

What Characters Know: Projected Knowledge and Projected Co-Presence

Richard J. Gerrig, Susan E. Brennan, and Justina O. Ohaeri

State University of New York at Stony Brook

What inferences do readers make about “who knows what” in narrative worlds? We introduce the concepts of *projected knowledge* and *projected co-presence* to describe circumstances in which readers infer that characters possess information presented, for example, only in narration. Our experiments examine one type of evidence readers use to project knowledge. In Experiment 1, readers used characters’ utterances as evidence to revise their judgments about characters’ awareness information presented in the narration. Experiment 2 established that this effect is not due to the presence of just any utterance in the story. Experiment 3 demonstrated differential projection of knowledge for characters depending on whether they were speakers or addressees of the critical utterance. Experiment 4 suggested that readers make these inferences with limited reflection. Experiment 5 demonstrated that readers’ judgment times for characters’ knowledge is affected by the properties of the projecting utterance. We conclude that individuals are skilled in evaluating textual evidence to project knowledge and co-presence. © 2001 Academic Press

Key Words: text comprehension; narration; inferences; characters’ knowledge; projected knowledge; projected co-presence.

Consider this excerpt from Amy Bloom’s story “Silver Water,” in which the narrator, Violet, describes her sister Rose’s experience of schizophrenia (1992, pp. 44–45):

She had her first psychotic break when she was fifteen. She had been coming home moody and tearful, then quietly beaming, then she stopped coming home. She would go out into the woods behind our house and not come in until my mother would go out at dusk, and step gently into the briars and saplings and pull her out, blank-faced, her pale blue pullover covered with crumbled leaves, her white jeans smeared with dirt. After three weeks of this, my mother, who is a musician and widely regarded as eccentric, said to my father, who is a psychiatrist and a kind, sad man, “She’s going off.”

Justina Ohaeri is now at Lucent Technologies, Holmdel, New Jersey. We thank Heather Bortfeld and Michael Schober for useful discussions of this project. We also thank Felicia Romano, Jennifer Scholl, and Darron Vanaria for assistance in data collection. This material is based upon work supported by the National Science Foundation under Grant IRI-9711974 and Grant IRI-9402167. The full set of stimuli may be obtained at www.psy.sunysb.edu/sbrennan.

Address correspondence and reprint requests to Richard J. Gerrig, Department of Psychology, State University of New York, Stony Brook, Stony Brook, NY 11794–2500. E-mail: rgerrig@psych1.psy.sunysb.edu.

What does “She’s going off” mean? Without knowledge of Rose’s behavior, it could simply mean “She’s leaving.” Is that what Rose’s father will take it to mean? The reader has no direct evidence that the father is aware of the behavior that disambiguates his wife’s utterance. For it to be otherwise, Bloom would need a sentence such as “my mother observed this behavior, and my father observed this behavior, and they both were aware that each knew that the other had observed the behavior.” Such a sentence would ensure that readers had reason to believe the mother and father shared common ground (Clark & Marshall, 1981). We suspect, however, that the probability of such a sentence ever appearing in a narrative work is quite low. How and why, then, are readers likely to infer that the father is in the know? We suggest that readers are making two types of inferences: *projected knowledge* and *projected co-presence*.

To define these types of inferences, we use the series of excerpts from Eudora Welty’s (1941) short stories presented in Table 1. Consider the first example, which is excerpted from the initial paragraphs of “Lily Daw and the

TABLE 1

Examples of Projected Knowledge and Projected Co-Presence from Eudora Welty

1. "Lily Daw and the Three Ladies" (p. 3)

Mrs. Watts and Mrs. Carson were both in the post office in Victory when the letter came from the Ellisville Institute for the Feeble-Minded of Mississippi. Aimee Slocum, with her hand still full of mail, ran out in front and handed it straight to Mrs. Watts, and they all three read it together. Mrs. Watts held it taut between her pink hands, and Mrs. Carson underscored each line slowly with her thimble finger. Everyone else in the post office wondered what was up now.

"What will Lily say," beamed Mrs. Carson at last, "when we tell her we're sending her to Ellisville!"

"She'll be tickled to death," said Mrs. Watts, and added in a guttural voice to a deaf lady, "Lily Daw's getting in at Ellisville!" . . .

"Do you suppose they'll look after her down there?" Mrs. Carson began to carry on a conversation with a group of Baptist ladies waiting in the post office.

2. "Why I Live At the P.O." (pp. 98-99)

Just then something perfectly horrible occurred to me.

"Mama," I says, "can that child talk?" I simply had to whisper! "Mama, I wonder if that child can be—you know—in any way? Do you realize," I says, "that she hasn't spoken one single, solitary word to a human being up to this minute? This is the way she looks," I says, and looked like this.

Well Mama and I just stood there and stared at each other. It was horrible!

3. "Lily Daw and the Three Ladies" (pp. 6, 9, and 17)

[The three ladies are searching for Lily.]

Ed Newton was stringing Redbird school tablets on the wire across the store.

"If you're after Lily, she come in here while ago and tole me she was fixin' to git married," he said.

[They return home and find Lily arranging items in a trunk.]

"Go and tell us what you're doing, Lily," said Aimee Slocum.

"Packing silly," said Lily.

"Where are you going?"

"Going to get married, and I bet you wish you was me now," said Lily.

[After Lily has boarded the train for Ellisville, a stranger approaches Aimee Slocum.]

"Could you tell me, madam" he [a stranger] said, "where a little lady lives in this burg name of Miss Lily Daw?" . . .

"What do you want to know for?" Aimee asked before she knew it. . .

"We was only going to get married, that's all," said the man.

Aimee Slocum started to scream in front of all those people.

4. "Petrified Man" (p. 42)

"Is Mr. Pike a good dresser?" asked Mrs. Fletcher sceptically.

"Oh, well, yeah," said Leota, "but he's twelve or fourteen years older'n her [Mrs. Pike]. She ast Lady Evangeline about him."

"Who's Lady Evangeline?" asked Mrs. Fletcher.

"Weal, it's this mind reader they got in the freak show," said Leota. "Was real good. Lady Evangeline is her name, and if I had another dollar I wouldn't do a thing but have my other palm read."

