



## Planning causes and consequences in discourse

Julia Simner\*, Martin J. Pickering

*Department of Psychology, University of Edinburgh, 7 George Square, Edinburgh EH8 9JZ, UK*

Received 23 August 2003; revision received 29 April 2004

Available online 25 December 2004

### Abstract

We investigate the planning of cause and consequence in language production by examining participants' continuations to discourse fragments in four experiments. Our studies indicate how the content of the continuation, and the association between the continuation and prior text, are influenced by the nature of prior discourse. People tend to anchor upcoming utterances to information that is both temporally recent and textually recent (Experiment 1). Furthermore, when discourse fragments describe the cause of an event, continuations tend to describe its consequence, and vice versa (Experiments 2 and 3). Finally, people plan utterances based on prior causality information drawn from world knowledge about the typicality of events (Experiment 4). When this knowledge provides information about causes, people are more likely to plan an upcoming consequence. Hence, people seek to satisfy gaps in their discourse model of the unfolding narrative, and use features of textual and temporal recency to anchor their productions.

© 2004 Elsevier Inc. All rights reserved.

*Keywords:* Causality; Production; Typicality; Satisfied gap; Recency

How do people produce utterances longer than a sentence? The literature on language production has provided little systematic examination of how people combine the parts of a complex narrative into a coherent whole. Formal analyses of narrative structure (e.g., Kehler, 2000) indicate that each contribution to a discourse can link to prior context in various systematic ways. Similarly, research in text comprehension has shown how people assess these linking relationships during reading (e.g., Duffy, 1986). In this paper, we examine how such relations are planned during production, by asking how people decide which part of the narrative a new contribution should be linked to, and what kind of relationship should be captured by this link. We focus

on relationships of causality,<sup>1</sup> where a segment of text can serve as either the cause or the consequence of a related segment.

Research on language production focuses on words and sentences (e.g., Levelt, 1989) with some consideration of dialogue (Clark, 1992; Cleland & Pickering, 2003; Pickering & Garrod, 2004) but very little on narrative, except with regard to specific issues such as the effect of feedback by addressees (e.g., Bavelas, Coates, & Johnson, 2000). However, there is considerable research on the determination of narrative structure within the domain of text comprehension, and it is possible to draw inferences about production from this work. Readers

\* Corresponding author. Fax: +44 131 650 3461.  
E-mail address: [J.Simner@ed.ac.uk](mailto:J.Simner@ed.ac.uk) (J. Simner).

<sup>1</sup> To avoid ambiguity, we use the term *causality* to refer to relations of cause and/or consequence, and reserve the term *causal* to refer to relations only of cause.

impose structure on the texts they encounter in order to assimilate each event into a coherent narrative (Duffy, 1986; Fletcher, Chrysler, van den Broek, Deaton, & Bloom, 1995). In this pursuit, relations of causality (i.e., of cause and consequence) are particularly important. Hence, statements with many causality connections are recalled better than statements with fewer connections (Trabasso & van den Broek, 1985) and accessed from memory more rapidly (O'Brien & Myers, 1987). Readers allocate attention to the chain of causality in a story narrative (Fletcher, Hummel, & Marsolek, 1990) and make inferences to bridge any gap in that chain (Murray, Klin, & Myers, 1993). Evidence of a similar concern in production is suggested by the fact that people overwhelmingly prefer to continue a discourse with information that is related by causality, rather than by temporal sequence alone (Arnold, 2001; Stevenson, Crawley, & Kleinman, 1994; van den Broek, Linzie, Fletcher, & Marsolek, 2000). What is less clear, however, is the exact nature of the causality associations that people produce.

For a model of narrative planning, we start with the assumption that language users strive to make their narratives comprehensible and well formed. People (generally) do not produce unrelated utterances. Instead, language users structure their narratives, and make decisions about both the content and organization of the text (see Bock & Levelt, 1994). Second, we assume that if people wish to produce an utterance that will contribute to a coherent whole, they must monitor their own unfolding discourse. Levelt (1989) suggests that each new contribution made by a speaker is updated in the *discourse model* (Johnson-Laird & Garnham, 1980) of previously generated information, in an activity he calls *bookkeeping*. We assume that this evolving model is consulted during the planning of narrative texts, in order for the language producer to determine what information to convey next and how to convey it. Given this, and the importance of causality in language processing, we ask two questions about the construction of narratives: How do people decide which section of prior discourse to target for continuation (which we refer to as the *anchor*) and how do they decide on the causality content of that continuation? Answers to these two questions would provide the framework for an account of narrative production.

### Selecting an anchor

In a simple text expressing one proposition, such as *Beryl admired John*, a coherent continuation must link back to that proposition. But in longer narratives (e.g., *Beryl applauded John because she admired him*) each new contribution could, in principle, link to any previously expressed proposition (in this case, to *Beryl ap-*

*plauded John* or to *she admired him*). We refer to the point within the existing narrative to which the new contribution links as the *anchor*, and ask how people choose an anchor for each new contribution.

Intuitively, it seems that people would be most likely to select an anchor that is currently relevant, although a moment's reflection suggests that current relevance might be defined in two different ways. One possibility is by *temporal recency*, in that people might link their new contribution to the event in prior context that began most recently in the time-line described by the narrative. When considering discourse fragments involving causality relations, and because causes take place temporally prior to their consequences, temporal recency amounts to the claim that the anchor should be the discourse fragment's most temporally recent consequence. In pure comprehension accounts, this has been called the *current state strategy* (Fletcher & Bloom, 1988; Fletcher et al., 1990) in which readers aim to associate each upcoming event to 'the endmost state in the causal[ity] chain... which has antecedents in the preceding text, but no consequences' (Fletcher & Bloom, 1988, p. 238). To support their claims, Fletcher et al. showed that continuations to this clause are read particularly quickly, and that target words drawn from it are recognized quickly as sentence-final probe words. Although Rizzella and O'Brien (1996) later suggested that narrative comprehension may also be influenced by global goals, the key assumption is that for local coherence, upcoming text is associated to the discourse's most temporally recent consequence. van den Broek et al. (2000) found the same effect in production, by examining continuations to discourse fragments. They show that people routinely continue a narrative from the final consequence, even if this is not the final line of the discourse fragment.

