NORTHWESTERN UNIVERSITY

The Use of Causal Coherence Information in Disambiguating Pronouns

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Psychology

By

Eyal Sagi

EVANSTON, ILLINOIS

December 2011

ABSTRACT

The Use of Causal Coherence Information in Disambiguating Pronouns

Eyal Sagi

This dissertation looks at the people use causal coherence information is used to determine the referent of a pronoun in sentence pairs, such as *John accused Mark of stealing a car. He called the police*. The experiments reported here test the idea that the process of pronoun disambiguation is informed by the causal connection between the person denoted by the pronoun (e.g., the one who called the police) and each of its potential referents (John and Mark). Experiment 1 varied how likely it is that the event of the first sentence (e.g., the accusation) would cause the event of the second (the notification) for each of the two individuals, and required participants to choose the appropriate referent. Participants selected the antecedent that had high causal strength or split their votes if neither did. Experiment 2 used explicit connectives that were either causal (*and so, because*) or temporal (*before, after*) and demonstrated that participants were more confident in their use of causal information when the connective was causal than when it was temporal. Experiment 3 explored the influence of negation on participants' choice of referent. Participants' ratings were affected by negation, but their choices of referent and their confidence in their choices were not. This indicates that participants use causal information when choosing the appropriate referent rather than relying on likelihood of occurrence. Finally, Experiment 4 used a self-paced reading paradigm to explore the effect pronoun ambiguity and reference has on participants reading time of two types of causal coherence relations – *result* and *explanation*. Participants read sentence pairs connected by *result* relations fastest when an unambiguous pronoun referred to the object. They read sentence pairs

connected by *explanation* fastest when an unambiguous pronoun referred to the subject. The results of these four experiments suggest that the process of causal coherence information and the process of pronoun disambiguation interact. These findings are in line with models of inter-sentential coherence in which discourse coherence is achieved by relating adjacent discourse segments by means of *coherence relations* (e.g., Hobbs, 1985; Kehler, 2002).

ACKNOWLEDGEMENTS

I would like to thank my committee, Lance Rips, Dedre Gentner, and Sid Horton, whose comments and direction have been invaluable to me. I am also thankful for the thoughtful conversations I had with various members of Lance Rips' lab over the design and analysis of the experiments I report on here and especially my conversations with Winston Chang on the nature of causality and causal reasoning. Finally, I would like to thank the research assistants who helped collect the data I used in writing this dissertation – Rachel Bhagwat, Sam Johnson, Rachel Ranney, and Samantha Thompson.

Table of Contents

List of figures

List of tables

1 Introduction

Inter-sentential coherence is often conveyed by the use of various connectives, such as the English words *therefore*, *likewise*, and, *however*. The variety of meanings conveyed by these connectives suggests that there is more than one way to combine two sentences into a coherent unit. For example, the sentence pair (1) can be combined using "as a result" with the resulting utterance being (2):

(1)

a. The surgery was completed with no undue complications.

b. The patient recovered within a week.

(2)

a. The surgery was completed with no undue complications.

b. As a result, the patient recovered within a week.

The use of the connective "as a result" in this case makes the causal link between the two sentences explicit. However, even if the connective is omitted readers are likely to infer that the quick recovery of the patient not only follows the surgery but that it was quick *because* there were no complications during the surgery. Consequently, both (1) and (2) seem like coherent texts. While connectives are often employed in speech and in writing, they are frequently omitted when the specific relationship between the sentences is easily determined.

On the other hand, there are many cases in which the relationship between adjacent sentences is unclear. In sentence pairs such as (3) it might be more difficult to determine how the two sentences relate. Such cases often become easier to understand if the connective *nevertheless* is added between the sentences, as in (4), because the connective explicates the nature of the relationship between the sentences.

(3)

a. The attending physician was incompetent.

b. The patient recovered within a week.

(4)

a. The attending physician was incompetent.

b. Nevertheless, the patient recovered within a week.

Connectives like *nevertheless* signal the existence of a specific relationship between the sentences. However, these relationships can exist even in the absence of a specific connective. Hobbs (1985) argues that inter-sentential coherence is predicated on the existence of these relationships. In other words, in order for a text to be perceived as coherent readers must be able to identify the specific relations that pertain between different parts of it.

In this dissertation I investigate how world knowledge about causes and effects and the consequent inference of a *cause-effect* coherence relation (which I will term *result* after Kehler, 2004) interacts with the process of pronoun disambiguation. In what follows, I will first describe the general theoretical frameworks that deal with coherence relations. Next, I will present some empirical evidence supporting the claim that these relations are computed during the reading

process. I will then propose how the constraints imposed by coherence relations and pronoun referents might interact and lead to the selection of an interpretation for a sentence. Finally, I will describe the results from 4 experiments that explore the implications of this hypothesis.

These experiments demonstrate the link between general reasoning about causes and effects and the interpretation of texts that describe individual events, some of which might be causes or effects of other described events. It is important to draw a distinction here between the background knowledge that a reader draws upon during the process of reading and the information incorporated into the representation of the text. While the notion that a patient's quick recovery might be one result of a successful surgery probably exists as part of our general knowledge of the world, its relevance to a specific piece of discourse such as (1) needs to be determined when the text is read. Specifically, there are many possible inferences that can be drawn about a pair of sentences and some of these inferences are unlikely to hold simultaneously (e.g., that the surgery was open heart surgery and that it was performed by a neurosurgeon). In order for a text to be coherent, readers need to be selective about the inferences they choose to draw about a text.

That readers are selective in the inferences they draw is a basic assumption in many cognitive models of reading. For example, Mckoon and Ratcliff's (1992) Minimalist Hypothesis suggests that readers draw only the necessary and easily available inferences. Likewise, Kintsch's (1988, 1998) Construction-Integration model in which readers add to the propositions provided in the text based on a model of spreading activation. According to Kintsch, working memory capacity constrains the number of the inferences that a reader can carry over from one part of the text to the next. While the Construction-Integration model might predict that a reader would consider many possible inferences, only a few of those will actually be incorporated into

their representation of the text. Both of these theories argue that cognitive factors, such as effort and memory capacity, limit the inferences that readers draw as they progress through the text.

Because coherence relations are presumed to be a (relatively) small set of possible relations, they offer a basis for a cognitively and computationally constrained process of selecting and evaluating possible inferences that might contribute to local coherence. However, the literature in psycholinguistics has largely focused on the processing of context dependent bridging inferences. As an example of a bridging inference it can be noted that (1b) makes reference to a patient that is not directly referenced in (1a). In order to identify who the patient is the reader is required to make a bridging inference that links this reference with the (previously unmentioned) object of the surgery described in (1a).

Whereas coherence relations are generic and the same coherence relation can be applied to many different contexts, bridging inferences are specific to the context in which there are drawn. For instance, the surgery-patient relation in (1) requires an inference that is unlikely to occur in most texts since it requires a specific medical context while a *result* relation can occur in almost any text. In essence, coherence relations suggest a top-down approach to coherence where specific relations are imposed to link disparate parts of the text. In contrast, bridging inferences are more congruent with a bottom-up approach to textual coherence where inferences are drawn based on the text and relevant ones are used while irrelevant ones are discarded. As will be discussed later, it is possible for bridging inferences to incorporate aspects of coherence relations suggesting that the two types of inference might be related.

In this dissertation I focus on coherence relations rather than bridging inference. Specifically, I examine how causal information might inform the selection of a specific causal coherence relation. Moreover, following Kehler (2002), I propose that the process of identifying

specific coherence relations also involves and constrains the disambiguation of various sentential ambiguities, such as the determination of referents for pronouns.

For the purposes of this dissertation, I will make several simplifying assumptions. These assumptions are necessary because there is little empirical evidence relating to the processing of coherence relations by readers. Firstly, I will assume that the set of coherence relations contemplated by readers is limited. This assumption is in line with Kehler (2002)'s account of coherence relations which identifies 3 general classes of relations with a handful of specific relations in each class. Nevertheless, I believe it is unlikely that readers consider all of these relations at every junction in the discourse. Rather, it is likely that readers initially explore only a few likely relations are and only if those fail to generate a coherent discourse other avenues are explored in turn.

Secondly, many of the theories on coherence relations (e.g., Hobbs, 1985) assume that these relations are necessary for a discourse to be coherent and that such relations are the only way in which coherence is achieved. This assumption appears to be at odds with the ample psycholinguistic evidence that suggests bridging inferences play a significant role in achieving such discourse coherence (e.g., Singer at al., 1992). As will be discussed later, there appears to be a correspondence between coherence relations and some common bridging inferences. Moreover, since this dissertation focuses on causal information, it is likely that this assumption will hold for the materials used. I will return to both of these assumptions in the discussion and examine how the empirical evidence from the experiments reflects upon them.