Three Ladies." Early on in this excerpt, "Everyone else in the post office wondered what was up now." By the end of the excerpt, it seems clear from Mrs. Carson's utterance "Do you suppose they'll look after her down there?" that the letter's contents have been widely disseminated. That is the type of inference we call *projected knowledge*: readers use evidence from the text to project their own knowledge to the characters. Suppose, further, that readers infer that all the women at the post office are

mutually aware that all the other members have also been apprised of the news in the letter. We count this as an instance of *projected co-presence*: Readers infer that two or more characters have mutual knowledge of the information.

Instances of projected knowledge and projected co-presence like these are quite ubiquitous in narrative experiences. Even so, they do not appear in standard taxonomies of inferences (e.g., Graesser, Singer, & Trabasso, 1994; Magliano, Bagget, & Graesser, 1996). The purpose

of this article is to introduce these inferences and to explore ways in which readers use one type of evidence, characters' utterances, to project knowledge to characters.

We emphasize that these are distinct types of inferences. When readers project knowledge, they infer only that the character knows what they know. Readers cannot project co-presence between themselves and characters: *Avant-garde* exercises aside, characters do not know that readers exist. When, however, readers project knowledge to more than one character they may also project co-presence (for that knowledge) among those characters. It is also quite frequently the case that readers project co-presence among characters for knowledge that they themselves do not know. We see this in the first passage in Table 1. The group of women, but not the readers, know why it is that Lily Daw might be "tickled to death" to be sent to the Ellisville Institute for the Feeble-Minded of Mississippi. The second excerpt from Table 1 presents an even more dramatic instance of projected co-presence in the absence of reader knowledge. When readers get to the phrase "and I looked like this," it might seem, at first, to be a mistake. In fact, Welty makes no attempt to elucidate "like this." Readers are left only to infer that Mama and the narrator, Sister, have co-presence for Sister's facial expression.

We can also find circumstances in which two or more characters have the same knowledge, but an inference of co-presence is not warranted. Consider the trio of passages in Table 1's third example. In their search for Lily Daw—to share the good news about Ellisville—the three ladies learn that she has been announcing that she is getting married. When the three ladies catch up with Lily Daw, she makes the same announcement to them—she does not know that they know (and readers must track the lack of co-presence to experience the moment appropriately). The three ladies believe that Lily Daw has been the victim of a traveling xylophone player who has promised marriage in exchange for Lily's favors. However, at the end of the story, the xylophone player appears on the scene to collect his bride. Aimee Slocum and the xylophonist both possess the knowledge

that Lily Daw believed that she was going to be married, but it is essential for readers to understand exactly how deeply these two characters do not have co-presence. (In the nick of time, Lily Daw is retrieved from the Ellisville-bound train and delivered to her delighted xylophonist.) A proper experience of this story requires readers to attend quite carefully to the situations in which knowledge and/or co-presence may be projected.

Table 1 provides a final example that confirms the inferential nature of projected knowledge and projected co-presence. From the way in which the character Leota refers casually to "Lady Evangeline," the utterance "She ast Lady Evangeline about him" should prompt readers to project co-presence for Evangeline between Leota and Mrs. Fletcher (making this a further example of circumstances in which characters have co-presence of something readers do not know). However, Mrs. Fletcher's response invalidates the inference of co-presence. In designing her utterance, Leota appears to have misjudged the prominence of Lady Evangeline in the community. Accordingly, readers must cancel the inference.

We intend this series of examples to illustrate the ubiquity of both projected knowledge and projected co-presence. However, for our initial empirical analysis of these phenomena, we have elected to focus on one type of evidence readers use to project knowledge with some precision to particular characters. Specifically, we analyze circumstances in which information is provided to readers, but not explicitly to characters, in the narration. Consider this brief story:

Harry and Eric were sitting in their living room, reading the newspaper. Outside the door, a dog was barking.

Harry said, "It's your turn."

Do Harry and Eric hear the dog barking? Readers may be willing to project knowledge before Harry's utterance. That utterance, however, appears to provide definitive evidence: Readers would have to infer that both Harry and Eric know about the barking to make sense of the utterance. We suggest that it is very generally the case that characters' utterances cue readers

when to infer that characters are in possession of information in the narration. Our experiments pursue the prediction that readers can and do use characters' utterances to judge when these inferences are warranted.

We devised a series of five experiments to test the prediction that readers often use the content of utterances to infer what speakers and addressees are likely to know. In these experiments we tested that hypothesis by writing stories that presented information to the readers in the narration. Consider these paragraphs:

John and Steve were walking together to their morning class. They both hoped that the professor's lecture would be interesting.

The professor got to class a few minutes late. He had been up all night fighting with his girlfriend.

Do John and Steve know about the professor's fight with his girlfriend? It seems quite unlikely. Now consider this utterance:

After the lecture, John said to Steve, "When is it our turn to whine in class?"

Clearly an inference would be required for readers to conclude that John and/or Steve knows about the professor's fight. However, given appropriate standards for discourse coherence and utterance relevance (Grice, 1975; Sperber & Wilson, 1986), the inference seems warranted. We suggest that readers project knowledge through a mechanism such as *abduction*: They make the inferences that provide the best explanation for what they read or hear in a discourse (Hobbs, Stickel, Appelt, & Martin, 1990; Pierce, 1955). As Pierce proposed, "the surprising fact, C [e.g., John's utterance], is observed; but if A [e.g., John and Steve know about the Professor's fight] were true, C would be a matter of course, hence there is reason to suspect that A is true" (Pierce, 1955, as cited by Hobbs et al., 1990, p. 44). We assume, as well, that readers expect authors to follow normal practices of cooperation in depicting conversations (Gerrig, 1993).

Note that we opted against an even simpler demonstration. Consider this utterance:

After the lecture, John said to Steve, "Why does he think we care about his girlfriend?"

For readers to conclude in this case that the professor mentioned his fight requires less of an inference. The utterance still projects knowledge, but in a way that tells us less about readers' standards for evidence. We designed our experiments to demonstrate that utterances that do not make surface reference to knowledge from the narration can nonetheless project that knowledge.

