all telic intransitive verbs have an unaccusative structure (Tenny, 1992; Sanz, 2000; Folli, 2001, others). Sanz's and Friedmann et al.'s results support this claim by showing evidence of an NP trace following unaccusative verbs. However, the results do not rule out the possibility that telicity is primarily semantic and that the conceptual structure of a telic verb is what makes it have an internal argument.

8.9 Significance for Language Processing Models

The primary implication of the experimental results in this dissertation for language processing models is that inherent verb telicity should be included as a factor in them. Exactly how and whether this can be done in various preexisting models is discussed in detail in chapter 7. As explained earlier in this chapter, the models which do the best job of explaining the results are constraint-based models and Townsend and Bever's (2001) analysis-by-synthesis model, the Late Assignment of Syntax theory.

The finding that telicity affects on-line comprehension is relevant to a very old debate in the language processing field: Is syntactic processing independent? Does semantics have immediate effects on syntactic parsing? Developments in syntactic theory over at least the past decade have put a new twist on this debate. In recent syntactic

theory, some semantic features, such as telicity, past tense, and animacy in some languages are now considered to be syntactically relevant. Syntactically relevant semantic features are operative throughout the syntactic derivation. Does the existence of syntactically relevant semantic features settle or nullify the debate over whether syntactic processing is independent? My opinion is that it really does neither. The existence of syntactically relevant semantic features merely redraws the line between syntax and semantics (plus all of the other things that syntax has been argued to be independent of, such as real world knowledge and context.) One can still claim that syntactic processing is independent if telicity, animacy, and other syntactically relevant semantic information are considered part of the syntax. The new question is whether telicity and other syntactically relevant semantic information are used earlier during processing than real world knowledge and other information that is not syntactically encoded. In short, there's semantics (the syntactically relevant kind), and then there's semantics (the other kind, like real world knowledge.) Now the debate should be about whether one of these kinds of semantics has a role in processing that is privileged over the other.

8.10 Significance for Language Acquisition Research

The finding that verb telicity information is used during on-line language comprehension raises the question of how children learn whether a verb is telic or atelic. The linguistic data on telicity and event structure also raise the question of how children learn the combinatorial aspects of event structure, such as the fact that a bounded incremental theme, such as *the cookie* in *eat the cookie*, can make the event telic. Wagner (2001,

2002) has shown that children are sensitive to telicity information. She found that children between the ages of 2 and 5 often interpret telic events to be perfective and past tense and that they often interpret atelic events to be imperfective and present tense. Her research suggests that although young children have difficulty discriminating between tense and different kinds of aspect, they are sensitive to verb telicity information. Wagner's research suggests that children acquire at least some verb telicity information at a fairly young age. This raises two important questions for future research. The first question is whether children learn the telic versus atelic distinction from the constructions in which adults use them. The second question is whether children have an innate mechanism that predisposes them to be sensitive to the telic versus atelic distinction.

8.11 Significance for Computational Linguistics

The linguistic and psycholinguistic data in this dissertation suggest that telicity may have useful applications in computational linguistics. The linguistic data show that a verb's inherent telicity is related to the noun phrase arguments that are required and to the semantic roles that these noun phrases bear. This suggests that tagging verbs as telic or atelic in corpora, including inherent telicity in the lexical specifications of verbs, and including telicity in phrase structure rules or processing constraints could improve the performance of natural language processing systems.

8.12 Significance for Aphasia Research

The finding that telicity and event structure affect on-line comprehension raises the question of how and whether it is used by speakers with various types of aphasia. Gahl, Menn, Ramsberger, Jurafsky, Elder, Reweg, and Holland (in press) found that simple main clause sentences with alternating unaccusative verbs were rated as equal in plausibility to sentences with unergative verbs by aphasic speakers. Though unaccusativity and telicity are not the same, they are related since many syntacticians have argued that all telic intransitive verbs are unaccusative (Tenny, 1992; Arad, 1998; many others). All of Gahl et al.'s tasks were off-line, in the sense that they did not collect reaction times during an on-line comprehension task. An interesting line of future research would be to investigate whether there are differences between aphasic and non-aphasic speakers in when and whether they employ telicity and event structure information during comprehension.

The study of event structure may be particularly interesting in research on aphasia, because event structure lies at the crossroads of syntax and semantics. It has been claimed that lesions including Broca's area might be associated with more difficulty in processing syntactic structure (among other things) and that lesions including Wernicke's area might be associated with more difficulty in processing semantics or with impaired lexical access. Such a generalization is most likely an oversimplification; however, the study of event structure in aphasia could provide interesting new data on how and whether language areas in the brain might be specialized.

8.13 Summary and Conclusions

In summary, this dissertation proposed and evaluated the Event Structure Processing (ESP) hypothesis. The ESP hypothesis is that event structure information, specifically telicity, is used during on-line language comprehension. Chapter 1 provided the rationale behind the proposal as well the rationale behind the experiments used to test the hypothesis. Chapter 2 provided a description of event structure and telicity as well as the linguistic evidence for their existence. Chapter 3 presented the results from a questionnaire study that aims to make the categorization of verb phrases as telic or atelic more objective. Chapters 4, 5, and 6 provided experimental evidence that inherent verb telicity decreases the amount of processing difficulty associated with the reduced relative effect (RRE). The experimental results in chapter 4 are post-hoc analyses of prior self-paced reading experiments. Chapter 5 presents the results of auditory experiments using the speaker change monitoring paradigm. Chapter 6 presents the results of an experiment involving reading rather than listening using the word maze paradigm. Chapter 7 discusses the ESP hypothesis and the experimental results in relation to a number of pre-existing processing models. In this final chapter, I have summarized the major findings presented in this dissertation, proposed an account of the telicity and transitivity effects taken together, reviewed related research, and discussed the relevance of the research in relation to linguistics, sentence processing, and other research areas. The major research finding that is supported by the experimental results in this dissertation is that comprehension difficulty is decreased in reduced relative clause sentences with telic

verbs compared to sentences with atelic verbs. This finding supports the hypothesis that event structure information is used during on-line language comprehension.