I make no explicit assumptions regarding the cognitive representation of coherence relations other than that they exist. While I will follow Kehler's definition of these relations, I only require that causal information is the basis on which causal relations are determined but

take no position as to the process that integrates such causal information into a possible interpretation of the discourse. I believe that this latter requirement is reasonable as it underlies our basic notion of how causal information informs our understanding of texts and assists in establishing textual coherence.

I will, however, assume that interpretations based on distinct causal information are themselves distinct. This assumption is critical for most propositional accounts of meaning and I therefore do not believe it to be controversial. However, it is possible that readers do not form such complete interpretations in all cases. In that case, the processing model I outline later will need to be adapted. I suspect such an adaptation, while not trivial, would not affect the basic premise of the model as it applies to this dissertation. The different interpretations I deal with here are based on differences in the referents of pronoun. These differences do not lend themselves to overlapping representations.

Before presenting the experimental results that make up the bulk of this dissertation, I will first describe theories of coherence relations and briefly summarize the existing evidence regarding the processing of pronoun disambiguation. I will begin by describing a few of the main theoretical approaches to coherence relations developed by linguists. I then explore empirical evidence that suggests that these relations are computed by readers as part of the process of discourse comprehension. Next I will examine the relationship between bridging inferences and coherence relations. Then I will turn to the psycholinguistic literature that deals with pronoun disambiguation in general and the biases introduced by the causal implications of some verbs in particular. Finally, I discuss how coherence relations might be cognitively processed and identified.

1.1 Theoretical Approaches to Coherence Relations

The focus of this dissertation is on the interaction between causal coherence relations and pronoun disambiguation. To better understand this interaction it is important to first understand the nature of coherence relations and the role they play in discourse comprehension. In this chapter I describe the main theories and frameworks that have been used to explicate the notion of coherence relations and lay the foundation for exploring their effect on pronoun disambiguation.

There have been several attempts at identifying and cataloguing the type of coherence relations that can exist between adjacent sentences. Most of these efforts are based on examining linguistic sources and evidence, with the express purpose of creating a taxonomy of such relations (e.g., Asher & Lascarides, 2003; Hobbs, 1985; Kehler, 2002, 2004; Mann & Thompson, 1988). A few other accounts focus on the role coherence plays within the wider scope of discourse (e.g., Polanyi, 1985; Polanyi et al., 2003).

One influential account of coherence relations was proposed by Jerry Hobbs (1979, 1985). Hobbs presented a hierarchical theory of discourse structure that has formed the basis of many contemporary theories in linguistics. According to Hobbs, clauses are linked through relations (termed 'coherence relations') such as *Occasion*, *Explanation*, and *Elaboration*. For example, the relationship in (5) between (a) and (b) is that of an explanation – (b) provides an explanation as to the reason behind the occurrence of (a), whereas the relationship in (6) between (a) and (b) is that of a parallel – both introduce similar propositions. Furthermore, each type of coherence relation is either a *subordinating* relation or a *coordinating* one. In a subordinating

relation, one discourse segment dominates the other (e.g., statement (b) in (5) expands on information given in (a) and therefore (b) is a subordinate of (a)), while in a coordinating relation both are of equal standing (e.g., statements (a) and (b) in (6)).

(5)

a. Mary was late to work.

b. She overslept.

(6)

a. Mary likes John.

b. John likes Jane.

When applied to a pair of discourse segments, a coherence relation ties the two together into a unit through a process of composition, the result of which is a summary of the unit. In the case of a subordinating relation, the summary is basically the dominant assertion. A coordinating relation revolves around the identification of a common proposition that can be inferred from both segments. The discourse segment resulting from a coordinating relation is then generated through a generalization over its constituent segments. Because the result of the composition of two discourse segments results in the inference of a new, combined, segment the resulting structure is a hierarchy of discourse segments of ever-increasing complexity, culminating in a single overall representation of the entire discourse as a highly complex structure (such as the one depicted in Figure 1). For example the sentences in (7) could be described using the discourse structure depicted in Figure 1 (patterned after Hobbs, 1985, 51a-f).

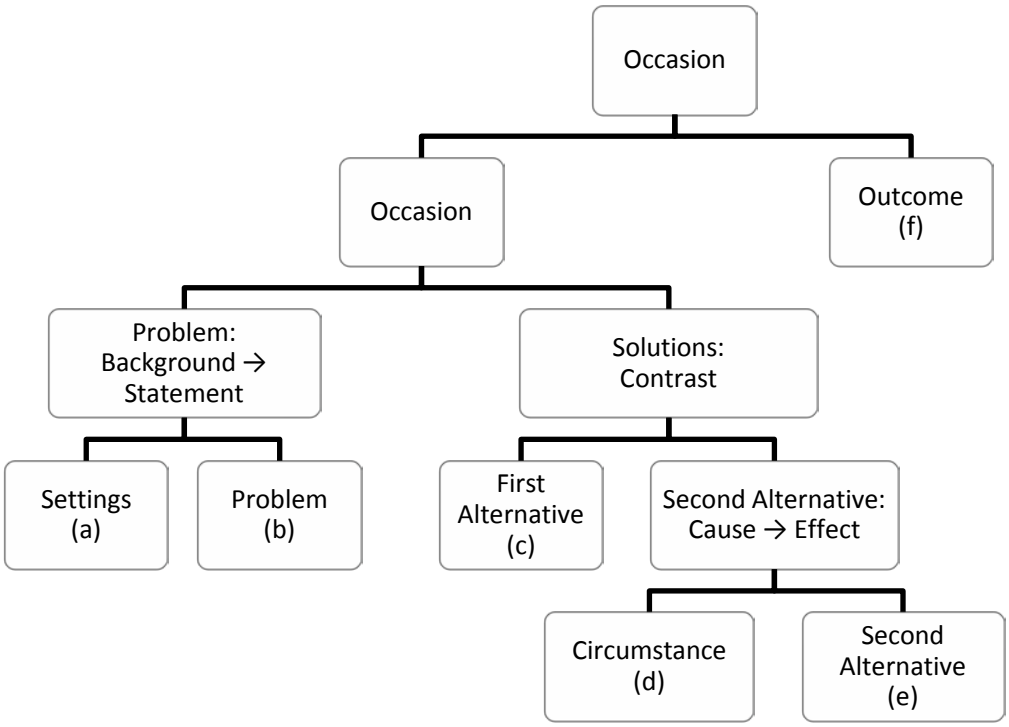

Figure 1 – A possible coherence-relation structure for a 6-sentence story (after Hobbs, 1985, Figure 5)[1]

(7)

a. John, a 30 year old college graduate, started out as a programmer at a company
b. He didn't like any of the people he was working with
c. His first thought was to quit and find a different job
d. But the manager offered him a substantial raise because he was such a hard worker
e. The manager suggested that John might be happier working with a different group
f. So John took the raise and moved to a different group instead of quitting

[1] It should be noted that this is probably not the only discourse structure that could be derived from the story. Hobbs does not claim that there is a unique structure associated with a discourse, but rather suggests a framework in which the structure of a discourse can be described.

Perhaps the most reasoned account of coherence relations was proposed by Kehler (2002, 2004). Kehler suggests that the taxonomy of coherence relations should follow the Humean classes of "connections among ideas" – *resemblance*, *contiguity in time and space*, and *cause or effect*. Kehler proposes that these coherence relations differ in the inferencing process underlying their evaluation. Within each of the three classes, different relations share the same inferencing process but differ in the specifics of their application. For instance, while a *result* relation uses causal information directly and entails that the second assertion be an effect of the first assertion, a *violated expectation* relation entails that the negation of the second assertion be the normal consequence of the first assertion. A similar distinction is drawn between the resemblance relations *parallel* and *contrast*.

Because Kehler's taxonomy is well organized and relatively simple, I will mostly refer to his terminology and definitions in this dissertation.[2] As such it is worth describing his definitions of the relations that will later be used in the experiments (the definitions below are taken from Kehler, 2004):

[2] It is important to note here that although this dissertation focuses on *cause-effect* relations, specifically the direct relation which Kehler terms *result*, my use of Kehler's definition does not generally detract from the applicability of the analysis and results to the other theories. This is the case because there is overall agreement among the theories regarding the categorization of cause-effect relations.