Experiment 1 demonstrates that the evidence provided by characters' utterances will, in fact, lead readers to make inferences that characters know critical information from the narration. Experiment 2 provides an additional control condition for and a replication of this experiment. Experiment 3 demonstrates that utterances can differentially project knowledge to speakers and addressees. Finally, Experiments 4 and 5 provide converging evidence from readers' judgments in a response-time paradigm. These speeded responses suggest that readers are able to make efficient judgments about circumstances in which these types of inferences are and are not warranted.

EXPERIMENT 1: UTTERANCES PROJECT KNOWLEDGE

In this initial experiment, participants read one of two versions of a series of brief stories. One version ended with the introduction of critical information in the narration. The second version added an utterance that we intended to project knowledge for that information. After reading each story, participants were asked to judge how likely it was that one or the other character knew the narration information. We predicted that participants' ratings would be considerably higher when the stories included the projecting utterances.

Method

Participants. Twenty-four undergraduates at the State University of New York at Stony Brook participated in the experiment. All were native speakers of English.

Materials. We wrote 20 brief stories (for examples, see Table 2). Each story began with a sentence or two of narration to set the scene. In the next one or two sentences, the narration

TABLE 2

Sample Experimental Stories

Narration

Frank and Larry had invested their life savings in a company called, "Pioneer Aviation." The value of their several thousand shares of stock had grown 1000% in three months.

Early one Friday morning in April, the president of Pioneer Aviation emptied out all the company's bank accounts and ran off to Brazil.

Projecting utterance

Frank said to Larry, "I'm glad you convinced me to play the stock market." [Experiments 1, 2, and 4]

Nonprojecting utterance

Frank said to Larry, "I just bought some software to help us track our investments." [Experiment 2]

Projects-speaker

Frank said to Larry, "You'd better be sitting down before we talk." [Experiments 3 and 5]

Projects-both

Frank said to Larry, "Aren't we just about the most unlucky guys on earth?" [Experiments 3 and 5]

Knowledge question

Frank [Larry] knows that the president has stolen the money.

Narration

Diane and Jack were celebrating their 50th anniversary. They had decided to throw a big party for themselves. Their three children and all of their grandchildren were flying into town for the weekend.

The couple's children had decided to pool their resources to get their parents one big gift. They had purchased a luxurious condominium for their parents in Florida.

Projecting utterance

Diane said to Jack, "We're lucky to have children who spoil us in our old age." [Experiments 1, 2, and 4]

Nonprojecting utterance

Diane said to Jack, "I can't wait to see how our grandchildren have grown." [Experiment 2]

Projects-speaker

Diane said to Jack, "I can't wait to see your face when the children announce our gift." [Experiments 3 and 5]

Projects-both

Diane said to Jack, "We're lucky to have children who spoil us in our old age." [Experiments 3 and 5]

Knowledge question

Diane (Jack) knows about the condo in Florida.

provided information that would, on the surface, not be available to the characters. The *narration only* stories ended with this potentially reader-only information. The *narration-and-projecting utterance* version of each story continued with an utterance that was intended to project knowledge. We planned for these projecting utterances not to make direct reference to the critical information. Accordingly, we asked two raters, students who were blind to our hypothesis, to read each story without its critical information. Consider this story with that information removed:

John and Steve were walking together to their morning class. They both hoped that the professor's lecture would be interesting.

After the lecture, John said to Steve, "When is it our turn to whine in class?"

The two raters indicated to us which stories they thought were incoherent. We amended problematic stories and then submitted them again to the student raters. The final stories were judged to be coherent.

We also wrote 20 filler stories to break the intended correlation between story structure and knowledge in the critical stories. That is, our prediction for the experimental items was that the versions of the stories including utterances would yield higher knowledge ratings than those without. Accordingly, we wrote *narration only* filler stories that we intended to yield high knowledge ratings and *narration-and-utterance* filler stories that we intended to yield low ratings.

For each story, participants were asked to

indicate how much they agreed with a statement about a character's knowledge (e.g., *Steve* knows the professor had been fighting with his girlfriend) on a scale ranging from 1 (*Definitely No*) to 9 (*Definitely Yes*). Participants made judgments about the knowledge of the speakers and addressees for half of the stories each. For the stories that terminated without an utterance, by "speaker" and "addressee" we mean the character who would have taken on that role in the full story.

Design and procedure. There were four versions of each item: Two stories (narration only vs narration-and-projecting utterance) were crossed with two characters (speaker and addressee) for each story. The experimental stories were randomly interspersed with the filler stories. One version of each story was distributed to one of four questionnaires using a Latin square design. Overall, readers saw an equal number of each type of item. Four stories appeared on each questionnaire page, and the pages were presented to each participant in a different random order.

The instructions began with a paragraph to motivate the experiment:

When you read a story, it is often important for you to keep track of both what you, the reader, know and what the characters know. For example, when you are reading a suspense story, you might know something that the characters do not—There's danger behind that door!—and those different states of knowledge contribute to your enjoyment of the story.

Participants were then instructed to read each story and use the 1-to-9 scale to indicate how much they agreed with each statement about the characters' knowledge.

Results and Discussion

The results support the prediction that utterances can project knowledge (see Table 3). We performed analyses with both participants ($F1$) and items ($F2$) as random variables. Knowledge ratings were reliably higher when a projecting utterance was added to the narration [$F1(1, 20) = 72.91$, $MS_e = 1.27$, $p < .001$; $F2(1, 16) = 44.20$, $MS_e = 1.75$, $p < .001$]. There was no main effect for character. That is, ratings did not differ overall for the speaker versus the

TABLE 3
Results of Experiment 1: Participants' Ratings of Characters' Knowledge

	Text version		Mean
	Narration only	Narration-and-projecting utterance	
Character			
Speaker	2.70	5.13	3.92
Addressee	3.13	4.63	3.88
Mean	2.92	4.88	

addressee (both F 's < 1). However the utterances, apparently, projected knowledge of the critical information less successfully for the addressees than for the speakers: The increase in ratings with the addition of the projecting utterance was 2.43 for the speakers but only 1.50 for the addressees. This interaction between story version and character was reliable in the item analysis and marginally so for participants [$F1(1, 20) = 3.76$, $MS_e = 1.39$, $p = .067$; $F2(1, 16) = 4.66$, $MS_e = 0.94$, $p < .05$]. This interaction suggests that readers were taking appropriate care with their judgments: It seems quite reasonable for readers to have greater confidence that speakers are in possession of the knowledge that is projected by their utterances. This is particularly true because, within the bounds of our stories, the addressees never have the opportunity to acknowledge or reject the utterance (e.g., an addressee might laugh knowingly or respond "Huh?") (cf. Clark & Brennan, 1991).¹ We also see in this result the distinction between projected knowledge versus projected co-presence. Our readers were more confident that they shared knowledge with the speakers than, apparently, that the speakers and addressees both were in the know.