Hence, both production and comprehension studies suggest that language producers select the temporally most recent expression to be the anchor. But it is possible that people also follow *textual recency*, where recently described events are preferred over less recently described events. If the ordering of prior narrative corresponds to the temporal ordering of events, then textual and temporal recency make the same predictions. But in sentences such as *Beryl applauded John because she admired him*, the applauding temporally follows the admiration. In this case, temporal recency predicts that the applauding event should be the anchor, whereas textual recency predicts that the admiring event should be the anchor. A possible influence of textual recency follows from our assumption that writers maintain a discourse model of the unfolding narrative, and that memory for language decays over time (e.g., Sachs, 1967). Hence, people are more likely to use linguistically recent context as anchors, compared to less recent context. Assuming, too, that producers are sensitive to the needs of the addressee, this consideration would also

promote a more successful process of comprehension. Although this account seems straightforward, it has not been championed in the literature on narrative processing. Fletcher and Bloom (1988) did not expressly look for a textual recency effect in comprehension, and van den Broek et al. (2000) found no evidence for it in production (although they did not identically match their discourse fragments across conditions, and so may have missed an effect).

As the literature already provides evidence for temporal recency, we can contrast two hypotheses: temporal recency (which is consistent with van den Broek et al., 2000, and with the current state strategy) and *temporal-textual recency* which predicts the additional influence of textual recency. In Experiment 1, we test between these hypotheses by having participants provide a continuation sentence to complex discourse fragments ordered CAUSE-CONSEQUENCE OR CONSEQUENCE-CAUSE. Temporal recency predicts a simple preference for the consequence as the anchor, and no difference in the number of continuations from each consequence across conditions (since it allocates no special status to textual recency). The temporal-textual recency hypothesis, however, predicts a significantly greater number of continuations to a consequence when it is also the most textually recent event (as in the sequence CAUSE-CONSEQUENCE) compared to when it is not (as in the sequence CONSEQUENCE-CAUSE).

### Selecting the content of the link

We next consider how writers plan the *content* of their continuations, in terms of the relationship between the continuation and the anchor. We again consider causality relations, and ask whether continuations tend to provide the cause or the consequence of the anchoring event. One study that addresses this issue with the discourse-completion task (van den Broek et al., 2000) concludes that writers display a uniform preference for producing consequences (i.e., with the anchor as the cause and the continuation as the consequence). We call this the *unconditional-preference hypothesis*, and test this against a more context-dependent account, based on the assumption that writers modulate their preferences according to the nature of preceding context (see also Arnold, 2001). In Experiment 2, we take the materials from van den Broek et al. (2000) and collect discourse completions to fragments of the passage presented either in or out of context. In this way, we investigate whether discourse completions show a uniform preference for consequential continuations or whether the presence of prior context influences production.

We then elaborate on our findings by showing exactly *how* writers make causality decisions based on context. To do so, we contrast the unconditional-preference

hypothesis with what we term the *satisfied gap hypothesis*, a model of narrative planning in which people seek to fill gaps they perceive in the developing discourse model. In a discussion of question-answering, Levelt (1989) suggests that a co-operative speaker will aim to fill in any gaps in the addressee's knowledge. We argue that a similar goal may govern choices in non-interrogative situations, and specifically, that people may produce causes or consequences according to what they regard as lacking from the preceding discourse (and hence, from the addressee's discourse model). An additional motivation comes from Hilton's (1990) theory of causal explanation, which examines how people explain the primary cause of an event that has more than one contributing factor, such as the derailment of a heavy train on a bent rail when the driver was speeding. In such cases, Hilton suggests that the speaker aims to 'close a gap' in the explainee's knowledge. In this case, the speaker would describe the cause as the excessive speed if the explainee knew about the bent rail, or the bent rail if the explainee knew about the speed. We suggest that similar motivations may govern choices about causality planning. If the discourse has already provided the cause of an event, people may be more likely to present its consequence. Equally, when the consequence has already been given, people may be more likely to present its cause. Finally, any instance of gap satisfaction can be interpreted as an application of Grice's maxim of Quantity (Grice, 1975). This states simply that the language producer should make her contribution just as informative as is necessary, and so should generate information of the type that is missing from the reader's discourse model. We test our satisfied gap hypothesis in three ways.

First, we analyze the content of our materials in Experiment 2, with the prediction that where prior discourse tends to provide the cause of an event, the writer will tend to add its consequence. We then test the theory more directly in Experiment 3, by examining the content of completions to an anchor clause that has been preceded either by its cause or by its consequence. The satisfied gap hypothesis predicts more consequences when the anchor is preceded by its cause than by its consequence, as people strive to provide information of the type that has not already been given in prior context.

Finally, we test the satisfied gap hypothesis with prior discourse that provides implicit rather than explicit causality information. Studies in language comprehension have suggested that judgments of causality are influenced by world knowledge, with people using information about what constitutes normal behavior (e.g., Corrigan, 1988, 1992; Hilton, 1990; Hilton & Slogoski, 1986). For example, Corrigan (1992) provided participants with typical and non-typical scenarios and asked them to rate whether the event was caused by the agent, the patient or neither. Participants were more likely to

attribute the cause to an actor (agent or patient) in typical scenarios (e.g., *The policeman chased the thief*) than in non-typical scenarios (e.g., *The policeman amused the thief*). Corrigan (1992) concluded that typical scenarios provide inherently more information about the cause of the event than non-typical scenarios. Given this, we can use event typicality to manipulate causality information in the following way. In Experiment 4, we present discourse fragments that describe either typical or non-typical events (e.g., *The doctor/waiter healed. . .*). Since typical events come with greater inherent cause (Corrigan, 1992) the satisfied gap hypothesis predicts that such events will produce more consequential continuations than non-typical events.

In summary, we present four studies that use the discourse-completion task to examine the nature of message-planning for causality in discourse production. In Experiment 1, we investigate how people select the anchor for their continuations, and in Experiments 2–4, we examine the content of those continuations.

### Experiment 1

In this experiment, we investigate the mechanism of message-planning that selects the event in prior discourse that is to serve as the anchor for upcoming contributions. We present discourse fragments for completion that describe the same events and causality relations, but which differ on the textual ordering of information (CAUSE-CONSEQUENCE vs. CONSEQUENCE-CAUSE). The temporal recency hypothesis predicts that people will prefer simply to anchor their continuations to prior consequences rather than causes. The temporal-textual recency hypothesis predicts that there will be an additional preference to attach to textually recent events. In other words, people will prefer to continue from consequences that are linguistically recent (in CAUSE-CONSEQUENCE) compared to less recent (in CONSEQUENCE-CAUSE).

#### Method

##### Participants

Thirty-two members of the University of Edinburgh community were paid £3 to participate. Participants were native speakers of English, with no history of dyslexia.