- Cause-Effect relations
    - **Result**: Infer P from the assertion of S_1 and Q from the assertion of S_2, where normally $P \rightarrow Q$.
    - **Violated expectation**: Infer P from the assertion of S_1 and Q from the assertion of S_2, where normally $P \rightarrow \neg Q$.
    - **Explanation**: Infer P from the assertion of S_1 and Q from the assertion of S_2, where normally $Q \rightarrow P$.
- Resemblance relations
    - **Parallel**: Infer $p(a_1,a_2,\ldots)$ from the assertion of S_1 and $p(b_1,b_2,\ldots)$ from the assertion of S_2, where for some vector of sets of properties q, $q_i(a_i)$ and $q_i(b_i)$ for all i.
    - **Contrast (i)**[3]: Infer $p(a_1,a_2,\ldots)$ from the assertion of S_1 and $\neg p(b_1,b_2,\ldots)$ from the assertion of S_2, where for some vector of sets of properties q, $q_i(a_i)$ and $q_i(b_i)$ for all i.

As an example, in (5) the presumed relation is one of *explanation*. The first sentence (S_1) is "Mary was late to work" and a possible propositional representation for this sentence (P) might be LATE(Mary, work). The second sentence (S_2) is "She overslept" and a possible propositional representation (Q) for it is OVERSLEPT(X) where X is a variable referring to a female individual. In order to identify that the relation between the sentences is *explanation* the

[3] Interestingly, Kehler's definition of relations in terms of logic distinguished between two different types of *contrast* relations. Because a contrast relation is defined as the negative counterpart of a *parallel* relation there is an ambiguity of what the negation operator applies to – the inferred proposition or its arguments. This distinction was originally described by Hobbs (1985) but is not present in Asher and Lascarides (2003).

reader would have to note, based on world knowledge and/or reasoning, that OVERSLEPT(*X*) normally results in LATE(Mary, work). As Kehler notes, for this inference to take place it is necessary that the reader also infer *X*=Mary.

The definitions used by Kehler explicate his notion of causality as used in his cause-effect relations. Kehler employs the implication operator (→) to denote the relation inferred from world knowledge. Specifically, the effect in a *result* relation is deemed as such because knowledge of its occurrence can be inferred from knowledge of the occurrence of its cause. Moreover, such inference is not required to be based on strict entailment, but rather it should obtain "normally". As such, this operator is more akin to a notion of causality based on conditional probability than one based on the mechanism of physical causes.

This notion of causality accords with the relaxed standards of causes and effects used in everyday language. For instance, Mary's oversleeping in (5) does not directly or necessarily lead to her being late to work (e.g., if she skips breakfast, she might still be at work on time). Nevertheless, it is possible to see how oversleeping is *likely* to lead Mary to be late and consequently, it is relatively straightforward to identify a probabilistic cause-effect relationship.

While the theories differ in how discourse relations are defined, it seems that there is a general agreement about several major types of coherence relations. All theories posit a set of relations that relate discourse segments as a cause and effect, as an elaboration or explanation, and as a list[4]. It might be that the reason theories differ in the specifics of the relation is due to

[4] While Sanders et al. (1992) lump relations such as *occasion* and *parallel* under the broad category of 'list', this might be due to the fact that the representation suggested was intended as an initial exploration of the

differences in the basis for categorization they use for discourse relations (e.g., logical, rhetorical, or property-based), but so far there is little agreement as to what that basis should be.

---

possibility of a property-based categorization of relations and was not intended to present a complete set of discourse relations.

1.2 Empirical Evidence in Support of Coherence Relations

While there is a substantial body of work dedicated to identifying and describing coherence relations, there is relatively little research into their psychological properties and the processes that might be responsible for their computation. In several experiments I (Sagi, 2004) used evidence from a priming experiment to argue that readers compute these relations when they read pairs of sentences. For instance, participants presented with a sentence using a *parallel* relation (e.g., (8)) read it more quickly when it was immediately preceded by a sentence using the same relation (e.g., (9)) than one that used an *explanation* relation (e.g., (10)). Similarly, participants read a sentence using an *explanation* relation more quickly when it was preceded by a sentence using an *explanation* relation than a sentence using a *parallel* relation. Sagi (2006, Experiment 1) replicated this priming effect using pairs of sentences instead of sentences with two clauses.

(8) Humans have two legs and dogs have four legs.

(9) Carpenters make furniture and lumberjacks cut down trees.

(10) People like chocolate because it is sweet.

Similarly, there is evidence that readers find it easier to comprehend sentences that employ more obvious causal relations than less obvious ones. Keenan et al. (1984) and Myers et al. (1987) presented participants with pairs of sentences and manipulated the strength of the causal relation between the first and second sentences in the pair. Their results show that the speed with which participants read the second sentence was correlated with the degree of causal

relatedness of the two sentences. While it is not obvious that this effect is related to the computation of a coherence relation, it seems likely that any causal inferences that the participants in the experiments drew were used to substantiate a causal coherence relation between the two sentences. Specifically, a plausible explanation of the effect is that the participants found it easier to establish the causal link between the sentences. Because Kehler's definition of the *result* relation involves precisely this type of reasoning, this explanation is congruent with readers automatically computing *result* coherence relations when applicable.

Another line of research that provides support for the cognitive utility of coherence relations is concerned with the interaction of coherence relations with pronoun disambiguation. To explore this relation, Wolfe et al. (2004) presented participants with two clause sentences in which the clauses were connected by a connective indicating either a *result* relation or a *parallel* relation. Each of these clauses used a transitive verb and the direct object of the second clause used an unambiguous pronoun to refer to either the subject or direct object of the first clause.

Wolfe et al. found that when a connective indicating a *parallel* relation was used participants were more likely to answer comprehension questions correctly when the pronoun referred to the direct object position (which paralleled the position of the pronoun). In contrast, when a connective indicating a *result* relation was used, participants were more likely to answer comprehension questions correctly when the pronoun referred to the entity in the subject position of the first clause. Similar effects were also found for the reading times of the pronouns. For example, participants who read the *result* clause pair "Fiona defeated Craig and so James congratulated her after the match but nobody took any notice" read the pronoun *her* faster and were more likely to answer the question "Did James congratulate Fiona?" correctly than participants who read the same sentence but with the pronoun *him* rather than *her* (i.e., "Fiona

defeated Craig and so James congratulated him after the match but nobody took any notice"). This effect was reversed when participants were presented with *parallel* clause pairs such as "Fiona complemented Craig, and James congratulated him."

However, as Rohde et al. (2006) note, the versions of the sentences that Wolfe et al.'s participants found more difficult in the *result* condition were also somewhat implausible (e.g., "Fiona defeated Craig and so James congratulated him after the match but nobody took any notice."). While it is possible that the difference observed by Wolfe et al. is due to the low plausibility of the direct object *result* condition, it seems unlikely that this effect is so powerful that it is solely responsible for the interaction in both the reading times and performance on the comprehension question because the opposite effect was found for the *resemblance* relation which does not suffer from this difference in plausibility. Even if the difference in plausibility accounted for all of the difference between the two conditions in the case of a *result* relation, the interaction observed by Wolfe et al. would still obtain.

In contrast with the study reported by Wolfe et al. (2004) which focused on language comprehension, Rohde et al. (2006) focused on language production. Rohde et al. asked participants to complete a text fragment that consisted of a sentence using a transitive verb followed by the ambiguous pronoun "he". There was a striking correlation between the type of coherence relationship that participants chose to use and which character they used the pronoun to refer to. For instance, when participants used an *explanation* relation they overwhelmingly constructed a sentence in which the pronoun referred to the subject of the original sentence (e.g., "Matt passed a sandwich to David. He felt sympathy for David, since he forgot his lunch at home"). In contrast, when the participants employed a *result* relation the ambiguous pronoun

often referred to the object of the first sentence. (e.g., “Matt passed a sandwich to David. He said thanks, and took a bite”).

Both Wolfe et al. (2004) and Rohde et al. (2006) present evidence of a relationship between coherence relations and the process of pronoun disambiguation. However, whereas Wolfe et al. focus on the reading process, Rohde et al. focus on the generative process involved in fragment completion. Moreover, Wolfe et al. identify the coherence relation to the reader through their use of a connective whereas Rohde et al. keep the relation ambiguous and allow their participants to choose it. Wolfe et al. also manipulated the congruence of the referent position between the two sentences whereas due to the limits of their design, Rohde et al. examined completions where the ambiguous pronoun only appears at the head of the sentence. The results from these two complementary studies argue for a link between coherence relations and the process of pronoun disambiguation, a link that exists in both the production and the comprehension phases of language.

