

The Fate of Completed Goal Information in Narrative Comprehension

MARK F. LUTZ AND GABRIEL A. RADVANSKY

University of Notre Dame

Most studies investigating the use of goal information in narrative comprehension have focused more on the increased availability of uncompleted or failed goal information relative to completed goal information. The current study does the opposite. Two experiments were conducted to determine the fate of completed goal information. Three possibilities were considered: (a) that completed goal information remains more available than neutral information, (b) that completed goal information decays to baseline, or (c) that completed goal information is suppressed below baseline. The results of the experiments indicate that completed goal information remains more available than neutral information. These results are related to current theories of how causal information affects narrative comprehension. © 1997 Academic Press

Recent research in narrative comprehension has found that the creation of causal inferences, guided by the goals of a story character, is an important aspect of comprehension (e.g., Albrecht & Myers, 1995; Fletcher & Bloom, 1988; Long & Golding, 1993; Suh & Trabasso, 1993; Trabasso, van den Broek, & Suh, 1989; van den Broek, 1995). Much of this research has focused on how causal relations influence processing when the protagonist's goal has not been completed. When completed goal information is considered in these studies, it is typically used as a comparison to illustrate that failed goal information is maintained and that it drives causal inferences, while completed goal information is not main-

tained because this information is no longer useful. There has been no specific consideration of the fate of completed goal information in the language comprehension process. That is the intent of this paper.

A causal structure of a text, derived by a reader, allows that person to more easily make inferences about actions that occur in a narrative. This causal structure can be thought of as the outline of the goals of the main characters in a story. Such a causal structure can include the goals themselves, the actions then taken to complete those goals, and the outcomes or achievements of those goals. For example, if a primary goal of a protagonist has not been achieved, that is to say, the protagonist has *failed* to achieve the goal, a reader may make causal inferences that relate subsequent actions to the purpose of achieving this goal (Dopkins, Klin, & Myers, 1993; Suh & Trabasso, 1993).

By making inferences, a reader organizes a narrative into a functional memory representation or situation model (Johnson-Laird, 1983, 1989; Kintsch, 1988; Radvansky & Zacks, in press; Trabasso et al., 1989). This situation model represents not only the people, setting, and events mentioned in the text, but also the causal inferences made by the reader (Graesser, Singer, & Trabasso, 1994). This representation

This research was done in partial fulfillment of the Master's degree for the first author at the University of Notre Dame. These experiments were supported in part by a Faculty Research Program grant from the Graduate School at the University of Notre Dame awarded to the second author. Some of the data presented here were presented at the 1995 meeting of the Midwestern Psychological Association in Chicago, IL. We thank Andrea Backscheider, Laura Carlson-Radvansky, and Brad Gibson for their comments on earlier versions of this paper. We also thank Jackie Curiel, Dan Olson, Francie Schmul, and Melissa Wein for their assistance in collecting the data. Address Correspondence and reprint requests to G. A. Radvansky, Department of Psychology, University of Notre Dame, Notre Dame, IN 46556. E-mail: Gabriel.A.Radvansky.1@nd.edu.

can be updated as the reader proceeds through the narrative and as new inferences are generated. As a result, this representation aids both in the integration of new information with previous information and in the recall of that information. Research has shown that the probability that specific information will be recalled on a subsequent memory test is influenced by the number of plausible causal connections (Black & Bower, 1980; Fletcher & Bloom, 1988; Myers & Duffy, 1990; Myers, Shinjo, & Duffy, 1987; Trabasso, Secco, & van den Broek, 1984; Trabasso & van den Broek, 1985). In general, an increased number of causal connections increases the probability of recall except at very high levels of interconnectivity, when recall may suffer because information is so interconnected that readers perform fewer elaborations on it, and therefore remember it less well (Myers & Duffy, 1990; Myers et al., 1987). Thus, the causal structure derived from a narrative, and the inferences generated, assist not only on-line language comprehension but also later memory retrieval.

Several theories have been put forward to explain when, where, and why inferences related to causal structures are made. For this paper, we will consider three recent views that can readily account for such inferences: the constructionist, resonance, and current state selection models. We are specifically interested in accounts that can explain the finding that information about a failed goal is more available than information about a completed goal. Based on these theories, we derive predictions about the fate of completed goal information. Specifically, does information about completed goals remain available, does it decay over time, or is it suppressed?

Constructionist View

Because our experiments are modeled after a method developed by Suh and Trabasso (1993), we first consider the position from which these studies developed, the constructionist view (see Graesser et al., 1994, for a review). This theory states that readers construct a situation model of a text to allow for

the on-line creation of global inferences. Readers are aided in making inferences by keeping available text information that they believe to be necessary for comprehension. This process is guided by general problem-solving operations that are governed by a search for meaning.

To analyze the potential causal connections in narratives in a systematic manner, Trabasso, van den Broek, and colleagues (Trabasso & van den Broek, 1985; Trabasso et al., 1989) developed the Recursive Transition Network (RTN) model. According to this theory, narratives are represented by a causal network of relevant clauses. Clauses can be categorized into settings, initiating events, goals, actions, outcomes, and reactions. An example of a story containing these different types of clauses is presented in Table 1a.

The *setting* introduces the characters and background conditions of the narrative (Sentence 1, introduction of a girl named Betty). The *initiating event* is the first event in the narrative and is usually a result of unstated causes (Sentence 2, Betty finds out about her mother's birthday). *Goals* are expressions of wanting to achieve desired states or avoid undesired states (Sentence 3, Betty wants to get her mother a present). Goals are usually followed by *actions*, which are attempts to attain a desired goal or avoid an undesired goal related to a particular state, activity, or object (Sentences 4 and 5, Betty goes to the store but finds everything is too expensive). Goal actions can result in successful or failed *outcomes* (Sentence 6, Betty could not buy her mother a present). Outcomes may then cause *reactions* which are responses to the outcome. These reactions can lead to other goals and actions, especially when a previous goal has not been satisfied (Sentence 7, Betty felt sorry because she did not get her mother a present).