The purpose of this experiment was to demonstrate that characters' utterances can project knowledge for information that is presented in the narration. Although the stories provided no direct evidence that the critical narration infor-

¹ We thank an anonymous reviewer for this substantial insight.

mation was available to the characters, readers were willing to infer possession of the knowledge when the utterance suggested this was the case. However, to solidify this point, we needed to demonstrate that it is not the mere presence of an utterance that leads readers' to provide more confident knowledge attributions. This was the goal of the next experiment.

EXPERIMENT 2: NOT ALL UTTERANCES PROJECT KNOWLEDGE

In Experiment 1 we demonstrated that the addition of appropriate utterances to brief stories radically changed readers' beliefs about characters' knowledge. Experiment 2 replicated Experiment 1 and provided an important control: We ruled out the possibility that the effect arises from the mere presence of an utterance. Some utterances should not project knowledge:

John and Steve were walking together to their morning class. They both hoped that the professor's lecture would be interesting.

The professor got to class a few minutes late. He had been up all night fighting with his girlfriend.

After the lecture, John said to Steve, "Let's study together for the exam."

The utterance "Let's study together for the exam" should not project knowledge of the professor's fight from the reader to either John or Steve. It is sufficiently relevant to the general topic of a classroom situation not to demand any further inferences.

In Experiment 2, participants read stories concluding with utterances that were intended to project or not project knowledge. Once again, readers assessed the knowledge of either the speaker or the addressee.

Method

Participants. Twenty-four undergraduates at the State University of New York at Stony Brook participated in the experiment. All were native speakers of English and none had participated in our previous experiment.

Materials. We wrote *nonprojecting utterances* for the 20 stories from Experiment 1. As in Experiment 1, we had two student raters read the stories without the critical narration information to ensure that interpretation of the non-

TABLE 4
Results of Experiment 2: Participants' Ratings of Characters' Knowledge

	Text version		Mean
	Projecting utterance	Nonprojecting utterance	
Character			
Speaker	5.63	2.98	4.31
Addressee	5.09	3.14	4.12
Mean	5.36	3.06	

projecting utterances did not rely on that narration. All stories appeared in their narration-and-utterance versions. We adjusted the filler items to equalize the number of (expected) high and low knowledge ratings. Readers provided knowledge ratings using the 1-to-9 scale from Experiment 1.

Design and procedure. There were four versions of each item: Two utterances (projecting vs nonprojecting) crossed with two characters (speaker and addressee) for each story. The experimental stories were randomly interspersed with the filler stories. One version of each story was distributed to one of four questionnaires using a Latin square design. Four stories appeared on each questionnaire page which were presented to each participant in a different random order. The instructions were identical to those for Experiment 1.

Results and Discussion

We predicted that knowledge ratings would be higher for projecting utterances than for nonprojecting utterances. As shown in Table 4, readers' responses bore out that prediction [$F(1, 20) = 120.15, MS_e = 1.06, p < .001$; $F(1, 16) = 45.53, MS_e = 2.32, p < .001$]. This result allows us to rule out the possibility that the mere presence of an utterance produced the ratings differences in Experiment 1. As in Experiment 1, there was no main effect for character (i.e., speaker vs addressee). Also, there was a tendency for the projecting utterance to be more successful for speakers than addressees—the increase in ratings was 2.65 for speakers but

only 1.95 for addressees—but this difference was not reliable [$F(1, 20) = 3.84$, $MS_e = 0.76$, $p = .064$; $F(1, 16) = 2.40$, $MS_e = 1.02$, $p > .10$]. Note that the mean ratings in Experiment 1 for the narration only versions (2.92) are roughly the same as those in Experiment 2 for the nonprojecting versions (3.06). This near equivalence supports the conclusion that both types of stories left readers equally reluctant to believe that the characters were in possession of information provided in the narration. In both experiments, the projecting utterances considerably boosted readers' knowledge attributions.

In both Experiments 1 and 2, we have seen traces of evidence that our utterances projected knowledge more successfully for speakers than for addressees. In Experiment 3, we turn directly to that effect.

EXPERIMENT 3: SPEAKERS AND ADDRESSEES

The purpose of Experiment 3 was to clarify the effects of our projecting utterances on readers' attributions of knowledge to speakers and addressees. As we suggested earlier, it is entirely appropriate, in ordinary circumstances, for readers to have greater confidence that speakers are in possession of knowledge projected by their utterances. Even given expectations that speakers aspire to be cooperative—i.e., they aspire to produce utterances for which addressees possess appropriate knowledge—it should still be the case that readers express a small amount of wariness toward speakers' estimates of addressees' knowledge (particularly because, as we also noted, the addressees in our stories never have the opportunity to respond to the critical utterances).

These considerations suggest that different utterances should increase or decrease readers' assessments that both the speaker and addressee possess information presented in the narration. Consider, yet again, the professor and his girlfriend:

John and Steve were walking together to their morning class. They both hoped that the professor's lecture would be interesting.

The professor got to class a few minutes late. He had been up all night fighting with his girlfriend.

After the lecture, John said to Steve, "You picked a bad day to sleep through class!"

In this case, we intended the final utterance to project knowledge only to John, the speaker. In Experiments 1 and 2 we had intended the utterances to project knowledge to both the speaker and addressee. In this experiment, we wrote utterances that we intended to project knowledge for just the speaker or for both characters.

Experiment 3 also gives us an opportunity to demonstrate that readers are taking reasonable care in the way they project that information mentioned in the narration is known to characters (cf. Gerrig, Ohaeri, & Brennan, 2000; Gerrig, Brennan, & Ohaeri, 2000; Keysar, 1994, 2000). If we succeed in increasing the difference between readers' attributions of knowledge to speakers versus addressees, that will suggest that readers are able to maintain distinctions between what they know and what some or all characters know.