##### Materials

To prepare our stimuli we first wrote 26 CAUSE-CONSEQUENCE sentences of the type: SUBJECT<sub>i</sub> VERB OBJECT<sub>j</sub> so PRONOUN<sub>i</sub> VERB PRONOUN<sub>j</sub>. (These items were based, in part, on the materials of Majid, Sanford, & Pickering, 2004.) For each sentence we then constructed a counterpart version with the inverse causality

structure (CONSEQUENCE-CAUSE) by swapping the position of the two verbs and changing *so* to *because*. This created 26 items pairs like (1) below:

- (1a) Beryl admired John so she applauded him. (CAUSE-CONSEQUENCE)  
 (1b) Beryl applauded John because she admired him. (CONSEQUENCE-CAUSE)

Next, we normed our materials to ensure that neither version was pragmatically more acceptable than its counterpart. For this, we placed all 52 sentences in a single list with 32 fillers that violated pragmatic conventions (e.g., *Rose comforted Peter so she cursed at him*). Half the fillers were structured as CAUSE-CONSEQUENCE, and the other half as CONSEQUENCE-CAUSE. Pseudo-randomization ensured that each CAUSE-CONSEQUENCE target was in the opposite half of the list from its CONSEQUENCE-CAUSE counterpart. Fourteen additional participants from the same population rated the naturalness of each sentence on a scale of 1–7 (where 7 represents highly natural). Twenty-four target items rated over 6.2 were selected, such that there was no difference between the ratings of the CAUSE-CONSEQUENCE and CONSEQUENCE-CAUSE groups (both means = 6.9; both *SDs* = .2; *t* < 1). The target sentences (shown in Appendix A) were placed into two lists, so that each list contained one version of each item, and 12 items from each condition.

#### Procedure

Sixteen participants were randomly assigned to each list, and presented with the materials in a double-sided, 2-page booklet, with approximately six items per page. They were told that the booklet contained a series of unrelated sentences describing different events. They were instructed to work through the booklet in order, and to add a one-sentence continuation that followed suitably from each of the sentences, in the space provided. The instructions were presented on the front of the booklet, and repeated by the experimenter.

#### Results

Completions were assessed by two independent coders to determine which of the two prior context clauses had been selected as the anchor. In the current experiment and in all other experiments reported here, the independent raters had an initial agreement rate of over 90%, and all initial disagreements were resolved by discussion. For (1a) above, an example of a clause 1 coding would be the continuation: *John was a very admirable person*, and an example of a clause 2 coding would be: *John acknowledged her applause*. Table 1 shows the number of continuations that used the cause versus the consequence as an anchor.

Table 1  
Total (and mean number of) continuations from causes and consequences in prior discourse; Experiment 1

Preamble structure	Anchor	
	Cause	Consequence
CAUSE-CONSEQUENCE	67 (2.8)	317 (13.2)
CONSEQUENCE-CAUSE	184 (7.7)	200 (8.3)
PROPORTION	.33	.67

Table 1 shows that consequences were overall more favored than causes (following van den Broek et al., 2000) and this effect is significant ( $Z = 6.71$ ,  $p < .001$ ). In addition to this, more continuations were anchored to the consequence when it was textually recent (i.e., in the CAUSE-CONSEQUENCE condition) than when it was not ( $t_1(31) = 7.2$ ,  $p < .001$ ;  $t_2(23) = 9.6$ ,  $p < .001$ ).

### Discussion

These results provide information about the stage of message-planning in which people make decisions about which event in the narrative to select for continuation. The effect of temporal recency (i.e., the preference for consequences as anchors) was modulated by an additional influence of textual recency, with a consequence in the second clause of the discourse fragment being more likely to be serve as the anchor than a consequence in the first clause. This supports the temporal–textual recency account, and may reflect the fact that memory for language decays sharply across clause boundaries (Sachs, 1967), making any association to earlier clauses more effortful.

### Experiment 2

In this experiment, we consider how people plan the content of continuations to narrative texts. We presented participants with discourse fragments either in a discourse context, or in no context. In the CONTEXT condition, a 295-word short story was taken from the materials of van den Broek et al. (2000, Experiment 1). The discourse appeared (as it had in this earlier study) as a coherent whole with 16 break points that prompt the participant for a continuation. In the NO-CONTEXT condition, the 16 sentences immediately prior to each prompt point were extracted, and presented in a random order that did not preserve continuity. In both conditions, participants provided 16 continuations (either at 16 breaks in a short story, or to 16 ostensibly unrelated sentences). The unconditional-preference hypothesis predicts a simple preference to produce consequences, with no difference across conditions. In contrast, differences between conditions would indicate that choices are based on the nature of prior discourse.

The primary aim of this experiment is to test the unconditional-preference account, and in fact, any difference between conditions would falsify this hypothesis. However, in anticipation of Experiments 3 and 4, we also predict the direction of that difference. The satisfied gap hypothesis states that language producers aim to provide information of the type that has not already been provided by prior context. To understand what predictions this makes for the current study, we analyzed the causality structure of the short story presented in the CONTEXT condition. Fig. 1 shows this story, with each of its component sentences indicated by subscript numbering, and with each of the 16 break-points represented by double slashes. Following Trabasso and Sperry (1985) we represent each sentence of the passage as a node in a causality network (see Fig. 2) with each causality relationship indicated by an arrow (pointing from the cause to the consequence).

Crucially, all arrows point from a lower number to a higher number. This means simply that the prior context of any given event (in this particular passage) serves to provide the cause of that event but not its consequence. Given this, the satisfied gap hypothesis predicts that writers in the CONTEXT condition will tend to produce consequences when prompted for a continuation, because prior context tends to have already provided a cause. In contrast, we should see fewer consequences in the NO-CONTEXT condition, because the absence of context means an absence of prior cause information, and hence less incentive for the participant to produce a consequence.

### Method

#### Participants

Twenty-four new participants from the same population were paid £3 to participate.

#### Materials

The materials in the CONTEXT condition were taken from Experiment 1 of van den Broek et al. (2000); see Fig. 1. For the NO-CONTEXT condition, the final sentence in each fragment was extracted. Thirteen of these fragments referred to the story protagonist, Sally, either by name or with an anaphor. These were changed to 13 different names and the resulting 16 sentences were randomized. This gave rise to 16 individual sentences that no longer formed a coherent passage.