If a reader encounters a clause that contains a goal, this can lead to the establishment of a goal plan, and any subsequent goals or sub-goals that fit the goal plan will be causally connected to the original goal. This is consistent with the general idea that information that

(a) Failed goal narrative

1. Once there was a girl named Betty.	(Setting)
2. One day Betty found that her mother's birthday was coming soon.	(Initiating Event 1)
3. She really wanted to give her mother a present.	(Goal 1)
4. She went to the department store.	(Action 1a)
5. She found out that everything was too expensive.	(Action 1b)
6. She could not buy anything for her mother.	(Outcome 1)
7. She felt sorry.	(Reaction 1)
8. Several days later, Betty saw her friend knitting.	(Initiating Event 2a)
9. She was good at knitting too.	(Initiating Event 2b)
10. She decided to knit a sweater.	(Goal 2)
11. She selected a pattern from a magazine.	(Action 2a)
12. She followed the instructions in the article.	(Action 2b)
13. Finally, Betty finished a beautiful sweater.	(Outcome 2)
14. She pressed the sweater.	(Action 3a)
15. She folded the sweater carefully.	(Action 3b)
16. She gave the sweater to her mother.	(Outcome 3)
17. Her mother was excited when she saw the present.	(Reaction 3)

(b) Completed goal narrative

1. Once there was a girl named Betty.	(Setting)
2. One day Betty found that her mother's birthday was coming soon.	(Initiating Event 1)
3. She really wanted to give her mother a present.	(Goal 1)
4. She went to the department store.	(Action 1a)
5. She found a pretty purse.	(Action 1b)
6. She bought her mother the purse.	(Outcome 1)
7. Her mother was very happy.	(Reaction 1)
8. Several days later, Betty saw her friend knitting.	(Initiating Event 2a)
9. She was good at knitting too.	(Initiating Event 2b)
10. She decided to knit a sweater.	(Goal 2)
11. She selected a pattern from a magazine.	(Action 2a)
12. She followed the instructions in the article.	(Action 2b)
13. Finally, Betty finished a beautiful sweater.	(Outcome 2)
14. She pressed the sweater.	(Action 3a)
15. She folded the sweater carefully.	(Action 3b)
16. She put it in the closet for the next time she went out.	(Outcome 3)
17. She was very happy.	(Reaction 3)

(c) Example of a neutral narrative

1. Once there was a girl named Betty.	(Setting 1a)
2. One day Betty found that her mother's birthday was coming soon.	(Setting 1b)
3. She bought the present she wanted to give her mother.	(Setting 1c [Goal 1])
4. Then, she went to the department store.	(Action 1a)
5. She found a pretty purse.	(Action 1b)
6. She bought herself the purse.	(Outcome 1)
7. She was very happy.	(Reaction 1)
8. Several days later, Betty saw her friend knitting.	(Initiating Event 2a)
9. She was good at knitting too.	(Initiating Event 2b)
10. She decided to knit a sweater.	(Goal 2)
11. She selected a pattern from a magazine.	(Action 2a)
12. She followed the instructions in the article.	(Action 2b)
13. Finally, Betty finished a beautiful sweater.	(Outcome 2)
14. She pressed the sweater.	(Action 3a)
15. She folded the sweater carefully.	(Action 3b)
16. She put it in the closet for the next time she went out.	(Outcome 3)
17. She was very happy.	(Reaction 3)

is interpreted as referring to a common situation is integrated into a common situation model (e.g., Radvansky, Spieler, & Zacks, 1993; Radvansky & Zacks, 1991; Zwaan, Magliano, & Graesser, 1995).

If a goal has not been satisfied, a hierarchy of failed or unsatisfied goals can be generated. If there is more than one goal or subgoal that is unsatisfied, the reader evaluates the current action against the most recent unsatisfied goal. When a goal or subgoal has been satisfied, the reader then evaluates actions against the next goal that has not been satisfied. This continues until the first goal of a goal plan has been satisfied. Therefore, if a goal has not been completed, this goal information should remain available. Accordingly, failed goals should be more available than successfully completed goals, because they are unsatisfied.

A set of experiments by Suh and Trabasso has demonstrated that unsatisfied goals are more accessible than completed or satisfied goals (Suh & Trabasso, 1993; Trabasso & Suh, 1993). They had participants read narratives with both failed and completed goal versions. Each version consisted of three separate episodes. An *episode* represents each series of actions or attempts that are taken in order to accomplish a desired goal or subgoal and result in some outcome. In the failed goal version of the story (Table 1a), Betty has a goal of getting her mother a birthday present. However, she is unable to get her mother a present in the first episode (Sentences 1–7), thereby failing to accomplish her goal (henceforth, Goal 1). So, in the second episode (Sentences 8–15), Betty's reasons for making a sweater should be attributed to the goal of getting her mother a present, because it has not been satisfied. In the third episode, which for this version may be thought of as a continuation of the first episode, Betty gives her mother the present (Sentences 16–18).

In the completed goal version (Table 1b), Betty buys her mother a present during the first episode (Sentences 1–7), thereby accomplishing this goal. In the second episode, Betty makes a sweater, but because she has

already purchased a present, the reasons for making it change (Sentences 8–15). In the third episode, Betty puts the sweater away in her closet (Sentences 16–18). In both versions, the second episode is identical, but the reasons attributed to the actions are different depending on whether the first goal has been satisfied. Accordingly, different story versions were designed to encourage different inferences depending on which version was read.

In order to determine whether people maintained the availability of the failed goal information, Suh and Trabasso (1993) probed for this information while people were reading narratives. People read three completed and three failed goal narrative versions, along with 10 filler narratives. Reading was self-paced. During the course of reading, each narrative was interrupted with two probe questions. The probes required people to assess whether a statement was true based on what had appeared earlier in the narrative. For example, a probe might ask, "Had Betty wanted to get her mother a birthday present?" Only the first probe on the experimental trials were used in the analysis, and the probe always referred to Goal 1 information. Response time to these probes was used as the dependent measure.

The central result of Suh and Trabasso's (1993) study was that people responded slower to probes that referred to completed goal information than to probes that referred to failed goal information. These results suggest that when reading a narrative, failed goal information is more available than completed goal information.

Resonance Model

Another theory that attempts to explain how readers make causal inferences is the resonance model (Albrecht & O'Brien, 1993; Albrecht & Myers, 1995; Dopkins et al., 1993; Klin, 1995; Myers, O'Brien, Albrecht, & Mason, 1994). According to this view, inference generation is primarily driven by more automatic memory retrieval processes. Specifically, information that is encountered makes contact with memory traces that are consistent

with the currently available information through a global matching mechanism (see Clark & Gronlund, 1996, for a review). Albrecht and Myers (1995, p. 1460) detailed three aspects of a text that influence whether an inference is made: (a) the interconnectedness of the to-be-accessed set of propositions, (b) the length of time since the information was last referenced, and (c) the degree to which the current information overlaps the information in the long-term memory trace.

The resonance model can be used to explain the pattern of results observed by Suh and Trabasso (1993). Because the Goal 1 sentence and the memory probes in the two versions are identical, the trace overlap in both the failed and completed goal conditions should be equivalent. However, in failed goal stories, information that occurs after Outcome 1 may activate other traces of previous story propositions, thereby allowing these sets of information to be interconnected. In contrast, when these traces become activated in the completed goal stories, they are much less likely to be interconnected, even if they are initially activated. As a result, the probed-for information is more accessible in the failed goal versions because it is more highly interconnected than the completed goal versions. Also, because the goal information is likely to be considered relevant for a longer period of time in the failed goal versions, less time will have elapsed since the last Goal 1 activation in the failed goal versions than in the completed goal versions. As such, this should also allow the Goal 1 information to be accessed faster in the failed goal versions.