Kehler et al. (2008) builds on these two studies and demonstrates the link between coherence relation and pronoun disambiguation even more strongly. In their first experiment, Kehler et al. used a disambiguation task similar to that used by Wolfe et al. while also varying the position of the ambiguous pronoun and the degree of syntactic parallelism. Their results showed the same pattern as those reported by Wolfe et al. even though their materials did not suffer from the plausibility issues Rohde et al. identified. Their next experiment replicated the results reported by Rohde et al. (2006) while their third and final experiment demonstrated how implicit causality biases capture a bias that only exists in the case an *explanation* relation is used. Kehler et al. then hypothesize that different coherence relations might give rise to a different set of biases similar to that of implicit causality. More generally, Kehler et al. (2008) describes how

many of the hypothesized biases in pronoun disambiguation can be explained using an account that focuses on coherence relations.

The evidence presented so far suggests that readers compute coherence relations as they read and that one factor in these computations is the plausibility of a particular relation (e.g., *cause-effect*). Moreover, as Kehler (2002, 2004) and Asher and Lascarides (2003) suggest, there seems to be a tight relationship between the process of computing a coherence relation and the process of pronoun disambiguation.

1.3 Coherence Relations and Bridging Inferences

As discussed in the previous chapter, there is very little empirical data on the processing of coherence relations. Nevertheless, it seems likely that the processes involved coincide with some of the processes that generate more general types of inferences during reading. Many studies have demonstrated that readers compute both predictive and bridging inferences as they progress through a text (e.g., McKoon & Ratcliff[5], 1992; Schmalhofer, McDaniel & Keefe, 2002; Singer et al., 1992).

Of specific interest here are local bridging inferences that help tie sentences together. For example, Singer et al. (1992) suggest that the pair of sentences in (11) requires the bridging inference (12) in order to become a coherence discourse. These inferences seem to hold a similar role to coherence relations as readers are presumed to compute them in "order to establish coherence between a just-read clause and a previously read clause" (Schmalhofer et al., 2002, page 106). Nevertheless, it is important to remember that while these two processes are likely related there are several important factors that differentiate the two types of inference.

---

[5] Interestingly, while McKoon and Ratcliff's Minimalist Hypothesis stresses the role of strategic (or bridging) inferences, it does allow for a restricted class of predictive inferences – "those based on easily available information" (pg. 441). As an example, on page 458 McKoon and Ratcliff suggest that one such predictive inference is "the predictable event of sitting after approaching a chair".

(11)

a. Sharon took the aspirins.

b. Her headache went away.

(12) CAUSE((TAKE, SHARON, ASPIRINS), (GO-AWAY, HEADACHE))

Firstly, bridging inferences are often represented as explicit propositions that are added to the common ground (e.g., Singer et al., 1992). In contrast, coherence relations are generally formulated as rules or schemas that relate propositions to each other (e.g., Asher & Lascarides, 2003; Kehler, 2002, 2004). Secondly, many theories of coherence relations argue that in order for a text to be coherent adjacent discourse segments *must* be connected by a coherence relation (Hobbs, 1985; Mann & Thompson, 1988). On the other hand, bridging inferences are only drawn when there is a gap in coherence. For instance, McNamarra et al. (1996, Experiment 1) reported that in their experimental texts "not a single bridging inference was required." Thirdly, it is generally implied that coherence relations form a singular link between adjacent discourse segments. That is, that one and only one such relation is computed between a pair of sentences. This is made most explicit by Asher and Lascarides' (2003) Nixon Diamond principle which states that when there is a conflict between two possible interpretations of the discourse, neither one is inferred. In contrast, bridging inferences are not required to exist between adjacent discourse segments for the text to be coherent, and it is also possible that more than one such inference may be drawn between a pair of sentences. For instance, in the construction phase of construction-integration (Kintsch, 1988, 1998) the process of drawing inferences does not limit the number of bridging inferences that can be drawn between a specific pair of textual propositions.

Finally, coherence relations are often considered to be a closed set of relations (cf. Kehler, 2002) while bridging inferences are semantically unbounded (cf. Sanders et al., 1992, for an interesting cognitive argument regarding the boundedness of the set of possible coherence relations). This difference is at least partly due to the generic nature of coherence relations. That is, the same coherence relation can pertain in a wide variety of situations (e.g., according to Kehler, 2004, a *result* relation is applicable whenever the event described in the first sentence caused the event described in the second sentence). In contrast, bridging inferences are frequently determined by the specifics of the content (e.g., Singer et al., 1992).

Essentially it appears that coherence relations often correspond to the operator that is at the core of a specific bridging inference (e.g., the *result* relation between the sentences in (11) is explicated in the bridging inference (12) by the use of the CAUSE operator). It is therefore not surprising that many of the formulations of coherence relations (e.g., Asher & Lascarides, 2003; Kehler, 2002, 2004) often entail specific bridging inferences. For instance, coherence relations in Asher and Lascarides' theory are defined based on their entailments. Similarly, Kehler's definition for coherence relations depends upon specific inferences being drawn. However, the inferences required by such coherence relations are often nothing more than generic relations that relate two existing propositions, as in Kehler's definition of *result* which requires two existing predicates (*P* & *Q*) to be in a cause-effect relation. This is similar to the example used by Singer et al. where they suggest that the bridging inference that links the two sentences in (11) is (12).

1.4 Pronoun Disambiguation

In the previous three chapters I described the notion of *coherence relations* and some empirical evidence supporting their cognitive reality. However, in order to evaluate the effect causal coherence relations might have on pronoun disambiguation, it is important to also understand how pronoun disambiguation proceeds more generally. In this chapter I provide a brief description of the history of psycholinguistic research into pronoun disambiguation. The next chapter then focuses on the empirical evidence regarding the influence of causal information on the disambiguation of pronouns.

The disambiguation of pronouns is complex and has a long history in linguistics and psycholinguists. Many different algorithms and processes have been proposed for this problem. For instance, Hobbs (1978) compares a simple model based on syntax parsing with a more complex, semantically-grounded model. While his semantic model is overall more successful than the simple syntactic one, the syntactic model performs remarkably well (over 80%).

In the psycholinguistics community, one factor that has often been argued to influence pronoun disambiguation is that of linguistic focus or markedness (e.g., Fletcher, 1984). The linguistic focus is frequently on the subject of the previous sentence. However, syntactic and semantic manipulations, such as using a passive construction, might affect this focus. This notion is embodied by theories such as centering theory (Grosz, et al., 1983, 1995). Gordon and Scearce (1995) use centering theory as a basis for arguing that the disambiguation of pronouns is primarily based on discourse processes and that world knowledge plays a secondary role in the process. Their primary finding is that when participants encounter an ambiguous pronoun they read the related verb phrase more quickly when it resolves the ambiguity in congruence with

centering theory and that the same expectations are present when the verb phrase is read first – if the verb phrase indicates a focus shift participants read the verb phrase faster when a name was used rather than a pronoun while when the verb indicated no focus shift participants read the verb phrase more quickly when a pronoun was used. For example, they presented the readers with passages such as (13) below. Critically, there were 4 versions of (13c) that varied the verb ("sent" or "received") and whether the pronoun *he* was used or the name of the actor was used (Bill or John, depending on the verb). Participants read the verb phrase of (13c) faster when *he* was used rather than Bill if the verb was "sent" and when John was used rather than *he* if the verb was "received".

(13)

a. Bill wanted John to look over some important papers.
b. He had to mail him the documents by Monday.
c. Unfortunately, he never sent the papers.
d. As a result, the whole deal fell behind schedule.

While the disambiguation of pronouns depends on many factors, it seems that a small number of simple heuristics and syntactic formulations such as those used by centering theory can account for a substantial proportion of them. The most notable of these heuristics seems to be that a pronoun is most likely to refer to the current focus of the discourse.

Nevertheless, as described in the previous chapter, results such as those described in Wolfe et al. (2004) and Rohde et al. (2006) imply that information about coherence relations is also an important factor. Moreover, these results suggest that while simple heuristics might be

useful, they describe some general biases in the process that are often overridden by more specific information that is provided in the text (e.g., the connective used by Wolfe et al.).