#### Procedure

Twelve participants were randomly assigned to each of two conditions. Those in the CONTEXT condition were presented with the story in a double-sided, 9-page booklet, with one fragment per side. Participants in the NO-CONTEXT condition were presented with 16 single sentences in a double-sided, 2-page booklet, with

There once was a girl named Sally<sub>1</sub>. One day, Sally’s class had ‘show and tell’<sub>2</sub>. Sally’s best friend, Cathy, showed a picture that she had painted<sub>3</sub>. Cathy told the class how she had made it<sub>4</sub>. Sally was jealous of her friend’s project<sub>5</sub>. // Sally wanted to do something special for show and tell<sub>6</sub>. // Sally got an arts and crafts book from the library<sub>7</sub>. That afternoon, she read the art and crafts book<sub>8</sub>. // The book had pictures of candles<sub>9</sub>. // Sally wanted to make a red candle<sub>10</sub>. // Sally found out that candles can be made by melting crayons and pouring them in a cup<sub>11</sub>. She found some crayons and a cup in the kitchen<sub>12</sub>. She put the crayons into the cup<sub>13</sub>. // There were not enough crayons to full the cup<sub>14</sub>. // Sally wanted to buy some crayons<sub>15</sub>. // Sally took some money from her bank<sub>16</sub>. She walked to the shopping mall<sub>17</sub>. Sally found the arts supply store<sub>18</sub>. She went into the arts supply store<sub>19</sub>. She talked to the sales lady about crayons<sub>20</sub>. The lady showed Sally a lot of boxes<sub>21</sub>. Sally picked out a large box of crayons<sub>22</sub>. She asked the lady how much the box cost<sub>23</sub>. The lady told Sally that the box cost £2.50<sub>24</sub>. // Sally bought the box of crayons<sub>25</sub>. // Sally took her new crayons home<sub>26</sub>. Sally melted the new crayons<sub>27</sub>. // Sally held a string in the cup<sub>28</sub>. The wax hardened quickly<sub>29</sub>. Sally had made a beautiful candle<sub>30</sub>. // Sally put her new candle in a holder<sub>31</sub>. She decorated the candle with a ribbon<sub>32</sub>. // the next day, Sally carried the candle to school<sub>33</sub>. Sally asked the teacher if she could be in show and tell<sub>34</sub>. The teacher agreed<sub>35</sub>. // Sally won the grand-prize blue ribbon for her candle<sub>36</sub>. // Sally was glad the teacher liked her show and tell project<sub>37</sub>. //

Fig. 1. Materials of CONTEXT condition in Experiment 2, where double slashes indicate break points at which participants were prompted for their continuation.




Fig. 2. Causality network for passage shown in Fig. 1, where numbered nodes correspond to sentence numberings in Fig. 1.

approximately five items per page. Participants in the CONTEXT condition were told they would have to add a single sentence at each of 16 break-points in a story, and those in the NO-CONTEXT condition were told they would be adding a sentence to each of 16 unrelated sen-

tences. (During debriefing, all participants in this latter condition confirmed that they had not noticed that the 16 sentences had originally been part of a single story.) Participants were asked to produce continuations that fit naturally with what had gone before. A space was

provided after each item for participants' continuations, and participants were asked to work through the booklet in order. Instructions were written on the front of the booklet, and repeated by the experimenter.

#### Causality coding

Participants' continuations were coded as 'cause,' 'consequence,' or 'other,' following the logical tests of Mackie (1980). The requirement of *necessity and/or sufficiency in the circumstances* (see below) characterizes both 'cause' and 'consequence' codings, while the distinction between the two is based on the criterion of temporal sequence: A 'consequence' must take place subsequent to the event in the discourse fragment, while a 'cause' must take place prior to it. Given these conditions, a 'consequence' coding of the continuation event (B, see below) can be defined as cases where either the following is true:

*Necessity:* In the circumstances, had the fragment event (A) not happened, then the continuation event (B) would not have happened.

*Sufficiency:* In the circumstances, had the continuation event (B) not happened, then the fragment event (A) could not have happened.

The criteria for a 'cause' coding of the continuation clause can be seen by reversing the events labeled A and B in the definitions above. Consider the fragment: *That afternoon, she read the arts and crafts book.* An example of 'cause' coding would be the continuation: *She had decided to set up a handy craft stall* (cause + necessity) while a 'consequence' would be: *She learned all about how to make things* (consequence + necessity + sufficiency). Non-causal, non-consequential continuations were coded as 'other' (e.g., *The living room door was wide open*). Coding was performed by two independent markers. In the current experiment and all others reported here, coders had an initial agreement rate of over 90%, and all disagreements were resolved by discussion.

#### Results

Table 2 shows the number of continuations of each type across condition. 'Other' responses were unaffected by the manipulation of context (both  $t_s < 1$ ) and given this, our main analysis compares the proportion of consequences in relation to causes plus consequences (by dividing the number of consequences by the sum of the causes and consequences, both by subjects and items).

Table 2

All continuations by type, and the proportion of consequences (in relation to causes plus consequences); Experiment 2

Prior context?	Causes	Consequences	Other	Proportion consequences
CONTEXT	27	112	53	.81
NO-CONTEXT	49	91	52	.65

Relatively fewer consequences were generated in the NO-CONTEXT condition than in the CONTEXT condition ( $t_1(22) = 2.7$ ,  $p < .02$ ;  $t_2(15) = 1.7$ ,  $p = .11$ ). The lower level of significance in the items analysis is discussed below.

#### Discussion

Participants generated more consequential (and fewer causal) continuations when the discourse fragment was presented in context than when it appeared in isolation. This difference (indeed, *any* difference) between conditions provides evidence against the unconditional-preference account, since it shows that people do not simply produce causality links according to a simple bias for consequences. Instead, we have shown that language producers are sensitive to context when providing continuations to discourse fragments. To account for the direction of our findings, we hypothesized that people would produce information that had not already been provided in prior context, and that this would bias them to produce consequences (since context tended to provide causes).

Our results missed significance by items using a (conservative) two-tailed test. We attribute this to the fact that the items in our CONTEXT condition (taken from van den Broek et al., 2000) differ on the degree to which they satisfy the experimental criterion (e.g., Items 1 and 9 are preceded by no causal information, and so function as equivalent to items in the NO-CONTEXT condition. Item 14, by contrast is a satisfactory candidate with one causal link, while Item 10 is an extreme candidate with three links).

As a final consideration, we point out that the materials of our study confound discourse context with discourse length, and indeed, it is possible to interpret our results as providing evidence that people may simply provide more consequences to longer discourses. In the following studies therefore, we investigate the satisfied gap hypothesis using items that are matched on length, and where causality information is systematically manipulated.

#### Experiment 3

In this experiment, we investigate more closely how the content of continuations is influenced by causality information in previous context, and do so in discourses

that are matched on length. We compare continuations to a clause (e.g., ...*she applauded him*) that is preceded either by its cause (*Beryl admired John so...*) or by its consequence (*Beryl pleased John because...*). We focus specifically on those continuations that treat the second clause as the anchor (because the first clause differed between conditions). The satisfied gap hypothesis predicts that continuations to this (second-clause) anchor should contain more consequences when it has been preceded by its cause than when it has been preceded by its consequence.

### Method

#### Participants

Twenty new participants from the same population were paid £3 to participate.