Current-State Selection Strategy

The third view of inference generation is the current-state selection strategy (Bloom, Fletcher, van den Broek, Reitz, & Shapiro, 1990; Fletcher & Bloom, 1988; Fletcher, Hummel, & Marsolek, 1990). According to this view, the ease of comprehension is regulated by whether information is currently in working memory. This model assumes that working memory contains the most recently

processed proposition, and the proposition referring to the most recently unsatisfied causal antecedent. All other information is stored in long-term memory. Readers make on-line causal inferences only if the causal antecedent is currently in working memory. Otherwise, readers must retrieve that information from long-term memory.

The current-state selection theory can also be used to explain the results of Suh and Trabasso (1993). Specifically, according to this view, readers will retain the goal information in working memory in the failed goal versions because no causal consequent for this information would have been encountered. As such, this information is highly accessible. Conversely, in the completed goal versions, because a causal consequent would have been encountered earlier in the text, people would no longer hold this information in working memory. As a result, the goal information would be less accessible in the completed goal versions than in the failed goal versions because people need to engage in a long-term memory search.

FATE OF COMPLETED GOAL INFORMATION

While it has been shown that failed goal information is more available than completed goal information, it is unclear whether this difference is a result of the heightened availability of the failed goal information, a decreased availability of the completed goal information, or both. Most theorists have suggested that the difference in response time is a result of the greater availability of the failed goal information (e.g., Graesser et al., 1994; Suh & Trabasso, 1993; Trabasso & Suh, 1993), but there is no direct evidence to support this idea. The current experiments explored how the three theories that we have outlined here could explain the difference in the availability of failed and completed goal information. Each of these theories will be considered in light of how failed and completed goal information is related to neutral information.

For the purpose of our experiments, *neutral*

information refers to information that is included in a narrative but is not likely to be a major part of the causal structure of that narrative. As such, it is “dead-end” information (e.g., Black & Bower, 1980; Trabasso et al., 1984; Trabasso & Sperry, 1985; Trabasso & van den Broek, 1985). In the present experiments, the neutral information refers to the same information that is probed for in the failed and completed story versions, such as the idea of Betty wanting to buy her mother a present. Because of this, the same memory probes can be used in all three story versions. However, in contrast with the failed and completed goal versions, when this information is presented in the neutral versions, it is presented so that there are no causal antecedents elsewhere in the text, nor has a goal been established that will need to be satisfied later. As such, the information stands alone relative to the rest of the story and therefore should be less salient in the context of the story. Although there have been previous studies of how neutral information relates to failed and completed goal information using story recall (e.g., Trabasso et al., 1984) and importance ratings (e.g., van den Broek, 1988), ours is the first on-line measure. Neutral information in the current experiments was established by presenting the information in the form of a completed goal outcome. Thus, this information could be thought of as a transient goal that served as additional setting information in the story.

An example of a neutral narrative is presented in Table 1c. In this story, in the first episode (Sentences 1–7), the reader is told that Betty has already completed the goal of getting her mother a present and then goes on to complete the goal of buying a purse. It is unlikely that this purse would be interpreted as a present for her mother because it has already been stated that she had bought her a present. In the second episode (Sentences 8–15), Betty decides to knit a sweater. This information is also not likely to be interpreted as being related to getting a present for her mother. Finally, in the third episode (Sen-

tences 16–18), Betty puts the sweater away in her closet. So, in this story, buying a present for her mother is not used to establish reference with an earlier part of the text, nor is it likely to be maintained in order to make subsequent causal inferences. Given these three text versions (neutral, completed goal, and failed goal), we can now consider what the three theories predict about the fate of completed goal information.

According to the constructionist view, goal information should be more available than neutral information, even after the goal has been satisfied. Information is less available in the neutral versions than the failed goal versions because it does not need to be kept available for future causal inferences. Because the neutral information would be designated as dead-end information, it would be less likely to be included in a reader’s causal structure of the text. If the additional assumption is made that dead-end information is less available than any information that contributes to interconnections among clauses in the causal structure (as suggested by data from free-recall studies (e.g., Black & Bower, 1980; Fletcher & Bloom, 1988; Myers & Duffy, 1990; Myers et al., 1987; Trabasso et al., 1984; Trabasso & van den Broek, 1985), then this view would predict that neutral information would be less available than completed goal information. Therefore, probes for failed goal information should be responded to faster than the completed goal information which should be responded to faster than the neutral information (Failed < Completed < Neutral). This would be consistent with the idea that there is both a heightened availability of failed goal information and a reduced availability of completed goal information.

According to a resonance model, the information in the neutral version is less likely to be integrated into a reader’s representation of the story than the information in either the completed or failed goal versions. As such, it should be less available in the neutral versions. Like the constructionist view, the resonance model predicts that probes for failed

goal information should be responded to faster than the completed goal information which should be responded to faster than the neutral information (Failed < Completed < Neutral).

According to the current-state selection strategy, the information in both the neutral and completed goal versions would no longer be needed for future comprehension because there is a causal consequence, and so this information is no longer maintained and would be removed from working memory. As a result, completed goal information is predicted to be at a baseline level of availability, just as neutral information. In contrast, failed goal information is needed for future comprehension and so it remains available. As such, neutral and completed goal information would be equally available, yet less available than failed goal information. Therefore, response times to completed goal and neutral information should be equivalent, and both should be slower than the failed goal information (Failed < Completed = Neutral). This would suggest that the difference in the availability of failed and completed goal information is due only to an increased availability of the failed goal information.

Another possibility to consider that is not addressed directly by any of the theories is that there could be a suppression of the completed goal information below the level of availability of the neutral information. Such suppression of the completed goal information could reduce the possibility that a reader will make erroneous inferences that would relate new information to an already completed goal. Furthermore, because failed goal information is likely to be needed to draw future inferences related to the goal structure of the narrative, it would be more available than neutral information. Therefore, failed goal information would be responded to faster than the neutral information, which in turn would be responded to faster than completed goal information which would be suppressed (Failed < Neutral < Completed). This would be consistent with a notion that there is both a heightened availability of the failed goal information

and a reduction in the availability of the completed goal information. Although, none of the theories that we are considering here explicitly implements some sort of suppression mechanism, such a process could be added to the constructionist and resonance models without changing the general spirit of those models.