Method

Participants. Twenty-eight undergraduates at the State University of New York at Stony Brook participated in the experiment. All were native speakers of English and none participated in our earlier studies.

Materials. We wrote utterances that we intended to project knowledge for the speaker or for both characters. In 9 (of 20) cases (as in the professor example) we used the original utterance from Experiment 1 as the utterance to project both (i.e., we retained those utterances that had, in fact, produced nearly equal ratings for speakers and addressees in Experiment 1). To equalize the number of (expected) high and low knowledge ratings, we adjusted the filler items so that 10 were meant to generate high ratings and 10 low ratings. Readers provided ratings using the 1-to-9 scale from the earlier experiments.

Design and procedure. There were four versions of each item: Two utterances (*projects-speaker vs projects-both*) crossed with two characters (speaker and addressee) for each story. The experimental stories were randomly interspersed with the filler stories. One version of each story was distributed to one of four

TABLE 5

Results of Experiment 3: Participants' Ratings
of Characters' Knowledge

	Text version		Mean
	Projects-speaker	Projects-both	
Character			
Speaker	6.23	6.34	6.29
Addressee	4.60	5.76	5.18
Mean	5.41	6.05	

questionnaires using a Latin square design. Four stories appeared on each questionnaire page which were presented to each participant in a different random order. The instructions were identical to those for the earlier experiments.

Results and Discussion

The intention of Experiment 3 was to demonstrate that readers can differentiate between utterances that project knowledge to a speaker versus those that project knowledge to both a speaker and addressee. As shown in Table 5, readers are able to make these distinctions. The speaker–addressee difference between knowledge ratings for the projects-speaker utterances was 1.63 points; the difference for the projects-both utterances was 0.58 [$F(1, 24) = 9.94$, $MS_e = 0.77$, $p < .005$; $F(1, 16) = 8.18$, $MS_e = 0.66$, $p = .011$]. In this experiment, readers gave higher ratings overall when the utterance projected both characters [$F(1, 24) = 7.67$, $MS_e = 1.48$, $p = .011$; $F(1, 16) = 6.30$, $MS_e = 1.28$, $p < .05$]; they also made stronger knowledge attributions overall to the speakers than the addressees [$F(1, 24) = 15.89$, $MS_e = 2.16$, $p = .001$; $F(1, 16) = 43.89$, $MS_e = 0.56$, $p < .001$].

This latter result makes plain that we were not able to eliminate the speaker–addressee difference in the projects-both case. Although we were able to make the difference larger than in Experiments 1 and 2 when we intended our utterances to project knowledge for the speaker alone, we did not succeed at making the difference smaller when we intended our utterance to project knowledge for both characters (compare

the 0.58 difference in this experiment to differences of 0.50 and 0.54 in Experiments 1 and 2 respectively). As we suggested earlier, this might be a consequence of the lack of addressee acceptance of the projecting utterances (e.g., “You’ve got that right!”). In any case, Experiment 3 demonstrated that the projection of knowledge to speakers and addressees is, to a certain extent, under control of a character’s exact utterance. As such, the experiment also demonstrates that readers are able to appreciate differences between their own knowledge and the knowledge of different characters.

Experiments 1 through 3 demonstrated that readers use characters’ utterances as evidence to project knowledge. However, each of these experiments allowed participants to review the brief texts for as long as they wished. We wanted to ensure that readers’ judgments follow a similar pattern when they make those judgments with fewer opportunities for deliberation. In our final pair of experiments, we asked participants to read these same brief stories on a computer screen and make their judgments as swiftly as possible. Experiment 4 provides a replication of the impact of the projecting utterance (Experiment 1). Experiment 5 provides a replication of the differential projection to speakers and addressees (Experiment 3).

EXPERIMENT 4: SPEEDED JUDGMENTS FOR PROJECTING UTTERANCES

In Experiment 4, participants were asked to make speeded yes/no judgments about characters’ knowledge. We predicted that participants’ speeded judgments would parallel those from Experiment 1.

Method

Participants. Twenty Stony Brook undergraduates participated in this study for class credit. All participants were native speakers of English. None had participated in our earlier studies.

Apparatus. The experiment was run on two IBM-compatible 486 personal computers that recorded agreement responses and response times. Participants were seated in front of a MicroScan color monitor with their hands rest-

ing on the keyboard. They used buttons on the keyboard to make appropriate responses. The stories were displayed in the center of the screen in standard upper- and lowercase type.

Materials. We used the 20 stories from Experiment 1. Each story had a *narration only* version, in which the story ended with the critical narration, and a *narration-and-utterance* version, in which the story continued with a character's utterance. The knowledge statements from that experiment (e.g., Steve knows the professor had been fighting with his girlfriend) served as the targets for yes/no verification. Finally, we used the filler items from Experiment 1 to break the correlation between story length and the expected yes/no response.

Design and procedure. As in earlier experiments, there were four versions of each item: Two stories (narration only vs narration-and-projecting utterance) were crossed with two characters (speaker and addressee) for each story. One version of each story was distributed to one of four story lists using a Latin square design.

Participants began with five practice stories and knowledge sentences to become acquainted with the stimuli format and keyboard controls. Each story began with the words "Press NEXT for the next story" appearing on the monitor. Participants pressed the "Y" key, labeled NEXT, to proceed to the next story. The whole story appeared on the screen. Participants were instructed to press the spacebar after they had fully understood the story. When they pressed the spacebar, a beep sounded from the computer, and a knowledge sentence appeared. Participants then were required to press either a YES (i.e., "I agree"; "?") or NO (i.e., "I disagree"; "Z") key in response to the sentence. Participants were informed that we were recording their response times, but that they should try to be both swift and accurate. Each participant saw the 20 experimental items and 20 filler items intermixed in a different random order.