1.5 Causality and Pronoun Disambiguation

In the previous chapter I described factors that affect pronoun disambiguation as a general process. However, when it comes to the specific question of interest for this dissertation, the interaction of causal coherence and pronoun disambiguation, there is a body of research that is more directly relevant and explores the effect of specific types of causal information on pronoun disambiguation. Specifically, there is a well established link between the causal implications of verbs and the process of pronoun disambiguation (e.g., Brown & Fish, 1983; Au, 1986; Stewart, Pickering & Sanford, 2000; and see Rudolph and Forsterling, 1997, for an extensive review). The literature focuses on the implicit effects that a verb has on the disambiguation of a pronoun in the following sentence. For instance, when Garvey et al. (1976) presented participants with sentences like "George telephoned Walter because he wanted sympathy", their participants were more likely to attribute the ambiguous pronoun *he* to George than to Walter. Conversely, when the verb "telephoned" was replaced with the verb "criticized" (i.e., "George criticized Walter because he wanted sympathy"), participants were more likely to attribute the pronoun to Walter than George.

In the literature, such differences are often considered to be the result of the semantic roles associated with the verb (e.g., Brown & Fish, 1983; Au, 1986). In the case of transitive state-denoting verbs, the cause of a sentence is often considered to be the person in the *stimulus* role (such as *like* in which the object is the stimulus and the subject is the experiencer and *disappoint* in which the opposite is true). For example, when completing a sentence fragment such as "Mary likes Lisa because she …" participants are more likely to assume that *she* refers to Lisa and complete the sentence accordingly (e.g., "She smiles a lot"). In contrast, in action verbs

the *agent* is generally preferred to the *patient* as the referent of an ambiguous pronoun (e.g., readers will prefer to attribute the cause of the event described by the verb *cheat* to be the subject than the object). For example, when completing a sentence fragment such as "Mary cheated Lisa because she…" participants are more likely to assume that *she* refers to Mary and complete the sentence accordingly (e.g., "She took more than her fair share").

Rudolph and Forsterling suggest an additional distinction within action verbs between verbs with *agent-patient* roles and those with *agent-evocator* roles. They note that when an *evocator* role is present, that role is often the preferred cause of the event (e.g., events described using the verb *praise* are frequently considered to be caused by the object rather than the subject).

While there is a substantial body of literature associated with the implicit attribution of causes, there is very little research into the consequences of events. In one major study Au (1986) found that when participants completed sentence fragments such as "Mary admired John so …" they were more likely to refer to the person filling the *experiencer* role (Mary in this case). When completing fragments that involved an *agent* and a *patient*, such as "Mary criticized John so …," participants were more likely to refer to the person filling the *patient* role (John in this case). Both of these relations held regardless of the grammatical role of the individuals (i.e., whether a passive or active construction was used). When completing sentence fragments Stewart et al. (1998) found similar biases in completions involving the consequence of verbs when they used the connective 'because' at the beginning of the sentence to indicate that the first clause was the cause of the second (e.g., "Because John loves Mary …").

Crinean and Garnham (2006) examined the completion biases reported by Stewart et al. and noted that for verbs in which the roles were *agent* and *evocator*, the *evocator* was the more

likely referent of the pronoun for both the cause and the consequence. This result is at odds with the pattern of reversal between causes and consequences that Au identified for sentences involving *agent-patient* and *experiencer-stimulus* roles. However, Pickering and Majid (2007) note some issues with the Crinean and Garnham's analysis. Most notably, they argue that the evocator role is "only called the Evocator because people tend to provide a reason for the event that is framed in terms of it rather than in terms of the agent" (page 785). Pickering and Majid further argue that the role of *evocator* has no basis in semantics.

Essentially, Pickering and Majid argue that Crinean and Garnham's use of the *evocator* role is circular as it is based on the data they use to explain. Nevertheless, it is important to note that Crinean and Garnham appear to be aware of the contentious nature of the *evocator* role they propose as they suggest that the verbs they identify "also possess the properties of psychological state verbs". They seem to be suggesting the *evocator* role as a new semantic role that is based on specific psychological properties of the verbs. Nevertheless, the inconsistency noted by Pickering and Majid (e.g., the inclusion of the verbs 'complimented' but not of 'flattered' or 'apologized') does seem to detract from this proposal.

Overall, it is obvious that verbs carry some information regarding their causes and their consequences. During the process of reading, this information is used to guide the reader's interpretation of the text. However, in order for this information to be useful, the reader must first determine the coherence relation that obtains between the previous sentence and the current one.

As Au (1986) and Crinean and Garnham (2006) note, the pattern of pronoun disambiguation for consequences is different from that of causes. Furthermore, as Rohde (2006) notes, participants vary their choice of referent for their sentence completions based on the

coherence relation they use. This result points to an interdependence between pronoun disambiguation and coherence relation. Therefore, determining the coherence relation is likely a prerequisite to identifying the appropriate referent to the pronoun. However, experiments involving implicit causality and consequentiality generally use explicit connectives. The use of an explicit connective essentially constrains the participants to use a specific coherence relation in their interpretation (e.g., the use of the connective *because* would identify the sentence pair as involving an *explanation* relation). As a result, the issue of the interaction between coherence relations and verb causality biases has not yet been fully addressed by this literature.

1.6 Processing Coherence Relations

Now that I have described in some detail the notion of coherence relation and provided some details on the process of pronoun disambiguation I turn to presenting a theory that tries to explain how the two interact during reading. I should note that while in this dissertation I restrict the discussion to the case of causal coherence, I believe the theory outlined below can be easily generalized to apply to other types of relations (e.g., *resemblance* relations).

Kehler's definition for the group of causal relations relies on the existence of two propositions *P* and *Q* that are causally related to one another. For instance, in the case of a *result* relation *P* normally causes *Q* whereas the reverse is true when the relevant relation is *explanation*. Therefore in order for readers to apply this relation they need to be able to identify this relation and determine that it is the most relevant relation between the two sentences.

However, such an evaluation depends on the specific interpretation of the two sentences. As Kehler (2002) notes, the relevant coherence relation can change depending on the chosen interpretation of the linguistic utterances involved. Most notably, this is the case for ambiguous pronoun references. For example it is possible to interpret (14) in at least two ways, depending on the resolution of the ambiguous pronoun "he":

- A *result* relation:
    - John was rushing to get Mark to the hospital *and therefore* John got into an accident.
- An *explanation* relation:
    - John was rushing to get Mark to the hospital *because* Mark got into an accident.

(14)

a. John raced to get Mark to the hospital.

b. He was badly wounded in an accident.

Consequently, readers need to consider different possible interpretations of the utterances (varying both the pronoun referent and the coherence relation independently) and determine which ones are coherent and which of the coherent ones is the most relevant. Even if readers often rely on a default relation whenever plausible, they still need to identify the interpretation that best fits with that relation. Moreover, readers cannot know in advance which interpretation is more likely to yield an appropriate relation. Once the different interpretations are evaluated readers will then need to identify which of these is the most likely.

This process therefore likely involves two stages – an evaluation stage and a decision stage. In this it is akin to the *causal continuer model* proposed by Rips et al. (2006) to account for tracking the identity of objects over time. According to the causal continuer model, when trying to identify which of several objects is the temporal continuer of a previously encountered object, it is necessary to first identify how close each possible continuer is to the original object (in Rips et al.'s model this closeness is based on causal relatedness). Once this evaluation is completed, the continuer is chosen as the object that is closest to the original, assuming that it is close enough. However, if none of the objects are close enough (based on a threshold of causal relatedness) then none of the objects are considered to be the continuer of the original object.

It seems reasonable to hypothesize that a similar mechanism governs the decision between alternative interpretations and coherence relations. One possibility is that readers first evaluate the coherence provided by the possible relations and the bridging inferences they

require. They then choose the relation that provides the most coherence. However, in some cases none of the evaluated relations will provide a sufficient level of coherence. In that case no relation is drawn and the text is judged incoherent. Alternatively, motivated readers might then backtrack or attempt to draw more complex inferences which might enable them to make the text coherent. This process fits well with models, such as the Landscape model (van den Broek, et al., 1999), that assume that readers attempt to maintain an "adequate" level of coherence in a text as they read.

Of course, the outline I provided here defines the general properties of a cognitive model but does not specify many of the details required to implement such a model. For the purposes of this dissertation this level of detail is sufficient in that it constrains the processing of coherence and makes testable predictions with regards to the relationship between coherence relations and pronoun disambiguation. Specifically, it makes the claim that coherence-based interpretations are one avenue that readers use to determine the referent of pronouns and that readers choose among the interpretations based on their plausibility. A complete model would require a much greater degree of specificity about many of the details I glossed over, such as how the interpretations are generated and evaluated.