PRESENT EXPERIMENTS

Because the neutral stories offer a baseline from which to assess the fate of completed goal information, we begin by describing a pilot study that verifies the neutrality of these story versions. Specifically, we assessed whether people would stop thinking about the neutral story information soon after it was mentioned using a verbal protocol analysis. In Experiment 1, the availability of neutral, failed, and completed goal information was assessed to determine which of the three theories best accounts for the fate of completed goal information. Experiment 2 examined a possible bias in Experiment 1 based on the notion of differential decay times.

Pilot Study

The purpose of the pilot study was to assess whether people would make inferences at the appropriate points of the story depending on the condition. People in this study were asked to "think aloud" while they read each story by typing their thoughts into a computer. This think aloud procedure has been adopted to help assess what inferences readers make (e.g., Magliano & Graesser, 1991; Suh & Trabasso, 1993; Trabasso & Suh, 1993).

For the failed or completed goal versions, we would expect that the pattern of inferences would follow that observed by Suh and Trabasso (1993). Specifically, for the failed goal versions, there should be a high rate of Goal 1 mentions initially. This rate should then decrease while the protagonist pursues a subgoal. Then, after the subgoal is completed, the rate of Goal 1-related mentions should increase again. For the completed goal versions, there should be a high rate of Goal 1 mentions until the goal is completed, at which point

mentions should decrease below the failed goal version rate and never rise again. Finally, if the neutral versions lead people not to think of the Goal 1 information after it has occurred in the story, then it should be mentioned less often than in the completed goal versions soon after it is presented. Note that this information is setting information in the neutral version but will be referred to as Goal 1 information to ease comparison with the other story versions. After the goal has been achieved in the completed goal versions, there should be no difference in the number of mentions between the completed and neutral versions.

Method

Readers. Forty-eight native English speakers were recruited from the University of Notre Dame subject pool in exchange for partial class credit.

Materials. There were 12 experimental narratives. Eight of these were those used by Suh and Trabasso (1993) and the other 4 were new narratives. Each narrative had three versions: (a) failed goal, (b) completed goal, and (c) neutral.¹

The narratives ranged from 13 to 17 sentences in length ($M = 14.4$, $SD = 1.4$). Each narrative contained three episodes. *Failed goal narratives* had the following structure: (a) in Episode 1, Goal 1 was introduced but not attained, (b) in Episode 2, Goal 2, a sub-

goal of Goal 1, was introduced and attained, and (c) in Episode 3, the actions lead to a successful outcome of Goal 1. Because of the completion of Goal 1, Episode 3 in the failed goal version can be considered a resumption of Episode 1. In the failed goal version of the story in Table 1a, Betty wants to buy her mother a present at the department store, but everything is too expensive (Episode 1). Betty decides to knit a sweater (Episode 2). Finally, she gives her mother the sweater, thereby resolving Goal 1 (Episode 3/resumption of Episode 1).

Completed goal narratives had the following structure: (a) in Episode 1, Goal 1 was introduced and attained, (b) in Episode 2, Goal 2 was introduced and attained, and (c) in Episode 3, the actions lead to an outcome of Goal 3. For example, in the completed goal version in Table 1b, Betty wants to get her mother a present, so she goes to the department store and buys her mother a purse (Episode 1). Betty then decides to knit a sweater (Episode 2). Betty puts the sweater away in her closet (Episode 3).

Neutral narratives were created through alterations of the completed goal versions. First, the Goal 1 information was introduced as already having been completed and, as such, serves as additional setting information. Second, the outcome of Goal 1 was changed in order to maintain story coherence. For example, "Betty bought her mother a purse" was changed to "Betty bought herself a purse" because in the neutral version, Betty had already bought her mother a present. Neutral narratives had the following structure: (a) in Episode 1, Goal 1 was presented as an outcome of a goal the protagonist had already accomplished and another task was completed in the outcome sentence, (b) in Episode 2, Goal 2 was introduced and attained, and (c) in Episode 3, the actions lead to a successful outcome of a final goal. For example, in the neutral version in Table 1c, Betty bought her mother a present, as she wanted to, and then buys herself a purse at the department store (Episode 1). Betty decides to knit a sweater

¹ One awkward characteristic of the narratives used in the pilot study, unlike the stories in the tables, was that the story protagonist's name was mentioned in most of the story sentences. It has been shown that the mention of a proper name is processed somewhat differently than a pronoun, perhaps leading to some cost in processing (Gordon & Chan, 1995; Gordon, Grosz, & Gilliom, 1993; Vonk, Hustinx, & Simons, 1992). This should not pose a difficulty in interpreting the data from the pilot experiment because (a) this aspect of the stories was present in all versions, (b) the repeated name penalty seems to be confined to more on-line processing, and (c) because the data for the failed and completed goal conditions parallel that observed by Suh and Trabasso (1993). The stories used in Experiments 1 and 2 mentioned the protagonist's name much less often, and pronouns were used more often to make the stories seem more natural.

(Episode 2). Betty puts the sweater away in her closet (Episode 3). By presenting the Goal 1 information in Episode 1 as an outcome (Betty had already bought the present), this information should not be maintained and used to make inferences later in the story because the action has been completed.

In addition to the 12 experimental stories, there was also a practice story. This practice narrative was a story about a character named Ivan who learned how to use a sword to protect his village from a dragon. This story was 14 sentences long and had a structure similar to the failed goal story versions. Protocols for this story were not recorded.

Design and procedure. These narratives were assessed using verbal protocols. The procedures were similar to those used by Suh and Trabasso (1993, Exp. 1) except that the readers typed their responses into a computer rather than vocalizing them. Although this difference in the manner of responding may have produced some differences in what people report, as will be seen, the resulting pattern of data for the failed and completed goal versions closely paralleled that reported by Suh and Trabasso.

At the beginning of each story, a warning screen consisting of the word "STORY" was presented in yellow to inform the person that a new story was about to be presented. People read narratives one sentence at a time, at a self-paced presentation rate. After reading each sentence, they typed their understanding of the narrative into the computer and then pressed the enter key. Before moving on to the next sentence, the preceding sentence and the person's protocol were erased. Each person first received a practice narrative, followed by three experimental narratives, one of each type in a random order. The 12 narratives were cycled through the readers so that all narratives and their versions were read equally often.

For the statistical analyses in this and the subsequent experiments, those tests marked with a subscript 1 used subjects as the random variable and those with subscript 2 used items

as the random variable. Also, unless otherwise mentioned, a rejection level of $p < .05$ was used for all statistical tests.

Results and Discussion

The general pattern of results supports the idea that people had the greatest proportion of Goal 1 information mentions in the context of failed goal narratives ($M = .47$; $SD = .20$), fewer in completed goal narratives ($M = .33$; $SD = .16$), and the least in the neutral narratives ($M = .24$; $SD = .13$). These overall mean proportions significantly differed from one another (all $ps < .05$).