Results and Discussion

We eliminated trials on which readers' response times fell more than 3 standard devia-

TABLE 6

Results of Experiment 4: Participants' Judgments of Characters' Knowledge and Response Times to Make Those Judgments

Measure	Narrative only		Narrative-and-projecting utterance	
	Speaker	Addressee	Speaker	Addressee
Percentage "yes"				
Responses	21.4	22.4	67.7	51.5
Decision times	2.66	2.75	2.56	2.64
"Yes"				
Responses	3.41	2.91	2.63	2.71
"No"				
Responses	2.45	2.71	2.41	2.56

Note. Decision times are participants' latencies to provide either a "yes" or "no" response. Decision times are reported in seconds.

tions above the mean. This resulted in a loss of 2.0% of the data.

The results support the prediction that participants' speeded judgments would parallel the patterns of responses from the earlier experiments. As shown in Table 6, the percentage of "yes" responses (i.e., the percentage of occasions on which readers agreed that a character possessed the critical narration information) in each category is consistent with participants' ratings from Experiment 1 (cf. Table 3). In fact, the correlation for the items between the two sets of data (i.e., percentage "yes" responses and mean ratings for each version of each story) was reliable ($r = .62, p < .001$). As in the earlier experiments, participants agreed much more often that the characters were in possession of knowledge when a projecting utterance was added to the narration [$F1(1, 16) = 62.63, MS_e = 0.05, p < .001; F2(1, 16) = 59.35, MS_e = 0.04, p < .001$]. There was also some evidence supporting a main effect for character: Participants were more likely to believe that speakers possessed knowledge than addressees [$F1(1, 16) = 3.05, MS_e = 0.03, p = .10; F2(1, 16) = 6.06, MS_e = 0.04, p < .05$]. The interaction between these two factors, however, was not reliable [$F1(1, 16) = 1.52, MS_e = .06, p > .10; F2(1, 16) = 3.71, MS_e = 0.04, p < .10$].

Because of the different mixtures of "yes" and "no" responses for each story version, we

did not have any particular predictions for decision times. As shown in Table 6, no strong patterns emerge in these data. Participants took roughly the same amount of time to respond either “yes” or “no” for each judgment (all F 's < 1.70). Table 6 also provides the response times separately for “yes” and “no” responses. Because the number of observations in each of these means is quite variable (as a function of the percentage of “yes” responses), we didn't carry out any statistical analyses. It is worth noting, however, that the longest response times were for those rare occasions on which a participant attributed knowledge to the characters in the absence of the projecting utterance.

Overall, these data converge nicely with the data from Experiment 1. Apparently, readers' intuitions about “who knows what” were highly similar in circumstances in which they had the opportunity to reflect and those in which they made speeded judgments. We conclude that people are able to project knowledge with relative ease.

EXPERIMENT 5: SPEEDED JUDGMENTS FOR SPEAKERS AND ADDRESSEES

Experiment 5 used the response time paradigm to examine the contrast between utterances we intended to project both speaker and addressee versus those we intended to project the speaker alone. Our goal, as with Experiment 4, was to demonstrate a pattern of responses that echoed readers' nonspeeded ratings. However, for Experiment 5 we also are able to make predictions about patterns of response times. Specifically, we suggest that readers' judgments should be particularly lengthy when an utterance projects the speaker alone, but the readers are asked about the *addressee's* knowledge: Rather than being able to rely on inferences already made, readers will have to engage an effortful and time-consuming decision process. In contrast, when an utterance has already projected knowledge to the addressee, readers should be able to respond relatively quickly that the addressee is in possession of the knowledge. We expect, furthermore, that readers' decision times will be roughly equivalent in instances in

which we ask about the speakers' knowledge. Experiment 5 allows us to test these predictions.

Method

Participants. Twenty-four Stony Brook undergraduates participated in this study for class credit. All participants were native speakers of English. None had participated in our earlier studies.

Apparatus. The experiment was run on one Pentium II and one IBM-compatible 486 personal computer that recorded agreement responses and response times. Participants were seated in front of a MicroScan color monitor with their hands resting on the keyboard. They used buttons on the keyboard to make appropriate responses. The stories were displayed in the center of the screen in standard upper- and lowercase type.

Materials. We used the 20 stories from Experiment 3. Each story had a *projects-speaker* and *projects-both* version. As in Experiment 4, the knowledge statements from earlier experiments (e.g., Steve knows the professor had been fighting with his girlfriend) served as the targets for yes/no verification. We used the same filler items as in Experiment 3.

Design and procedure. As in earlier experiments, there were four versions of each item: Two stories (*projects-speaker* vs *projects-both*) were crossed with two characters (speaker and addressee) for each story. One version of each story was distributed to one of four story lists using a Latin square design.

Participants began with five practice stories and knowledge sentences to become acquainted with the stimuli format and keyboard controls. Each story began with the words “Press NEXT for the next story” appearing on the monitor. Participants pressed the “Y” key, labeled NEXT, to proceed to the next story. The whole story appeared on the screen. Participants were instructed to press the spacebar after they had fully understood the story. When they pressed the spacebar, a beep sounded from the computer, and a knowledge sentence appeared. Participants then were required to press either a YES (i.e., “I agree”; “/?”) or NO (i.e., “I disagree”; “Z”) key in response to the sentence.

TABLE 7

Results of Experiment 5: Participants' Judgments of Characters' Knowledge and Response Times to Make Those Judgments

Measure	Projects-speaker		Projects-both	
	Speaker	Addressee	Speaker	Addressee
Percentage "yes" Responses	70.4	39.0	70.6	59.2
Decision times "Yes"	2.79	3.57	2.92	3.15
Responses "No"	2.84	3.98	2.88	3.26
Responses	2.69	3.31	3.01	2.99

Note. Decision times are participants' latencies to provide either a "yes" or "no" response. Decision times are reported in seconds.

Participants were informed that we were recording their response times, but that they should try to be both swift and accurate. Each participant saw the 20 experimental items and 20 filler items intermixed in a different random order.

Results and Discussion

We eliminated trials on which readers' response times fell more than 3 standard deviations above the mean. This resulted in a loss of 1.0% of the data.