The experiments that follow test several predictions that arise out of this hypothesis regarding the nature of the cognitive apparatus. The experiments focus on the *result* relation because that relation is defined in all of the taxonomies of coherence relations in a fairly uniform manner. Moreover, this relation is based on causal information which is relatively easier to quantify and manipulate.

These experiments are summarized in Table 1.

Experiment 1 examines how different causal inferences influence the disambiguation of a pronoun. Experiment 2 extends this by exploring whether causal information plays a more important role in causal relations than in relations involving temporal contiguity. Experiment 3 looks at the effect of negation in a cause-effect relation on the interpretation of an ambiguous pronoun. Specifically, it tests the hypothesis that negating an effect results in the readers assuming a *violated expectation* relation rather than a relation that is not causal in its nature (e.g., *temporal*). The latter might be expected because an event is unlikely to be perceived as the cause of the absence of its normal effect. Finally, Experiment 4 looks at the effect of ambiguity on the process of choosing the appropriate coherence relation.

While prior research provides ample demonstration for the computation of general bridging inferences, these inferences were generally considered as specialized inferences that are tailored to the specific context. In contrast, theories of coherence relations propose that many of the inferences readers draw as they read are selected from a small list of pre-existing relations. This difference suggests a processing difference – readers need only *select* a relation from a list of coherence relations rather than generate a brand new inference based on the context. The experiments proposed here explore some of the processing implications of this selection process.

However, it is important to note that the model I outlined above is only one of several possible models of the process readers might use to identify coherence relations. One assumption that is inherent in the simple model is that readers first construct the interpretations and then choose among them, there are other models that could produce the same effects. However, it is possible that readers construct interpretations and compare them in parallel. This type of parallel processing would be more complex but might be reasonable given the apparently parallel nature of the brain.

Another important assumption in the model above is that all possible interpretations are computed. Due to cognitive constraints that is unlikely to be the case. It is far more likely that only a small subset of interpretations are computed and compared. If we assume that only a small subset of interpretations are computed at first, their failure might lead to more interpretations being computed and compared in an iterative process. Such a model would require an additional mechanism that determines which interpretations should be computed at any given point and allows for a greater degree of cognitive flexibility.

However, the focus of this dissertation is on causal information which is a common and highly salient source of coherence in discourse. The results of the experiments below would therefore be unlikely to distinguish a more complex iterative or parallel model from the simpler mode I outlined. As such, I compare the results with the performance of this simple model with the expectation that such a comparison would be equally valid regardless of the specifics of the model. I leave the task of identifying the specifics of the model as an open question for future studies.

2 The Experiments

Table 1 – Summary of Experiments

Design	Independent variable(s)	Dependent variable(s)	Prediction
Experiment 1	Causal relatedness and pronoun disambiguation in the *result* relation		
Pronoun disambiguation	Causal relatedness between head and tail sentences	Pronoun disambiguation (Subject/Object), Confidence in response	Pronoun disambiguation will show a sharp *S* curve based on causal relatedness. Confidence will be a linear function of causal relatedness
Experiment 2	The limits of causality		
Pronoun disambiguation for causal vs. temporal relations	Causal relatedness, connective (causal/temporal)	Pronoun disambiguation (Subject/Object), Confidence in response	Causal connective will replicate experiment 1. Temporal connective will show a decrease in levels of confidence because causal information is less relevant.

Experiment 3	Causality and Negation		
Pronoun disambiguation for violated expectation relations	Causal relatedness	Pronoun disambiguation (Subject/Object), Confidence in response	Following Kehler, should replicate Experiment 1. Alternatively, if the most causal interpretation is chosen effect should reverse (because causality is negated) or disappear (because causality is no longer relevant).
Experiment 4	Time course of causal coherence		
Word reading times for result and explanation relations	Ambiguity of pronoun, Coherence relation (result/explanation), Target reference of pronoun	Verb and Verb Complement reading times	A 3-way interaction between the independent variables. Fastest reading times for unambiguous pronouns that indicate the subject for explanation relations and object for result relations.

2.1 Experiment 1 – Causal Likelihood and Pronoun Disambiguation

Experiment 1 looks into the relationship between causal relatedness and pronoun disambiguation in the context of a *result* coherence relation. Specifically, I hypothesize that when a "good enough" relation is found, the inferences that follow from it will also be drawn. This will result in a specific interpretation of semantic ambiguities. For instance, an ambiguous pronoun will be resolved. In contrast, when no interpretation yields a "good enough" relation the disambiguation of an ambiguous pronoun would be at chance.

Experiment 1 tested this hypothesis by examining the relationship between ratings of causality and pronoun disambiguation in the *result* relation. Participants in this experiment were presented with pairs of sentences that varied in the degree to which the second sentence is a likely effect of the first. In each pair, the first sentence mentioned two people and the second sentence used an ambiguous pronoun to refer to one of them (e.g., "John accused Mark of stealing a car. He called the police", see Table 2 below for more examples). Following Rips et al.'s (2006) causal continuer model, I hypothesized that when the causal relatedness of one interpretation is high enough, the disambiguation will follow the inference drawn from that interpretation. However, when the two sentences are not causally related the pronoun disambiguation will be at chance.

2.1.1 Predictions

Following the hypothesis, I predicted that the proportion of disambiguation would show an interaction between the level of causal relatedness and the preferred interpretation of the pronoun (based on the high causality condition sentence pair). While this interaction is fairly straightforward and intuitive, it is important to validate it as it forms the basis for the experiments that follow. Moreover, while there is abundant evidence for biases in the interpretation of pronouns, there is virtually no empirical evidence supporting the strong intuition that pronoun disambiguation is sensitive to degree of causal relatedness.

2.1.2 Design

Experiment 1 had the form of a 2x2 factorial.

The dependent variable was the proportion of object choices in the disambiguation of the ambiguous pronouns.

The independent variables were the causal relatedness of the two sentences (high vs. low) and the preferred interpretation of the pronoun (subject vs. object).

2.1.3 Materials

The materials for this experiment consisted of 16 initial (*head*) sentences and 32 final (*tail*) sentences. Each head sentence describes an event involving 2 named human entities, a subject and an object. For instance, the head sentences given in Table 2 introduce John (the subject) and Mark (the object). Each tail sentence (e.g., ‘He called the police.’ from Table 2) describes a second event but refers to only one agentive entity and only using the appropriate initial pronoun (*he* or *she*). These sentences are divided into 16 groups each consisting of a head

sentence and 2 tail sentences. Appendix A provides a complete list of all the sentences used in this experiment.

Within each group one of the tail sentences forms a highly likely causal continuation of the head sentence when the pronoun is resolved to refer to the *subject* of the initial sentence (but not the object). Similarly, the other tail sentence is a highly likely causal continuation of the head sentence when the ambiguous pronoun is resolved as referring to the *object* of the initial sentence (but not the subject). For instance, if John accused Mark of stealing a car, it is much more likely that John would call the police than Mark. In contrast, following the accusation it is much more likely for Mark to call a defense attorney than for John to do so.

These groups are further paired into 8 sets such that each set consists of 2 groups. The tail sentences of one group are unlikely to be causal continuations of the head sentence of the other group. For instance, accusing or being accused is unlikely to cause someone to take out a credit card.

Consequently there are 3 possible types of head and tail sentence pairing:

1. Head + Subject Tail
2. Head + Object Tail
3. Head + Low Causal Likelihood Tail (Head taken from one group, tail from the other)

2.1.3.1 Preliminary Causality Rating Study

To validate these pairings, a group of participants was asked to rate the likelihood that the event described in the head would cause the event described in the tail sentence. For this rating task the subject of the tail sentences was described unambiguously using a name rather than a

pronoun. Participants were therefore asked to rate the causal likelihood of each tail sentence twice, once with the subject of the head sentence as the subject of the tail sentence and once with the object of the head sentence as the subject of the tail sentence. A sample set of materials is given in Table 2.

Table 2 – Sample materials for Experiment 1

	Head	Tail type	Tail
Group A	John accused Mark of stealing an expensive car.	Subject continuation	He called the police.
		Object continuation	He called a defense attorney.
Group B	John sold Mark a brand new suit.	Subject continuation	He asked for payment.
		Object continuation	He took out a credit card.

The questions participants were asked to answer when rating the causal likelihood of the sentence pair for Group A's subject continuation were as follows[6]:

- If John accused Mark of stealing a car, how likely is that to cause John to call the police?

  1 2 3 4 5 6 7 8 9

- If John accused Mark of stealing a car, how likely is that to cause Mark to call the police?