The pattern of Goal 1 mentions across story positions is illustrated in Fig. 1. In the figure, "G" stands for a goal sentence, "O" stands for an outcome sentence, "I" stands for the intervening sentences between the goal and outcome sentences, and "B" stands for the sentences between the Outcome 1 and Goal 2 sentences. The values at the "I" and "B" positions represent averages across all sentences falling into those categories. This was done because the stories differed in the number of sentences in each of these categories.

The pattern of data for the failed and completed goal story versions closely follows that reported by Suh and Trabasso (1993). In particular, for the failed goal stories, Goal 1 mentions start off high (G1 through O1) and then decrease as the protagonist switches to a subgoal (B through I2). Then, when the subgoal is completed, the initial goal returns to the fore (O2 through O3). For the completed goal stories, although mentions start off high (G1 through O1), after the goal has been successfully completed (B through O3), mentions drop off dramatically and remain very infrequent. Participant mentions of the goal do not differ in the failed and completed goal versions up through the "B" section (all $ps > .10$). After that point, there are fewer mentions of the Goal 1 information in the completed goal versions than the failed goal version (all $ps < .05$, except for position I3 in the items analysis, $t_2(11) = 1.65$, $p = .13$). This is consistent with the idea that people continue to




FIG. 1. Mentions of the Goal 1 information for the failed goal, completed goal, and neutral story versions as a function of story position. "G" stands for a goal sentence, "O" stands for an outcome sentence, "I" stands for the intervening sentences between the goal and outcome sentences, and "B" stands for the sentences between the Outcome 1 and Goal 2 sentences. The values at the "I" and "B" positions represent averages across all sentences falling into that category. This was done because the stories differed in the number of sentences in each of these categories.

maintain the Goal 1 information in the failed goal version to help them understand the character's actions and to look for a subsequent resolution of that goal. However, in the completed goal versions, this information becomes less relevant to understanding the character's actions, and so it is not maintained after the goal has been completed.

As for the neutral stories, although the proportion of Goal 1 mentions was highest in this condition the first time the goal is mentioned at G1 (in fact, significantly more often than the failed goal versions, $t_1(47) = 2.42$, $t_2(11) = 2.00$, $p = .07$, but not the completed goal versions, $t_1(47) = 1.00$, $t_2(11) = 1.15$, $p > .10$), the number of mentions drops off dramatically after that. From section I1 through section B, the proportion of Goal 1 mentions is less in the neutral versions than in the failed and completed goal versions, all $ps < .05$, except for the difference between the completed goal and neutral narratives at I1, $t_1(47)$

$= 1.51$, $p = .14$, $t_2(11) = 1.52$, $p = .16^2$. After section B, the proportion of Goal 1 mentions in the neutral versions was always significantly below the rate in the failed goal versions, all $ps < .05$, but not significantly different from the rate of mentions in the completed goal versions, all $ps \geq .18$. This is consistent with the idea that in the neutral story versions, although people may continue to think about the Goal 1 information immediately after it is mentioned (e.g., at position I1), it is removed from the stream of processing faster than in-

² In a subsequent analysis we combined the data from the failed and completed goal versions at positions G1 and I1 because these versions were identical up to that point. The combined data were then compared against the neutral version. While the proportion of Goal 1 mentions at position G1 occurred more often for the neutral ($M = .94$; $SD = .25$) than the combined versions ($M = .82$; $SD = .30$), $t(47) = 2.04$, at position I1 there were fewer mentions in the neutral ($M = .39$; $SD = .39$) than in the combined versions ($M = .60$; $SD = .33$), $t(47) = 2.90$.

formation in the completed and failed goal versions until this goal is attained by the character in the completed goal versions. After this point Goal 1 mentions are very infrequent.

In sum, the data from this pilot study suggest that the Goal 1 information in our neutral story versions is not thought about much beyond the time it is first mentioned. This presumably occurs because it does not need to be maintained to help make subsequent goal-related inferences. As such, it would seem that story versions of this type would be an appropriate baseline with which to compare the availability of Goal 1 information in the failed and completed goal versions.

Experiment 1

Experiment 1 sought to determine which of the three theories—the constructionist view, resonance model, or current-state selection strategy—best accounts for the fate of the completed goal information. To get a baseline against which to measure the availability of the completed goal information, the neutral story versions were used. These neutral story versions contained the same information that was probed for in the completed and failed goal versions but were designed so that they would lead readers to consider it as dead-end information. As a reminder, the constructionist and resonance models would be supported by data that show that completed goal information is more available than neutral information, whereas the current-state selection strategy would be supported by data that show that neutral and completed goal information are equally available. This experiment also tests whether a suppression version of either the constructionist or resonance models should be considered.

Method

Readers. Thirty-six native English speakers were recruited from the University of Notre Dame subject pool and were given \$5 or partial class credit.

Materials. There were 5 practice, 18 experimental, and 30 filler stories used in this experi-

ment. These stories included some of the narratives from the pilot study as well as some new narratives. The new experimental narratives had the same general structure as those used in the pilot study. The experimental narratives ranged from 13 to 17 sentences in length ($M = 14.3$, $SD = 1.1$). The practice narratives ranged from 13 to 15 sentences in length ($M = 13.8$, $SD = 0.8$), while the filler narratives ranged from 11 to 15 sentences in length ($M = 13.4$, $SD = 1.0$). The practice and filler narratives had a structure similar to that of the experimental narratives. This was done so that people could not anticipate which narratives were the experimental narratives.

All narratives had two memory probes. These probes took the form of a question (e.g., “Did Betty want to buy her mother a present?”). For the experimental narratives, the same probes were used for all three versions. Furthermore, in the experimental narratives, there were two types of probes. The first was the critical probe that asked about Goal 1 information. The critical probe always occurred in the third episode, just before the successful outcome of Goal 1 in the failed goal version. This placement was based on the results of Suh and Trabasso (1993), which showed the greatest difference in response time between the completed and failed goal versions at this point of the narrative. This critical probe served as a measure of the availability of Goal 1 information. The correct answer to these probes was always “yes.” The second probes in the experimental narratives were fillers and asked about some portion of the story that was unrelated to the Goal 1 information. These filler probes could be either true or false.

For the filler and practice stories, both probes were fillers and they occurred at different locations of the story. The locations of the filler probes were designed to counterbalance the late placement of the critical probes, so that people could not anticipate when in the narrative the first probe would occur. Furthermore, the first filler probes were counterbalanced so they could be either true or false. Also, across all narrative types, the first probe

was located approximately one-third of the time in the first third of the narrative, one-third of the time in the middle third of the narrative, and one-third of the time in the last third of the narrative. Across all narratives there were an equal number of true and false probes.