The results support the prediction that participants' speeded judgments would parallel the patterns of responses from the earlier experiments. As shown in Table 7, the percentage of "yes" responses (i.e., the percentage of occasions on which readers agreed that a character possessed the critical narration information) in each category is consistent with participants' ratings from Experiment 3 (cf. Table 5). The correlation between the two sets of data was reliable ($r = .60, p < .001$). In analogy to Experiment 3, the speaker-addressee difference between "yes" responses was larger for the projects-speaker utterances (31.4%) than for the projects-both utterances (11.4%) [$F(1, 20) = 8.91, MS_e = 0.027, p < .007; F(1, 16) = 10.67, MS_e = 0.019, p < .005$]. Participants also made stronger knowledge attributions overall to the speakers than the addressees [$F(1, 20) = 19.91, MS_e = 0.053, p < .001; F(1, 16) = 27.25,$

$MS_e = 0.033, p < .001$]; they gave somewhat higher ratings overall when the utterance projected both characters [$F(1, 20) = 5.20, MS_e = 0.052, p = .034; F(1, 16) = 4.21, MS_e = 0.50, p = .057$].

Critically, in Experiment 5 we were able to make definitive predictions about the time course of readers' judgments: The data in Table 7 confirm that readers are slower to make their judgments about addressee's knowledge in circumstances in which the utterance was intended to project only the speaker. The difference between responses to the speaker and addressee target sentences was larger for the projects-speaker condition (776 ms) than for the projects-both condition (230 ms) [$F(1, 20) = 4.28, MS_e = 390397, p = .052; F(1, 16) = 6.08, MS_e = 248747, p < .025$]. Participants also made quicker judgments overall about speakers' knowledge [$F(1, 20) = 21.61, MS_e = 268879, p < .001; F(1, 16) = 15.82, MS_e = 314574, p < .001$]. There were no reliable differences between the judgment times for the projects-speaker versus projects-both conditions (both F 's < 1.6).

These data support the conclusions we drew in Experiment 3: Even with relatively less time for reflection, readers are responsive to the exact nature of the utterance when projecting knowledge to characters. We found evidence for this reader responsiveness both in the pattern of "yes" judgments and in judgment times. The results of this experiment also reinforce our conclusion from Experiment 3 that readers are making reasonably subtle distinctions about "who knows what": They are not projecting their own knowledge to characters in a blanket fashion.

GENERAL DISCUSSION

We intended these experiments to demonstrate one type of evidence readers use to project knowledge. In Experiment 1, stories presented critical information in the narration. When the stories were truncated after that narration, readers were reluctant to attribute knowledge to the characters. When, however, one of the characters produced an utterance that could be interpreted as relevant to the critical infor-

mation, readers' knowledge ratings were reliably higher. Experiment 2 provided a replication and ruled out the possibility that this effect was merely a consequence of the mere presence of any additional utterance. Experiment 3 followed up the hint from the previous pair of experiments that the utterances projected knowledge more solidly to the speakers. We demonstrated that utterances could differentially project knowledge to just the speaking character or to both characters. Experiments 4 and 5 provided evidence that readers' speeded judgments paralleled the knowledge judgments. Experiment 5, in addition, demonstrated consistency between judgments and judgment times.

We take this series of experiments to provide strong evidence that readers project knowledge in an appropriate and subtle fashion. In these experiments, the characters' use of utterances suggested to readers that the characters were aware of information in the narration without any *explicit* evidence to that effect. We also believe that the utterances often gave rise to projected co-presence. We must acknowledge, however, that our methodology does not allow us to conclude that readers came to believe that information was mutually known to the two characters. We expect, even so, that readers do project co-presence. We can bolster that intuition a bit more by turning back to a type of text that we thought was too trivial to warrant empirical attention. Recall that the stories we used in these experiments specifically avoided circumstances in which the surface content of the projecting utterances alluded to the critical information. Recall the version of the professor story in which the projecting utterance did have overlapping content:

John and Steve were walking together to their morning class. They both hoped that the professor's lecture would be interesting.

The professor got to class a few minutes late. He had been up all night fighting with his girlfriend.

After the lecture, John said to Steve, "Why does he think we care about his girlfriend?"

In this case, it seems enormously likely that readers would both project knowledge from themselves to each character and also co-presence between John and Steve. Furthermore, it

would seem quite easy to confirm that readers will agree that there are textual circumstances in which characters possess mutual knowledge (i.e., an inference of co-presence is warranted) in the absence of reader knowledge (cf. Table 1, example 2). Future work will pursue these intuitions empirically.

Our experiments, in any case, have illustrated one set of circumstances in which readers are likely to project knowledge, at least. We have not, however, offered a more general account of all the textual circumstances in which readers are likely to make these inferences. We believe there is one category of textual events that is likely to prompt readers automatically to project knowledge and, if there is more than one character in the scene, co-presence. Recall our brief story:

Harry and Eric were sitting in their living room, reading the newspaper. Outside the door, a dog was barking.

Harry said, "It's your turn."

It seems quite likely, on our view, that readers will almost always infer it to be the case that characters know about perceptually salient events. In this case, we believe that readers would project knowledge and co-presence for the dog's bark. (As we noted earlier, Harry's utterance likely confirms the inference.) We suspect that readers automatically infer that the content of perceptual events (e.g., words but not necessarily meanings) are available to all characters who are physically present in the environment.

For the types of projecting utterances that we have explored in our current experiments, we believe that it will not be possible to make broad claims about what is "almost always" the case. To write our projecting utterances, we imagined ourselves in the situations and invented utterances that we believed would function appropriately. We counted on our readers to assume that the authors were cooperative and the texts were constructed in such a fashion that the characters' utterances were taken, when possible, to be relevant to the narration. Our experiments demonstrated that readers make these inferences without any direct overlap of content

between narration and utterance. We suspect, however, that in many texts there would be such textual overlap. The opening quotation from “Silver Water” presents an example of this type. The utterance “She’s going off” makes more or less explicit reference to the types of behavior that have been described in the narration. We suspect, in light of the diversity of textual circumstances, that readers will employ a range of strategic and nonstrategic processes to project knowledge and co-presence.