  1 2 3 4 5 6 7 8 9

2.1.4 Procedure

Participants were presented with a paper packet containing the 16 sentence pairs in one of 4 predetermined random orders. Each participant read 4 sentence pairs of each of the 4 possible combinations of head and tail (e.g., the group A head sentence in Table 2 can be paired with the group A subject continuation, the group A object continuation, the group B object continuation, or the group B subject continuation). Furthermore, there were 4 counterbalancing conditions such that each head is associated with each possible tail in one of the conditions. This resulted in 16 possible combinations of order and counterbalancing conditions. Each participant was randomly assigned to one of these combinations.

After reading each sentence participants were asked to circle the name referred to by the pronoun *he* in the second sentence. The two names appeared immediately below the sentence pair with the name of the subject on the left and the object on the right (the same order as they appear in the sentence). After making their choice participants were asked to indicate their level of confidence in their response on a 9-point likert scale.

[6] Obviously, when presented to participants these two questions will not appear one after the other.

2.1.5 Participants

Forty-one Northwestern undergraduate students participated in this experiment. Sixteen participants provided causal ratings and 25 participants completed the pronoun disambiguation task.

2.1.6 Results

One participant chose the subject disambiguation for all 16 sentences and this data was therefore discarded.

The mean causal rating for the subject and object interpretations of each head-tail combination was computed for each participant. Similarly, the proportion of object responses for each head-tail sentence combination was also calculated for each participant. The overall means for each condition are provided in Table 3.

Table 3 – Mean ratings and disambiguation proportions in Experiment 1 (standard deviations given in parenthesis)

Causal Likelihood	Tail type	Mean Causal Likelihood Rating		% of Object Disambiguations	Mean Confidence
		Subject	Object		
High	Subject continuation	6.10 (0.76)	1.90 (1.52)	0.01 (0.05)	7.68 (1.11)
	Object continuation	2.85 (1.66)	6.39 (1.07)	0.97 (0.08)	7.65 (1.08)
Low	Subject continuation	1.69 (1.01)	1.59 (0.96)	0.30 (0.26)	4.36 (1.69)
	Object continuation	1.58 (1.13)	2.13 (1.13)	0.53 (0.27)	4.20 (1.86)

2.1.6.1 Ratings of Causal Likelihood

As expected, the ratings of causal likelihood validated the assignment of sentence pairs to condition. Sentence pairs in the high causality condition were rated higher on average than sentences in the low causality condition. Likewise, in the high causality condition sentence pairs involving a subject continuation had higher ratings for the subject interpretation than the object interpretation while sentence pairs involving an object continuation had higher ratings for the object interpretation than the subject interpretation. These differences in rating are greatly reduced in the low causality condition.

These results were confirmed using separate analyses of variance with participants (F_1) and sentence pairs (F_2) as random effects. Participants rated high causality pairs higher than low causality pairs ($F_1(1, 15) = 209.68$, $MSE = 1.00$; $F_2(1, 15) = 86.42$, $MSE = 2.42$, $p < .001$ in both cases). The analyses also showed the predicted triple interaction. Participants rated the causal likelihood of sentence pairs higher for high causality sentence pairs when the pronoun interpretations matched the tail type. In contrast, for low causality sentence pairs ratings were uniformly low and the pronoun interpretation made little difference ($F_1(1,15) = 53.66$, $MSE = 1.88$; $F_2(1, 15) = 107.52$, $MSE = .94$, $p < .001$ in both cases).

There was no main effect of the pronoun interpretation used ($F < 1$ in both cases). However, there was a significant main effect of tail type. Participants rated the causal likelihood of sentence pairs higher when an object continuation was used than when a subject continuation was used ($F_1(1, 15) = 10.53$, $MSE = .54$, $p < .01$, $\eta^2_p = .41$; $F_2(1, 15) = 4.57$, $MSE = 1.25$, $p = .049$, $\eta^2_p = .23$). That is, participants rated the causal likelihood of sentence pairs involving “He called a defense attorney” higher overall than those involving “He called the police”. This effect was not predicted and I have no explanation for it.

2.1.6.2 Disambiguations

The pattern of pronoun disambiguation followed my predictions and agreed with the causal likelihood ratings. As shown in Table 3, when presented with a sentence pair involving high causal likelihood, participants identified the subject as the referent of the pronoun *he* with a high degree of confidence when a subject continuation was used (e.g., “John accused Mark of stealing an expensive car. He called the police.”). In contrast, when an object continuation was

used as part of a high causal likelihood sentence pair, participants identified the object as the referent of the pronoun *he* with a similarly high degree of confidence (e.g., “John accused Mark of stealing an expensive car. He called a defense attorney.”). This pattern was less evident and confidence was likewise lower in the case of sentence pairs involving low causal likelihood (e.g., “John sold Mark a brand new suit. He called the police.”).

As above, separate analyses of variance for participants and sentence pairs were used to validate the observed trends. These analyses revealed the hypothesized interaction between the level of causality (high or low) and the type of continuation (subject or object). Participants were more likely to disambiguate the sentence pairs based on the type of continuation in the high causality condition than the low causality condition ($F_1(1, 23) = 124.12$, $MSE = 0.026$; $F_2(1, 15) = 74.96$, $MSE = 0.028$, $p < .001$ in both cases). Because of the difference in the high causality condition, participants were also likely to overall disambiguate the sentence pairs based on their type of continuation ($F_1(1, 23) = 256.79$, $MSE = 0.033$; $F_2(1, 15) = 212.14$, $MSE = 0.027$, $p < .001$ for both). However, there was no overall effect of causality ($F_1(1, 23) = 2.85$, $MSE = 0.045$, $p = .11$; $F_2(1, 15) = 2.25$, $MSE = 0.098$, $p = .15$). These results are consistent with the hypothesis that participants are likely to use evidence for a strong causal link between two events when determining the appropriate referent for an ambiguous pronoun.

In the model outlined earlier, I proposed that pronoun interpretation followed from the choice of coherence relation in a 2-step process. The first step involves the evaluation of the coherence of different possible interpretations while the second step involves choosing the most relevant relation. Following Rips et al. (2006)’s causal continuer model I hypothesized that this information would only come into play when a certain threshold of coherence is crossed.

Because in this experiment participants are faced with a choice between two competing causal interpretations, a likely model of their choices can be generated using Luce's choice axiom (Luce, 1977). Specifically the likelihood that participants will choose one interpretation over the other can be estimated as $Pr(Obj) = \frac{CausalRating_{Obj}}{CausalRating_{Subj} + CausalRating_{Obj}}$ [7]. The likelihood that at least one of the interpretations is beyond threshold of causal relatedness will be represented by the greater of the two causal relatedness ratings for the sentence: $T = max(CausalRating_{Subj}, CausalRating_{Obj})$. In its simplest form, the model above translates to a multiple regression model with *Pr*(*Obj*) and T as independent variables. The model predicts that when the causal relatedness (T) is high participants choice will follow *Pr*(*Obj*) whereas when it is low participants choice will not depend on it.[8]

In the experiment, participants indicated the referent of the pronoun of the tail sentence as either the subject or the object of the head sentence. However, as the variation in confidence ratings indicate, their certainty in their choice varied. It is likely that in repeated trials their choices of referent would show some variation, and that this variation would be proportional to

[7] This model assumes that the likelihood of interpretations based on information other than causes and effects is minimal. However, if we assume that causal relations are preferred over other types of relations then once the likelihood of a causal interpretation crosses the relatedness threshold this assumption is likely to hold. Moreover, in this experiment there was little information other than causal information available for participants to base their inferences on.

[8] While it is possible that these two variables are related in a nonlinear fashion, a linear regression provides a first order approximation of such a relation. As such, this fit was chosen because it involves the least number of assumptions.

their confidence in their response. To account for this variation, the independent variable for the regression weighted the actual choice made by the participant by their confidence in that choice. Specifically, their choice of referent was assigned a value of 1 or -1 depending on their choice of referent, object or subject, respectively. This value was then multiplied by the participants' indicated confidence (ranging from 1 indicating very low confidence to 9 indicating maximum confidence). This resulted in each choice being represented by a number between -9 (high confidence in a subject choice) to 9 (high confidence in an object choice). Participants' responses were averaged for each sentence pair, resulting in 64 data points that were used in the regression model.