Procedure. Narratives were presented, one sentence at a time, on a Gateway 2000 486 computer. Each sentence appeared in white halfway down a black screen, in 40-column mode. Readers were told to try to comprehend each sentence as it appeared. Presentation was self-paced. People pressed the space bar with their left hand to advance to the next sentence. Readers were told that they would be interrupted by probe questions during the course of reading, and that they should answer each probe question as quickly and as accurately as possible. Responses were made by pressing one of two buttons on a computer mouse with their right hand. The left button was marked with a "Y" for "yes, this is true," whereas the right button was marked with an "N" for "no, this is false." Readers were instructed to always keep their right hand on the mouse. The probes were presented in red to distinguish them from the rest of the story.

People were given 5 practice narratives to familiarize themselves with the procedure but were not told that these narratives were practice. The experimental and filler narratives were given in an order that was randomized for each reader, with the constraint that each person receive all 18 experimental narratives, 6 of each version. Before the experiment, it was decided that any response times to the probe questions that were shorter than 500 ms and longer than 10 s were to be identified as anticipations and lapses of attention, respectively, and would be excluded from the analyses.

After each story, people were given feedback on the accuracy of their responses. They were then given a distracter task of counting backward by 3s (e.g., 364, 361, 358, . . .). This continued until 10 s had elapsed. The word "STORY" appeared on the screen as a

signal to write down the last number they had counted to and to proceed with the next story by pressing the space bar. The time spent on the distracter task served to decrease a person's memory of the previous story, thus reducing the chance of that information interfering with the next story. Each experimental session lasted approximately 1 h.

Results

Response times. None of the response time data met the criteria for exclusion from the analysis in Experiment 1. Recall that the RT data (as well as the error rate data) reflect responses to an identical probe question in all story versions. Failed goal information was responded to faster ($M = 1994$ ms, $SD = 512$) than completed goal information ($M = 2136$ ms, $SD = 549$), $t_1(35) = 2.25$, $t_2(17) = 2.31$. This supports the notion that failed goal information is more available than completed goal information and is consistent with Suh and Trabasso (1993). Furthermore, responses to the neutral information took longer ($M = 2387$ ms, $SD = 614$) than to completed goal information, $t_1(35) = 2.83$, $t_2(17) = 3.39$. This overall pattern (i.e., Failed < Completed < Neutral) is consistent with the constructionist and resonance theories.

Error rates. Readers made an average of 3.2% ($SD = 8.8$) errors on the failed goal stories, 1.9% ($SD = 5.3$) errors on the completed goal stories, and 12.5% ($SD = 10.8$) errors on the neutral stories. An analysis revealed that more errors were made on the neutral story versions than the failed, $t_1(35) = 3.80$, $t_2(17) = 1.86$, $p = .08$, and completed goal story versions, $t_1(35) = 5.30$, $t_2(17) = 2.18$. The error rates did not significantly differ between the failed and completed goal story versions, $t_1(35) < 1$, $t_2(17) = 1.37$, $p = .19$. These results support the notion that the neutral information is less available than both completed and failed goal information.

Discussion

The results of Experiment 1 replicate the findings of Suh and Trabasso (1993). Specifi-

cally, failed goal information was more available than completed goal information. Moreover, this experiment also found that completed goal information was more available than neutral information. The error rate data demonstrated a similar pattern with fewer errors being made in the failed and completed goal versions than the neutral versions. This suggests that the information was more available in the failed and completed goal versions than in the neutral versions. This pattern of data also removes the need to consider a suppression version of either the constructionist or resonance models.

These results are consistent with the constructionist model, which states that readers attempt to causally relate information previously presented in the text with the current text information. As people read, they make inferences as part of the development of a causal structure for the narrative. Information from an uncompleted goal should be made available whenever it is the most recent uncompleted goal. This accounts for the response time difference in the failed and completed goal versions. If an additional assumption is made that any information that is used in drawing causal linkages is more available than other information, then this view is able to account for the response time difference in the completed and neutral versions.

These results were also consistent with the resonance model, which predicts that when information in a narrative is more integrated, that information is more likely to activate the memory traces from earlier in the narrative. As a result, the more integrated the information is, the more accessible that information should be. The most integrated narrative version was the failed goal version, followed by the completed goal version, with the neutral version being the least integrated. This description conforms to the pattern of response times observed in Experiment 1.

The results of Experiment 1 do not support the current-state selection strategy. According to this model, goal-related information is either more available, as when it is being held

in working memory, or is in a less available state in long-term memory. It does not predict the graded pattern of availability observed in Experiment 1. Because the completed goal information was more available than the neutral information, this view was not supported.

Experiment 2

In Experiment 1, completed goal information was more available than neutral information. However, an examination of the narratives reveals a bias that could have potentially obscured a finding in favor of the current-state selection strategy. Specifically, in the completed goal versions, the Goal 1 information was relevant further into the story than it was in the neutral versions. For example, if the probed-for information was presented in Sentence 3 for both versions, while this information was finished in Sentence 3 of the neutral version, it was not finished until Sentence 5 in the completed goal version when that goal was resolved. Therefore, the number of sentences and the amount of reading time that occurred between the end of the use of the probed-for information and the probe location were not equal. Simply put, information might be more available in the completed goal version because it had less time to decay from working memory.

In order to test this counterexplanation, in Experiment 2 the distance between the probe and the end of the probed-for information was equated by adding additional lines to the completed goal story versions (see Table 2). The extra lines (e.g., X1, X2, and X3) made the distance between the probe and the end of the probed-for information equivalent. To avoid confusion, the story versions with the extra lines were called *long completed goal versions*, while the old completed goal versions were called *short completed goal versions*.

For Experiment 2, if the difference in probe distance does have an effect, then the response time advantage of the completed goal versions relative to the neutral versions should be reduced or eliminated in the long completed goal versions but not in the short completed

TABLE 2
AN EXAMPLE OF A LONG COMPLETED GOAL NARRATIVE

1. Once there was a girl named Betty.	(Setting 1)
2. One day Betty found that her mother's birthday was coming soon.	(Initiating Event 1)
3. She really wanted to give her mother a present.	(Goal 1)
4. She went to the department store.	(Action 1a)
5. She found a pretty purse.	(Action 1b)
6. She bought her mother the purse.	(Outcome 1)
7. Her mother was very happy.	(Reaction 1)
X1. Betty went back to the mall to meet a friend.	(Action X)
X2. Her friend had on a new sweater.	(Setting X)
X3. She wondered where her friend had bought the sweater.	(Reaction X)
8. Several days later, Betty saw her friend knitting.	(Initiating Event 2a)
9. Betty was also good at knitting.	(Initiating Event 2b)
10. She decided to knit a sweater.	(Goal 2)
11. She selected a pattern from a magazine.	(Action 2a)
12. She followed the instructions in the article.	(Action 2b)
13. Finally, Betty finished a beautiful sweater.	(Outcome 2)
14. She pressed the sweater.	(Action 3a)
15. She folded it carefully.	(Action 3b)
16. She put it in the closet for the next time she went out.	(Outcome 3)
17. She was very happy.	(Reaction 3)

Note. Sentences with numbers beginning with an "X" signify the additional lines added to equate probe distance with the neutral narratives.

goal versions. Such a result would lend support to the current-state selection strategy. However, if the other theories are correct, then completed goal information should be clearly more available than neutral information, even in the long completed goal story versions.