#### Materials

We began with 53 candidate anchor clauses, each containing an agent-patient verb (selected from the materials of Majid et al., 2004) and pronouns of different genders (e.g., *She applauded him*). Two preambles were generated for each clause, one providing the cause of the target event (*Beryl admired John*) and the other providing its consequence (*Beryl pleased John*). Each preamble was one-clause long, and contained two names of different genders and a transitive verb. We then ensured that the preamble pairs were matched on causality (i.e., that the cause preamble was as causal as the consequence preamble was consequential). To do this, we paired the preambles and their respective target clauses to create 106 two-clause statements, which were then randomized. The list was presented to 20 further participants from the same population as the experimental group, half of whom were asked whether the first event could provide a cause of the second, and the other half, whether it could provide a consequence. Ratings were made on a 1–7 scale, with higher numbers corresponding to more positive judgments. We selected 26 target items such that the causal preambles had a causal rating greater than 5 (mean = 6.3;  $SD = .4$ ), the consequential preambles had a consequential rating greater than 5 (mean = 6.2;  $SD = .5$ ), and there was no significant difference between the groups ( $t < 1$ ). Finally, the conjunction *because* or *so* was inserted between the two clauses of each sentence to make the causality relationship explicit. This provided 26 pairs of items like (2) below (see Appendix B).

- (2a) Beryl admired John so she applauded him. (CAUSE PREAMBLE)  
 (2b) Beryl pleased John because she applauded him. (CONSEQUENCE PREAMBLE)

The items were placed in two lists such that each list contained one version of each item and 13 items from each condition. To each list, we added 30 filler items that did not express causality relations, and which consisted of two-clause sentences employing the conjunctions *before*, *after*, and *while* in ten items each.

#### Procedure

Ten participants were randomly assigned to each list. The procedure was identical to Experiment 1, and the coding method was identical to Experiments 1 and 2.

#### Results

In our analyses, we examined continuations from an anchor clause (the second clause) which was identical in content across conditions, but which had occurred with either a cause preamble, or consequence preamble. (As noted above, no comparisons were made of continuations from the first clause, since this differed in content across conditions.) Before examining the content of participants' continuations, we note first that the second clause was chosen as the anchor more often in the CAUSE PREAMBLE condition than the CONSEQUENCE PREAMBLE condition ( $t_1(19) = 12.7$ ,  $p < .001$ ;  $t_2(25) = 11.1$ ,  $p < .001$ ). This supports the findings of Experiment 1, and van den Broek et al. (2000), and reflects the preference for anchors to be consequences.

Next we examined the content of continuations to these (second-clause) anchors. Table 3 shows the total number of continuations broken down by type, and their corresponding percentages (given the different numbers of second-clause anchors in each condition).

As before, the relative number of 'other' responses did not differ (both  $ts < 1$ ) and so the main analysis considered the proportion of consequences (in relation to causes plus consequences) across conditions. There were relatively more consequences in the CAUSE PREAMBLE condition versus the CONSEQUENCE PREAMBLE condition ( $t_1(19) = 2.3$ ,  $p < .05$ ;  $t_2(25) = 2.5$ ,  $p < .02$ ). Hence continuations tended to contain information of the type

Table 3

Continuations by type (with percentages) from anchor clause (clause 2) and proportion of consequences (in relation to causes plus consequences); Experiment 3

Preamble	Causes	Consequences	Other	Proportion consequences
CAUSE	24 (11%)	165 (74%)	34 (15%)	.87
CONSEQUENCE	20 (21%)	62 (64%)	15 (15%)	.76

that had not already been provided about the anchor in prior context.

### Discussion

This experiment showed a preference for consequences to serve as the anchor, in accord with Experiment 1. More importantly, it provided support for the satisfied gap hypothesis. Participants preferred more consequential continuations when the (second-clause) anchor had been preceded by its cause compared to when it had been preceded by its consequence. We suggest that in planning their production this way, participants were aiming to provide information missing from their model of the anchoring event.

### Experiment 4

In this experiment, we provide a further test of the satisfied gap hypothesis, by considering a case in which information provided about causes and consequences is presented implicitly. In so doing, we can determine whether the satisfied gap mechanism extends to cases where causality is expressed without overt markers (such as the conjunctions *so* and *because*). We presented short narratives for completion, which described either typical or non-typical events. We manipulated event typicality by presenting actions performed either by typical agents (e.g., *The doctor healed Jack*) or non-typical agents (e.g., *The waiter healed Jack*) (see McRae, Ferretti, & Amyote, 1997, for a discussion of agent typicality). Following Corrigan (1992), we assume that typical events come with greater inherent cause than non-typical events. Given this, the satisfied gap hypothesis predicts that participants will produce more consequences in their continuations to the former compared to the latter.

### Method

#### Participants

Twenty new participants from the same population were paid £3 to participate.

#### Materials

Fifty-seven agent-patient verbs were selected from Majid et al. (2004) and presented for norming to 23 participants from the same population as the experimental

group. Each verb was accompanied by five different noun phrases which had been selected to provide examples of both typical and non-typical agents for the verb (e.g., for the verb *healed*, the agents were: *nurse*, *waiter*, *clarinetist*, *potter*, and *doctor*). Participants rated how common it is for each type of person to perform the action described by the verb, on a scale from 1 (very uncommon) to 7 (very common). Items were presented in a 12-page booklet, with approximately five items per page, and with the instructions written on the front of the booklet. Forty-six verbs were selected, for which the most typical agent scored higher than 6 (mean = 6.5;  $SD = .3$ ) and the least typical agent scored lower than 2 (mean = 1.4;  $SD = .2$ ); see Appendix C. We then constructed experimental items by combining each verb with its most typical agent or its least typical agent (in the *typical* and *non-typical* conditions, respectively) and a proper name patient, as shown in example (3) below:

(3a) The doctor healed Jack. (TYPICAL)

(3b) The waiter healed Jack. (NON-TYPICAL)

The materials were placed in two lists, such that each list contained one version of each item and 23 items from each condition.

#### Procedure

Ten participants were randomly assigned to each list. The materials were presented in an 11-page booklet, with approximately six items per page. All other aspects of the procedure were identical to those of Experiment 1, and the coding method was identical to Experiment 2.

#### Results

Table 4 shows the number of continuations by type. The number of ‘other’ responses did not differ across conditions ( $t_1(19) = 1.7$ ,  $p > .1$ ;  $t_2(45) = 1.5$ ,  $p > .1$ ). Our main analysis considered the proportion of consequences (in relation to causes plus consequences), and as predicted, people produced relatively more consequences when the discourse fragment described typical versus non-typical events ( $t_1(19) = 2.4$ ,  $p < .05$ ;  $t_2(45) = 2.0$ ,  $p < .05$ ).