Note that our experiments suggest that many instances of projected knowledge and co-presence will require readers to revise their initial representations. In our stories, readers most often did not credit the characters with the critical information from the narration until after the projecting utterances. Thus, it seems likely that however readers had initially represented that critical information, they had to amend those representations. There is, unfortunately, a major difficulty with testing this prediction. We would need to devise a methodology that allows us to differentiate between what readers know and what characters know. Because moment-by-moment paradigms largely rely on probes with single words (e.g., cross-modal lexical decision, Stroop interference paradigms, and word recognition) it is not clear how to examine readers’ representations of characters’ knowledge separately from readers’ representations of their own knowledge. To study the precise time course with which inferences are generated by projecting utterances, we would have to overcome this methodological hurdle.

Our experiments emphasize how important it is to understand how language users maintain models of “who knows what.” We have shown that readers are able to make subtle distinctions between, for example, what they know and what characters (e.g., narrative speakers and addressees) know. It would be important, in addition, to document whether those distinctions are maintained over longer texts and whether they are encoded into long-term memory representations. Lea, Mason, Albrecht, Birch, and Myers (1998) provided evidence that suggests that readers do, in fact, track “who knows what” over the course of longer texts. In their experi-

ments, readers encountered stories in which either one or both characters were aware of the actions of the first character. For example, in one story either Jane, but not Gloria, or both Jane and Gloria knew that Jane was going to have dinner with her cousin. At the end of the text, the *second* character made reference to that action (e.g., Gloria asked Jane, “Did she play you old disco records?”). As documented by their reading times, participants noticed the inconsistency when the second characters made reference to events of which they were—by the evidence of the text—officially unaware. It would be instructive to use Lea et al.’s methodology to seek similar inconsistency effects in circumstances in which readers may or may not project knowledge or co-presence.²

Within the text processing literature, researchers have begun to document readers’ quite impressive ability to correctly differentiate the sources (e.g., the narrator vs a character) of information in real stories (Graesser, Bowers, Bayen, & Hu, 2000; Graesser, Swamer, & Hu, 1997). Researchers have also begun to consider how readers construct models of the narrator as a source of information (Dixon & Bortolussi, 1996). Attention has also been focused on circumstances in which readers’ inferences of “who knows what” may not be warranted (cf. Gerrig, Brennan, & Ohaeri, 2000; Gerrig, Ohaeri, & Brennan, 2000; Keysar, 1994, 2000). Our analysis of projected knowledge and projected co-presence suggests that such potentially unwarranted inferences should be interpreted within the broader context of how readers manage their models of character and narrator knowledge.

In this article, we have suggested that readers are often required to make inferences—to project knowledge and co-presence—to build a model of how knowledge is distributed to characters. We suspect, in fact, that good authors often purposefully leave readers with some inferences to make, as part of the fun of experiencing stories. Consider a final moment from “Lily Daw and the Three Ladies.” By the time

² Again, we thank an anonymous reviewer for this very interesting suggestion.

the three ladies have convinced Lily to board a train for Ellenville, the Victory Civic Band had “assembled without any orders” (Welty, 1941, p. 15) to see her off. Readers must project knowledge and co-presence to and among the members of the Victory Civic Band—and, in doing so, encode an enriched model of how quickly and efficiently knowledge might be shared in a small Southern town. To understand how textual moments like this one work so amusingly well, we must expand theories of text processing to encompass inferences about “who knows what.”

REFERENCES

- Bloom, A. (1992). Silver water. In *The best American short stories 1992* (pp. 43–51). Boston: Houghton Mifflin.
- Clark, H., & Brennan, S. (1991). Grounding in communication. In L. Resnick, J. Levine, & S. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). American Psychological Association: Washington, DC.
- Clark, H. H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. Webber, & I. Sag (Eds.), *Elements of discourse understanding* (pp. 10–63). Cambridge, UK: Cambridge Univ. Press.
- Dixon, P., & Bortolussi, M. (1996). Literary communication: Effects of reader–narrator cooperation. *Poetics*, **23**, 405–430.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford Univ. Press.
- Gerrig, R. J. (1993). *Experiencing narrative worlds*. New Haven, CT: Yale Univ. Press.
- Gerrig, R. J., Brennan, S. E., & Ohaeri, J. O. (2000). What can we conclude from speakers behaving badly? *Discourse Processes*, **29**, 173–178.
- Gerrig, R. J., Ohaeri, J. O., & Brennan, S. E. (2000). Illusory transparency revisited. *Discourse Processes*, **29**, 137–159.
- Graesser, A. C., Bowers, C. A., Bayen, U. J., & Hu, X. (2000). Who said what? Who knows what? Tracking speakers and knowledge in narrative. In W. van Peer & S. Chatman (Eds.), *Narrative perspective: Cognition and emotion*. Buffalo, NY: SUNY Press, in press.
- Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, **101**, 371–395.
- Graesser, A. C., Swamer, S. S., & Hu, X. (1997). Quantitative discourse psychology. *Discourse Processes*, **23**, 229–263.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol.3. Speech acts* (pp. 41–58). New York: Academic Press.
- Hobbs, J. R., Stickel, M., Appelt, D., & Martin, P. (1990). *Interpretation as abduction*. (Technical Note 499). Menlo, CA: SRI International.
- Keysar, B. (1994). The illusory transparency of intention: Linguistic perspective taking in text. *Cognitive Psychology*, **26**, 165–208.
- Keysar, B. (2000). The illusory transparency of intention: Does June understand what Mark means because he means it? *Discourse Processes*, **29**, 161–172.
- Lea, R. B., Mason, R. A., Albrecht, J. E., Birch, S. L., & Myers, J. L. (1998). Who knows what about whom: What role does common ground play in accessing distant information? *Journal of Memory and Language*, **39**, 70–84.
- Magliano, J. P., Bagget, W. B., & Graesser, A. C. (1996). A taxonomy of inference categories that may be generated during the comprehension of literary texts. In R. J. Kreuz & M. S. MacNealy, *Empirical approaches to literature and aesthetics* (pp. 201–220). Norwood, NJ: Ablex.
- Pierce, C. S. (1955). Abduction and induction. In J. Buchler (Ed.), *Philosophical writings of Pierce* (pp. 150–156). New York: Dover Books.
- Sperber, D., & Wilson, D. (1986). *Relevance*. Cambridge, MA: Harvard Univ. Press.
- Welty, E. (1941). *A curtain of green & other stories*. San Diego: Harcourt Brace Jovanovich.

(Received July 13, 1999)

(Revision received June 1, 2000)