Experiment 2 also served as a test between the resonance and constructionist models. Specifically, the resonance model states that information availability is mediated by the amount of time that has elapsed since the last time that information was referenced (Albrecht & Myers, 1995). This is because, after the information is no longer being actively used, the strength of a memory trace decays over time. As a result, information in the short completed goal versions should be more available than information in the long completed goal versions.

Such a prediction is not made by the constructionist view. This model argues that information availability is mediated more by how it fits within the causal structure a reader creates. Because causal inference making dur-

ing narrative comprehension is guided by an effortful and strategic process, information can be kept available depending on its status in the causal structure and would be less affected by more automatic aspects of memory. Because the information in the short and long completed goal information conditions plays an identical role in the comprehension of the story, it should be equally available in these two versions.

Method

Readers. Thirty-six native English speakers participated at the University of Notre Dame for partial class credit. The data from an additional subject were replaced due to a computer malfunction.

Materials and procedure. The materials for Experiment 2 consisted of 6 long completed goal, 6 short completed goal, and 6 neutral versions of a set of 18 stories. There were no failed goal story versions presented. Also, the distractor task was changed so that people were solving three-digit addition problems (e.g., 628

+ 353 = ?) for 10 s. Otherwise, the methodology was the same as in Experiment 1.

Results

Response times. Based on our rejection criteria, one probe question response time (over 22 s long) was removed from the analyses but was not counted as an error. Response times for the short ($M = 1960$ ms; $SD = 547$) and long completed goal versions ($M = 1978$ ms; $SD = 496$) were significantly faster than for the neutral versions ($M = 2295$ ms; $SD = 702$), $t_1(35) = 3.08$, $t_2(17) = 2.12$, and $t_1(35) = 3.76$, $t_2(17) = 2.50$, respectively. There was no difference in the availability of information in the short and long completed goal versions, both $ts < 1$, even though the probe question occurred about 3 sentences further into the story in the long story versions than the short story versions. Because there was no difference in the availability of information in the long and short completed goal versions, these results are more consistent with the constructionist view than the resonance model. Furthermore, because completed goal information was always more available than neutral information, these results are inconsistent with the current-state selection strategy.

Error rates. Subjects in Experiment 2 made an average of 2.8% ($SD = 7.5$) errors on the short completed goal stories, 1.4% ($SD = 4.7$) errors on the long completed goal stories, and 13.4% ($SD = 14.3$) errors on the neutral stories. Subjects made more errors on the neutral stories than the short, $t_1(35) = 4.60$, $t_2(17) = 2.26$, and the long completed goal stories, $t_1(35) = 5.11$, $t_2(17) = 2.50$. The error rates did not significantly differ between the short and long completed goal versions, both $ts < 1$. This pattern of data is also more consistent with the constructionist view than the resonance model because there is no difference between the short and long completed story versions. These data are also inconsistent with the current-state selection strategy.

Discussion

The results of Experiment 2 are more easily reconciled with the constructionist view of

narrative comprehension than with the resonance model. In the constructionist view, the creation of causal inferences during narrative comprehension is guided less by automatic memory processes and more by the causal structure created by the reader as part of the situation model of the text. Information that is part of a causal chain, although completed, has a special representational status. This is consistent with other research that has found that causal information in general is recalled better (e.g., Black & Bower, 1980; Fletcher & Bloom, 1988; Myers & Duffy, 1990; Myers et al., 1987; Trabasso et al., 1984; Trabasso & van den Broek, 1985) and is rated as being more important (Trabasso & Sperry, 1985; Trabasso & van den Broek, 1985; van den Broek, 1988). Because the creation of a causal structure of a text in narrative comprehension is guided more by high-level processes, it can be less affected by lower level factors, such as the amount of time from when it was last mentioned.

The results from Experiment 2 are inconsistent with the resonance model, which states that narrative comprehension is influenced primarily by general memory processes that operate automatically. Specifically, information availability should decrease with the amount of time that has elapsed since it was last used. This was not observed here. The pattern of results found here deviates from that found in a recent study reported by Albrecht and Myers (1995) which used similarly sized differences between the long and short completed goal story versions. In that study, reading times for a sentence containing a global coherence inconsistency were less likely to be affected when there was a long distance between the goal information than when there was a short distance.

While we cannot precisely determine the nature of the discrepancy between our results and those of Albrecht and Myers (1995), two possibilities loom prominently. First, the nature of the task was different. In our experiment, subjects responded to memory probes, whereas in the Albrecht and Myers study,

reading time for a critical sentence was measured. Second, the role of the completed goal information in the context of the stories varied as well. In our experiment, the completed goal information was a major part of the theme of the stories up until the goal was completed. In contrast, in the Albrecht and Myers study, based on the 2 (out of 24) stories provided in their Appendix, it appears as though the completed goal was less central to the theme of the story. However, without further study, these possibilities should be regarded as speculative.

GENERAL DISCUSSION

The present study investigated the fate of completed goal information during story comprehension. Previous research had shown that failed goal information is more available than completed goal information (e.g., Suh & Trabasso, 1993), yet it was unclear whether that difference was the result of a heightened availability of the failed goal information, a decreased availability of the completed goal information, or both. There is now evidence that even when a goal is completed, this information remains more available than neutral information. So, the difference in the availability of failed and completed goal information is due to both an increased availability of failed goal information and a decreased availability of completed goal information (relative to failed goal information).

These experiments tested three theories of narrative comprehension. The constructionist view (e.g., Graesser et al., 1994) is consistent with the results from both experiments. This theory states that readers engage in a search after meaning during the comprehension process. In line with this, people make causal inferences about why various actions in a narrative are mentioned. The creation of these causal inferences is guided in part by the introduction of a goal and whether this goal had been completed. Completed goal information is more critical to the causal structure, in terms of connecting different discourse clauses, whereas the neutral information can be con-

sidered dead-end information. If one assumes that the processes that allow goal structure information to be more available than other (e.g., dead-end) information during recall are operating during comprehension as well, this view predicts that there will be a graded function of availability, with goal structure information being more available than neutral information, and failed goal information being more available than completed goal information. Furthermore, because narrative comprehension is guided primarily by high-level, strategic processes, it is less influenced by more automatic processes. This conforms to the pattern of data reported here.