#### Discussion

Our results show that people are more likely to produce consequential continuations to events with typical

Table 4

All continuations by type, and proportion of consequences (in relation to causes plus consequences); Experiment 4

Fragment type	Continuation type			Proportion consequences
	Cause	Consequence	Other	
TYPICAL	146	217	97	.60
NON-TYPICAL	176	165	119	.48

rather than non-typical agents. These results support previous findings that real-world knowledge about event typicality can influence causality reasoning in language processing (Corrigan, 1992; also Hilton, 1990; Hilton & Slogoski, 1986). Corrigan (1992) showed that typical events are judged by readers as having a greater inherent cause than non-typical events. In terms of the satisfied gap hypothesis, this increase in inherent cause should come with an increase in the number of consequences planned for upcoming discourse, and this is what we found. Thus, a consequence was more likely in the TYPICAL than in the NON-TYPICAL condition.

### General discussion

Our findings provide information about two aspects of narrative planning: the choice of anchor in prior discourse, and the content of the continuation generated. First, we showed that people prefer to continue a narrative from previous consequences rather than causes (Fletcher & Bloom, 1988; Fletcher et al., 1990; van den Broek et al., 2000) and from consequences that are linguistically recent rather than less recent. These findings support the temporal–textual recency hypothesis, which states that people prefer to anchor their continuations to events that are both temporally recent (i.e., to consequences over causes) and textually recent (i.e., to final rather than non-final clauses). We suggest that our novel result, the textual recency effect, might follow from the capacity limitations of memory for language (Sachs, 1967). Given that clause boundaries are significant barriers to memory retrieval, it seems reasonable to conclude that people might find it difficult to link a new contribution to a more remote clause. Additionally, people may realize that such a link would be unhelpful for their addressee. This finding accords with other preferences for textually recent contexts during language processing. For instance, anaphors with local antecedents are read faster and more accurately than those whose antecedents are even only marginally more remote (e.g., Murphy, 1985, 1990). It seems then that both comprehenders and producers prefer to make discourse links between elements that are textually close (at least when these latter are completing narratives they have recently comprehended).

Our remaining studies provide evidence for the satisfied gap hypothesis, which states that people plan causality information of the type that has not already been provided in previous discourse. We showed that people produce more consequences when the fragment for completion is presented within a context (providing causal information about events thus far) than without context (Experiment 2). Any difference between conditions argues against the unconditional-preference ac-

count, by showing that causality planning involves more complex considerations than a bias simply to produce consequences. We suggest that the direction of the effect (more consequences within vs. without context) supports the satisfied gap hypothesis. When participants anchor their continuation to the most recent discourse fragment of the story (at each prompt point), the context thus far has already provided the cause of that anchoring event. The satisfied gap hypothesis therefore predicts that continuations should indeed tend towards consequences.

This finding has implications for the practice of using completion norms to estimate the frequency with which linguistic forms occur in natural language. Many researchers in psycholinguistics have participants complete sentence fragments as an alternative to performing corpus counts, and the two methods for assessing frequency are generally treated as interchangeable. Our work implies that this assumption should be questioned, because it shows that people produce different types of structures according to whether the fragment is, or is not, embedded in a longer narrative. Hence, completions involving single sentence fragments may not reflect what happens in longer texts.

In Experiment 3, we provided clearer evidence for the satisfied gap hypothesis, by matching our discourse fragments on length, and systematically varying causality information. We examined the content of continuations to an anchor clause that had been preceded either by its cause or by its consequence. Participants were more likely to produce a consequence in the former case, suggesting that people seek to provide causality information of the type not already given. In Experiment 4, we showed that contextual information regarding causality can be drawn not only from explicit statements of cause and consequence, but also from more implicit sources. Following Corrigan (1992) we assumed that passages describing typical events carry more implicit cause than those describing less typical events. Accordingly, we found that discourse fragments describing actions performed by typical agents (e.g., *The doctor healed...*) gave rise to more consequential continuations than those performed by non-typical agents (e.g., *The waiter healed...*). This supports the satisfied gap hypothesis, and suggests that people use knowledge about the typicality of events to plan the causality structure of subsequent utterances.

It is perhaps appropriate at this point to briefly examine the claims of Hilton (1990; also Hilton & Slogoski, 1986) about the effects of event typicality on causal reasoning. Hilton has suggested that people tend to explain the cause of an event in terms of its most abnormal aspect. At first sight, this might appear to contradict the premise of Experiment 4, which assumed that cause is more clearly attributable when the agent is typical (Corrigan, 1992). However, the specific circum-

stances of Hilton's claims reveal no true incompatibility. Hilton (1990) distinguishes between *ascribing* cause and *explaining* cause. He found that people tend to ascribe cause to the most typical aspect of an event, but explain cause with respect to the less typical aspect. For instance, he refers to an unpublished study in which participants were given the sentence: *Maureen telephoned the operator*, and rated highly as an explanation: *... because the operator is the sort of person whom people phone*. In other words, participants ascribed cause highly to the typical entity. At this point, the satisfied gap hypothesis would predict that writers should produce more consequences when this typical event (with inherent cause) is presented in a discourse-completion task. However, Hilton instead asked his participants to write a *cause* continuation (with the prompt *because...*). Crucially, writers did not repeat the same type of information as before (i.e., they did not write *... because operators are the type of people we tend to call*). Instead, Hilton claims, the writer assumed her reader necessarily already knew this, and aimed to provide information not already known (see also Garnham & Crinean, 2004). This gave rise to causal explanations that focused on the least typical entity (i.e., Maureen; e.g., *...because she wanted an early wake-up call*). Hence, people do not explain cause with respect to the typical element even though they rate this element highly causal. Instead, they tend to mention what they assume is not already known to their addressee. Indeed, this aim to fill gaps in the addressee's knowledge is precisely the notion captured by our satisfied gap hypothesis.