The resonance model (e.g., Albrecht & Myers, 1995) is consistent with the results from Experiment 1 but has difficulty with the results of Experiment 2. This theory states that information availability is primarily guided by automatic processes that activate memory traces of information presented earlier in the text. The resonance model is consistent with the results of Experiment 1, primarily based on the principle of interconnectivity. Namely, there are fewer and fewer opportunities to interconnect information in the text as one moves from failed goal, to completed goal, to neutral text versions. However, the resonance model directly predicts a difference between the short and long completed goal versions in Experiment 2 that was not observed. The amount of time that has elapsed between the last time goal information was referenced should influence information availability, but it did not.

Finally, the current-state selection strategy (e.g., Fletcher & Bloom, 1988) states that information is either maintained in working memory or it is stored in long-term memory, depending on whether the goal has been resolved or not. As such, this position does not predict a graded amount of availability of information in the failed goal, completed goal, and neutral conditions. This model is inconsistent with the result of both experiments.

In summary, the results of the two experiments were most consistent with the construc-

tionist model concerning the fate of completed goal information. In narrative comprehension, once a goal has been stated and a causal structure linking this goal to subsequent actions, outcomes, and responses has been created, this information remains more available than neutral information. This is true even when the goal has been completed.

REFERENCES

- ALBRECHT, J. E., & MYERS, J. L. (1995). Role of context in accessing distant information during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 1459–1468.
- ALBRECHT, J. E., & O'BRIEN, E. J. (1993). Updating a mental model: Maintaining both local and global coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **19**, 1061–1070.
- BLACK, J. B., & BOWER, G. H. (1980). Story understanding as a problem solving. *Poetics*, **9**, 223–250.
- BLOOM, C. P., FLETCHER, C. R., VAN DEN BROEK, P., REITZ, L., & SHAPIRO, B. P. (1990). An on-line assessment of causal reasoning during comprehension. *Memory & Cognition*, **18**, 65–71.
- CLARK, S. E., & GRONLUND, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, **3**, 37–60.
- DOPKINS, S., KLIN, C., & MYERS, J. L. (1993). Accessibility of information about goals during the processing of narrative texts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **19**, 70–80.
- FLETCHER, C. R., BLOOM, C. P. (1988). Causal reasoning in the comprehension of simple narrative texts. *Journal of Memory and Language*, **27**, 235–244.
- FLETCHER, C. R., HUMMEL, J. E., & MARSOLEK, C. J. (1990). Causality and the allocation of attention during comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **16**, 233–240.
- GORDON, P. C., & CHAN, D. (1995). Pronouns, passives, and discourse coherence. *Journal of Memory and Language*, **34**, 216–231.
- GORDON, P. C., GROSZ, B. J., & GILLIOM, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, **17**, 311–347.
- GRAESSER, A. C., SINGER, M., & TRABASSO, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, **101**, 371–395.
- JOHNSON-LAIRD, P. N. (1983). *Mental models: Towards a cognitive science of language, inference and consciousness*. Cambridge, MA: Harvard University Press.
- JOHNSON-LAIRD, P. N. (1989). Mental models. In M. I. Posner (Ed.), *Foundations of cognitive science*. Cambridge, MA: MIT Press.
- KINTSCH, W. (1988). The role of knowledge in discourse comprehension: A constructive-integration model. *Psychological Review*, **95**, 163–182.
- KLIN, C. M. (1995). Causal inferences in reading: From immediate activation to long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 1483–1494.
- LONG, D. L., & GOLDING, J. M. (1993). Superordinate goal inferences: Are they automatically generated during comprehension? *Discourse Processes*, **16**, 55–73.
- MAGLIANO, J. P., & GRAESSER, A. C. (1991). A three-pronged method for studying inference generation in literary text. *Poetics*, **20**, 193–232.
- MYERS, J. L., & DUFFY, S. A. (1990). Causal inferences and text memory. In A. C. Graesser & G. H. Bower (Eds.), *Inferences and text comprehension*. New York: Academic Press.
- MYERS, J. L., O'BRIEN, E. J., ALBRECHT, J. E., & MASON, R. A. (1994). Maintaining global coherence during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **20**, 876–886.
- MYERS, J. L., SHINJO, M., & DUFFY, S. A. (1987). Degrees of causal relatedness and memory. *Journal of Memory and Language*, **26**, 453–465.
- RADVANSKY, G. A., SPIELER, D. H., & ZACKS, R. T. (1993). Mental model organization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **19**, 95–114.
- RADVANSKY, G. A., & ZACKS, R. T. (1991). Mental models and the fan effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **17**, 940–953.
- RADVANSKY, G. A., & ZACKS, R. T. (in press). The retrieval of situation-specific information. In M. A. Conway & S. E. Gathercole (Eds.), *Cognitive models of memory*. London: UCL Press.
- SUH, S. Y., & TRABASSO, T. (1993). Inferences during reading: Converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *Journal of Memory and Language*, **32**, 279–300.
- TRABASSO, T., SECCO, T., & VAN DEN BROEK, P. W. (1984). Causal cohesion and story coherence. In H. Mandl, N. L. Stein, & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 83–111). Hillsdale, NJ: Erlbaums.
- TRABASSO, T., & SPERRY, L. L. (1985). Causal relatedness and importance of story events. *Journal of Memory and Language*, **24**, 595–611.
- TRABASSO, T., & SUH, S. Y. (1993). Understanding text: Achieving explanatory coherence through on-line inferences and mental operations in working memory. *Discourse Processes*, **16**, 3–34.
- TRABASSO, T., & VAN DEN BROEK, P. W. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language*, **24**, 612–630.
- TRABASSO, T., VAN DEN BROEK, P. W., & SUH, S. Y.

- (1989). Logical necessity and transitivity of causal relations in stories. *Discourse Processes*, **12**, 1–25.
- VAN DEN BROEK, P. W. (1988). The effects of causal relations and hierarchical position on the importance of story statements. *Journal of Memory and Language*, **27**, 1–22.
- VAN DEN BROEK, P. W. (1995). Comprehension and memory of narrative texts: Inferences and coherence. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics*. New York: Academic Press.
- VONK, W., HUSTINX, L. G. M. M., & SIMONS, W. H. G. (1992). The use of referential expressions in structuring discourse. *Language and Cognition Processes*, **7**, 301–333.
- ZWAAN, R. A., MAGLIANO, J. P., & GRAESSER, A. C. (1995). Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 386–397.

(Received November 27, 1995)

(Revision received June 6, 1996)