Our studies have provided evidence for the two mechanisms of temporal–textual recency and gap satisfaction. We suggest that these mechanisms must work in tandem to guide the message-planning of discourse in production. Indeed, we have implicitly argued that gaps in narrative structure are defined according to the anchoring event. To see this, consider again the materials of Experiment 1, where otherwise identical discourse fragments are ordered either CAUSE-CONSEQUENCE or CONSEQUENCE-CAUSE. The satisfied gap hypothesis working alone would predict that people generate the same proportion of causes and consequences in both cases, since both fragments present exactly the same information (and hence the same gaps to be filled). However, the mechanism of temporal–textual recency necessarily highlights some gaps over others, simply by making people more likely to select certain anchors over others. Specifically, the effect of temporal–textual recency is to make consequential anchors attractive, and indeed *most* attractive when the fragment is ordered CAUSE-CONSEQUENCE. Since people are more likely to anchor to consequences in this particular condition (and because these anchors come with a context that has already described their cause)

they should generate fewer causal continuations than the reverse ordering. Indeed, a reconsideration of the data produced in Experiment 1 supports this. A post hoc examination of continuation types (with coding procedures identical to those used in Experiments 2–4) reveals that significantly fewer causes were generated when fragments were ordered CAUSE-CONSEQUENCE (mean 3.6) compared to the reverse ordering (mean 5.7:  $t_1(31) = 3.8$ ,  $p < .001$ ;  $t_2(23) = 3.78$ ,  $p < .01$ ). As we have suggested, this likely arose because people targeted consequences more often in this former condition, and this in turn limited the number of causes generated (as a result of 'gap satisfaction'). In summary then, we suggest that the planning of the content of a subsequent utterance will necessarily depend on a consideration of where that utterance is to be anchored.

An important question is the extent to which our results relate to the use of the discourse-completion method—a hybrid task, involving both production and comprehension. Specifically, the task involves the input of a second party (the experimenter) in the construction of the discourse, and participants must continue what has already been provided, rather producing a narrative that is entirely of their own making. It is not known how processing might be affected by such conditions, although the technique is common in studies of language production (e.g., Meyer & Bock, 1999; Pickering & Branigan, 1998; van den Broek et al., 2000). One problem is that ecological validity may suffer if the experimenter has prepared a fragment with language/style that the participant would not, herself, have produced. Naturally, all efforts to avoid this were taken (e.g., verbs were vetted during norming to ensure they were familiar to the participant population). A second problem is that adding to a supplied discourse, compared to a self-generated one, might simply engender inherently different mechanisms. However, this criticism equally holds for any method that elicits production using experimental stimuli (e.g., picture description; see Bock, 1996). Third, some decisions about what to say in upcoming utterances may be made very much earlier, or may be partially influenced by an overall plan. In this way, our relatively short discourse fragments (Experiments 2–4) might have exaggerated the influence of recent context. Nonetheless, studies have shown that local coherence plays a prominent role in both the comprehension and production of causality (Fletcher & Bloom, 1988; Fletcher et al., 1990; van den Broek et al., 2000). Indeed, this dominance can be seen in the causality network shown in Fig. 2, where the majority of causality associations are between sentences that are textually adjacent. Given these considerations, it may still be unwise to generalize from discourse completions to unprompted language production, although there is certainly no reason to believe that one act is more theoretically

important than another. Production that takes place in response to contributions by interlocutors may in fact represent a more usual form of communication (Pickering & Garrod, 2004) as in cases where dialogues unfold, or couples cooperatively describe an event, or indeed, where authors jointly prepare a manuscript. Discourse completion could therefore serve as a laboratory analogue of joint acts of production, as found in such real-world situations.

A more theoretically interesting question is whether our conclusions hold for all literary genres, or whether they are limited to the kinds of descriptive narratives that our participants have engaged in. One possibility is that the processes we describe might occur when people are relating an event, but not, for example, when they are describing a scene or producing a list (where propositions are much less likely to have clear relationships of causality). Additionally, the processes that we identify might be modified in tightly coupled dialogue, when the anchors and relationships of causality might hold within or between speakers. But in the construction of narrative discourses of the type examined here, we suggest that people are indeed directed by the considerations we have discussed. We argue that people determine the anchor both on the basis of the narrative's structure, and on the nature of the events described. We suggest that people seek to fill gaps in their narrative, and furthermore, that they consider an absence of either cause or consequence to be one such gap. More generally, we conclude that people do not simply have uniform preferences for producing causality relations, but rather, that they select between cause and consequence in response to fine-grained features of the discourse.

### Acknowledgments

This research was supported by the first author's British Academy Postdoctoral Fellowship (PDF/2001/345). Many thanks to Bryn Harris and Sarah Haywood for their help in data coding.

### Appendix A

Materials for the CAUSE-CONSEQUENCE condition of Experiment 1. In the CONSEQUENCE-CAUSE condition, the order of verbs is reversed and the conjunction is *because*.

1. Simon heard Anna so he answered her.
2. Beryl admired John so she applauded him.
3. Thomas disliked Debbie so he criticised her.
4. Arnold loved Jennifer so he helped her.
5. Jonathan feared Nicola so he obeyed her.
6. Bob disapproved of Sally so he reprimanded her.
7. Ian suspected Amy so he accused her.

8. Beth fancied Rob so she bantered with him.
9. Carol wanted Trevor so she chased him.
10. Claire trusted Gordon so she confided in him.
11. Gemma spotted Adam so she greeted him.
12. Jane detested Ted so she harassed him.
13. Ted loathed Jane so he harmed her.
14. Jennifer was angry with Arnold so she hit him.
15. Kate objected to Donald so she insulted him.
16. Donald disagreed with Kate so he interrupted her.
17. Stephen liked Liz so he joked with her.
18. Lorna despised Henry so she kicked him.
19. Peter hated Rose so he pushed her.
20. Sue opposed Luke so she slandered him.
21. Jack adored Susan so he stared at her.
22. Susan needed Jack so she stopped by him.
23. Suzanne respected Richard so she supported him.
24. Veronica feared for Callum so she warned him.

### Appendix B

Materials for Experiment 3. The first and second verbs show the CAUSE and CONSEQUENCE preambles, respectively. The space '...' is filled by *so* in the former, and *because* in the latter.

1. Beryl admired/delighted John... she applauded him.
2. Thomas disliked/provoked Debbie... he criticised her.
3. Arnold was worried for/thanked by Jennifer... he helped her.
4. Jonathan was scared of/pleased Nicola... he obeyed her.
5. Bob didn't like/influenced Sally's behavior... he reprimanded her.
6. Ian suspected/embarrassed Amy... he accused her.
7. Beth fancied/impressed Rob... she bantered with him.
8. Carol wanted/caught Trevor... she chased him.
9. Claire trusted/surprised... she confided in him.
10. Gemma noticed Adam arriving/made a good impression on Adam... she greeted him.
11. Simon heard/enlightened Anna... he answered her.
12. Jane detested/sickened Ted... she harassed him.
13. Ted hated/frightened Jane... he harmed her.
14. Jennifer was angry with/disliked by Arnold... she hit him.
15. Joan hated/frightened Barry... she hurt him.
16. Kate disliked/offended Donald... she insulted him.
17. Donald had enough of/offended Kate... he interrupted her.
18. Stephen wanted to make Liz relax/relaxed Liz... he joked with her.
19. Lorna was furious/in trouble with Henry... she kicked him.
20. Peter got angry with/scared Rose... he pushed her.
21. Sue hated/infuriated Luke... she slandered him.
22. Luke disliked/angered Sue... he snubbed her.
23. Jack fancied/annoyed Susan... he stared at her.
24. Susan needed/startled Jack... she stopped by him.
25. Suzanne admired/pleased Richard... she supported him.
26. Veronica was worried for/saved Callum... she warned him.

## Appendix C

Materials for Experiment 4. Agent-patient verbs accompanied by their typical/non-typical agents. Discourse fragments were structured: AGENT + VERB + PROPER NAME.

healed: doctor/waiter; applauded: fan/shopper; criticised: politician/busker; hit: boxer/make up artist; reassured: psychiatrist/thief; answered: telephone operator/monster; censured: complaints representative/tap dancer; judged: judge/fishmonger; praised: supporter/shoplifter; reprimanded: headmaster/typist; accepted: foster parent/dissident; accused: prosecutor/cellist; calmed down: psychiatric nurse/rioter; helped: lollipop lady/miser; placated: hijack negotiator/terrorist; approached: reception representative/recluse; bantered with: chat show host/monk; betrayed: double agent/superhero; chased: policeman/invalid; cheated: card shark/vicar; cheered: cheerleader/surgeon; cheered up: flower deliveryman/executioner; commanded: president/wimp; competed with: athlete/Buddhist; confided in: repentant sinner/Mafia hitman; defied: naughty pupil/teacher's pet; denounced: Greenpeace activist/undertaker; encouraged: personal trainer/cynic; greeted: reception committee member/mail bomber; harassed: stalker/apple-picker; harmed: gunman/charity worker; hurt: sniper/aid-worker; joked with: comedian/hangman; kicked: kick boxer/ambulance driver; killed: murderer/chiropractist; manipulated: con man/midwife; murdered: killer/ball boy; ordered around: Duchess/servant; protected: bodyguard/weakling; pushed: wrestler/parson; rushed to: first-aid expert/sunbather; stared at: artist/canoelist; stopped by: visitor/hermit; supported: campaign manager/patient; tied up: kidnapper/newspaper boy; warned: referee/baby.

## References

- Arnold, J. E. (2001). The effect of thematic roles on pronoun use and frequency of reference continuation. *Discourse Processes*, 31, 137–162.
- Bavelas, J. B., Coates, L., & Johnson, T. (2000). Listeners as co-narrators. *Journal of Personality and Social Psychology*, 79, 941–952.
- Bock, J. K. (1996). Language production: Methods and methodologies. *Psychonomic Bulletin & Review*, 3, 395–421.
- Bock, J. K., & Levelt, W. (1994). Language production. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 945–984). San Diego: Academic Press.
- Clark, H. H. (1992). *Arenas of language use*. London: University of Chicago Press.
- Cleland, A. A., & Pickering, M. J. (2003). The use of lexical and syntactic information in language production: Evidence from the priming of noun-phrase structure. *Journal of Memory and Language*, 49, 214–230.
- Corrigan, R. (1988). Who dun it? The influence of actor-patient animacy and type of verb in the making of causal attributions. *Journal of Memory and Language*, 27, 447–465.
- Corrigan, R. (1992). The relationship between causal attributions and judgements of the typicality of events described by sentences. *British Journal of Social Psychology*, 31, 351–368.
- Duffy, S. A. (1986). Role of expectations in sentence integration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 208–219.
- 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., Chrysler, S. T., van den Broek, P. W., Deaton, J. A., & Bloom, C. P. (1995). The role of co-occurrence, coreference, and causality in the coherence of conjoined sentences. In R. F. Lorch, Jr. & E. J. O'Brien (Eds.), *Sources of coherence in reading*. Hillsdale, NJ: Erlbaum.
- 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.
- Hilton, D. (1990). Conversational processes and causal explanation. *Psychological Bulletin*, 107, 65–81.
- Hilton, D., & Slogoski, B. R. (1986). Knowledge-based causal attribution: The abnormal conditions focus model. *Psychological Review*, 93, 75–88.
- Garnham, A., & Crinean, M. (2004). Amusing comedians can be difficult—implicit causality and role names. Poster presented at AMLaP (Architectures and Mechanisms for Language Processing) 2004, Aix-en-Provence, France.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Speech acts* (Vol. 3). New York: Academic Press.
- Johnson-Laird, P., & Garnham, A. (1980). Descriptions and discourse models. *Linguistics and Philosophy*, 2, 371–393.
- Kehler, A. (2000). Coherence and the resolution of ellipsis. *Linguistics and Philosophy*, 23, 533–575.
- Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
- McRae, K., Ferretti, T. R., & Amyote, L. (1997). Thematic roles as verb-specific concepts. *Language and Cognitive Processes*, 12, 137–176.
- Mackie, J. L. (1980). *The cement of the universe*. Oxford: Clarendon.
- Majid, A., Sanford, A. J., & Pickering, M. J. (2004). *Do interpersonal verbs lead to focus on causes or consequences?* Manuscript submitted for publication.
- Meyer, A. S., & Bock, K. (1999). Representations and processes in the production of pronouns: Some evidence from Dutch. *Journal of Memory and Language*, 41, 281–301.
- Murphy, G. L. (1985). Processes of understanding anaphora. *Journal of Memory and Language*, 24, 290–303.
- Murphy, G. L. (1990). Interpretation of verb phrase anaphors: Influences of task and syntactic context. *The Quarterly Journal of Experimental Psychology A*, 42, 674–692.
- Murray, J. D., Klin, C. M., & Myers, J. L. (1993). Forward inferences in narrative text. *Journal of Memory and Language*, 32, 464–473.
- O'Brien, E. J., & Myers, J. L. (1987). The role of causal connections in the retrieval of text. *Memory & Cognition*, 15, 419–427.
- Pickering, M. J., & Branigan, H. P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language*, 39, 633–651.
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27, 169–225.

- Rizzella, M. L., & O'Brien, E. J. (1996). Accessing global causes during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1208–1218.
- Sachs, J. S. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception & Psychophysics*, 2, 437–442.
- Stevenson, R. J., Crawley, R. A., & Kleinman, D. (1994). Thematic roles, focus and the representation of events. *Language and Cognitive Processes*, 9, 519–548.
- Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and importance of story events. *Journal of Memory and Language*, 24, 595–611.
- Trabasso, T., & van den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language*, 24, 612–630.
- van den Broek, P., Linzie, B., Fletcher, C., & Marsolek, C. D. (2000). The role of causal discourse structure in narrative writing. *Memory & Cognition*, 28, 711–721.