This regression model was statistically significant ($F(3, 60) = 58.35$, $MSE = 8.93$, $p < .001$) and explained 74% of the variance in the weighted proportion of object disambiguations. More importantly, when causal relatedness was high participants' choice of referent was more likely to follow the prediction of *Pr*(*Obj*) than when causal relatedness was low ($b=.33$, $t(60) = 4.46$, $p < .001$). Pr(Obj) was a significant overall predictor in this regression ($b=.64$, $t(60) = 8.58$, $p < .001$) but T was not useful as an independent predictor of participants' choices ($b=.05$, $t(60) = .42$, *n.s.*).

As predicted by the model, the likelihood of object interpretation was more predictive when the maximum causal relatedness was high than when it was low. This is best demonstrated by examining the correlation between the weighted proportion of object disambiguations and the predicted probability of object choices for different ranges of *T*. Table 4 presents the correlation coefficients between the likelihood of object interpretation and *Pr(Obj)* for each quartile (16 sentence pairs) of the data. For high values of *T* (the 3rd and 4th quartiles), the correlation

coefficient is very high (.94). The correlation coefficient drops precipitously for the 2nd quartile and is non-significant when the value of *T* is low (the 1st quartile).

Table 4 – Correlation between the measured likelihood of object interpretation and the predicted probability of object choice (*Pr(Obj)*) by quartile of the maximum causal likelihood (*T*). (* indicates a statistically significant correlation coefficient, + indicates a marginally significant correlation coefficient)

Quartile	Lower bound of quartile (value of *T*)	Upper bound of quartile (value of *T*)	Correlation with *Pr(Obj)*
1st	0	1.77	.21
2nd	1.77	4.72	.43+
3rd	4.72	6.41	.94*
4th	6.41	9	.94*

The results from this experiment provide empirical support for the intuition that pronoun disambiguation is sensitive to the degree of causal relatedness between sentences. When the two sentences presented had a high degree of causal relatedness participants were almost uniform in their choice of referent for the ambiguous pronoun. In contrast, when the degree of causal related between the sentences was low, participants' choice of referent varied.

Moreover, these results can be explained using the model of pronoun disambiguation I proposed which is based on two variables: A threshold of coherence and the likelihood of choice of one interpretation over another. If the degree of relatedness between the sentences is beyond the threshold of coherence, readers base their choice of referent on the information provided by coherence. However, if the degree of relatedness between the sentences does not cross the threshold, readers are less likely to use the degree of causal relatedness to inform their choice of referent for an ambiguous pronoun.

This experiment tested a model in which causal information plays a pivotal role in readers' choice of a referent for ambiguous pronouns. The results of this experiment demonstrate that when a pair of sentences describes events that have a high degree of causal relatedness, that causal relatedness affects the referent readers choose for a pronoun. However, when the pair of sentences had a low degree of causal relatedness these effects were greatly reduced. Experiment 2 will test a related prediction of this model – That causal information is more relevant when the relationship between the sentences is causal than when it is simply temporal.

2.2 Experiment 2 – The Limits of Causality

In Experiment 1 participants were not provided with any information regarding the nature of the relationship between the two sentences. As a result they were free to use whatever information they could deduce from their world knowledge. In the case of the materials presented, this often led them to identify a cause-effect relationship between the sentences in a pair. However, the *result* relation is only one possible coherence relation and other coherence relations rely on different types of information. The usefulness of information such as cause-effect relationship depends on the specific coherence relation. For instance, if a connective identifies the relevant relation as *explanation*, participants should prefer interpretations in which the 1st sentence is the effect of the 2nd sentence rather than its cause (as is the case with the *result* relation). That is, causal information should be relevant for causal relations more so than for other relations.

Experiment 2 aims to test this hypothesis by manipulating the connective used to bind a pair of clauses into a single sentence. The connectives used varied on whether they explicitly imply a causal relationship between the events described and on the temporal order in which the two events described in the clauses occurred (i.e., which event occurred first). Consequently, four connectives were used (see Table 5).

The clauses had the same structure as the sentences used in Experiment 1 – The first clause (the head) introduced two people as the subject and object of a described event while the second (the tail) began with the ambiguous pronoun “he”. However, unlike Experiment 1, the type of ambiguity resolution was manipulated by changing the head sentence. Table 6 provides a sample material set from Experiment 2.

In one head type condition, the *result* connective was congruent with the interpretation of the subject as the referent of the pronoun and the *explanation* connective was congruent with the interpretation of the object as the referent (e.g., "John asked Mark for some study tips and so he aced the test" → he = "John" vs. "John asked Mark for some study tips because he aced the test" → he = 'Mark'). Conversely, in the other condition the *result* connective was congruent with the interpretation of the object as the referent of the pronoun and the *explanation* connective was congruent with the interpretation of the subject as the referent (e.g., "John tutored Mark for several hours and so he aced the test." → he = "Mark" vs. "John tutored Mark for several hours because he aced the test" → he = "John").

Table 5 – The connectives used in Experiment 2

		Causal Relation Type (Temporal Order)	
		Result (forward)	Explanation (backward)
Causal implication	Explicit	and so	because
	Implicit	before	after

The prediction of this experiment follows directly from the type of connective used – when the connective explicitly involves causality ("and so", "because") participants should be more likely to rely on causal information and more confident about their use of such information than when a temporal connective is used ("before", "after"). Specifically, a causal relation

implies a temporal relation and is therefore more constrained than pure temporal information. Following the Gricean maxims (Grice, 1975), using a temporal connective rather than a causal connective could be taken to indicate that the relation is decidedly *not* causal. As a result, when a temporal connective is used, both the proportion of object choices should be closer to chance (further away from uniform agreement about the interpretation which would be indicated by either 100% object choices or 100% subject choices) and the confidence of participants in their interpretations should decrease. Following Kehler, I will use the term *occasion* to refer to these temporal coherence relations.

Table 6 – Sample materials for Experiment 2

Head Type	Head	Tail
Result: Subject preferred *Explanation*: Object preferred	John asked Mark for some study tips	he aced the test.
Result: Object preferred *Explanation*: Subject preferred	John tutored Mark for several hours	he got an A on the midterm.

2.2.1 Predictions

From the hypothesis that causal information is more relevant to coherence relations that explicitly rely on it, such as *result* and *explanation*, than other relations (e.g., *occasion*), it follows that participants will be more inclined to use relevant causal information when a connective that indicates a *result* and *explanation* relation is present than when a connective

associated with an *occasion* relation is present. Consequently, causal information should influence participants' disambiguation of the ambiguous pronouns more when a causal connective is present. Likewise, participants should be more confident in their responses when a causal connective is used.

More specifically, according to Kehler connectives provide readers with the coherence relation that they should employ to understand the relationship between the sentences. Because different coherence relations are based on different types of information, the presence of a connective would therefore make a reader more likely to use sources of information that are congruent with that relation and less likely to rely on source of information that are less congruent with the relation. Connectives that indicate a causal relation should result in an increased reliance on causal information presented whereas connectives that indicate a relation that is not necessarily causal (e.g., *occasion*) are likely to induce participants to rely on sources of information other than causes and effects.

This is in line with the Gricean implications of using a connective that indicates *occasion* in lieu of a more constraining causal connective. Under the Gricean account the implications of such a choice of connective is that the more constraining relation (causal) is not applicable in this situation either because such a relation it is not true, or the speaker is not sure whether it is true. In this case, both the semantics of coherence relations and the pragmatics of it would make the same prediction – a reduced reliance on causal information when a temporal connective is used. Nevertheless, while performance on the disambiguation task should be more variable for *occasion* (i.e., the proportion of object disambiguations should be closer to chance), it is likely that the causal information present in the sentences will still result in significant disambiguation biases.

Similarly, the shift from causal connectives to temporal connectives should show a decrease in the level of confidence participants have in their interpretations. Because causal information represents only one type of information relevant to an *occasion* relation there is more likelihood that other sources of information, such as probabilistic biases, would also be judged relevant to the disambiguation of a temporal relation.

Participants are also more likely to focus on causal information when a causal connective is used. Moreover, when a sentence pair uses a temporal connective participants might take more time to decide whether the causal information they find is relevant. Both of these factors should result in a decreased reading time for sentence pairs using a causal connective compared to pairs that use a temporal connective.

Consequently, the type of head should have a stronger effect on the proportion of object choices when a causal connective is used than when a temporal connective is used. In addition, participants' confidence in their responses should be higher when the connective is causal than when it is temporal. However, confidence should not vary based on the type of head because a similar degree of causal relatedness is provided by both types. Finally, participants should read the sentence pairs that use a causal connective faster than those that use a temporal connective.

In terms of the model I outlined earlier, the interpretations of an occasion relation would likely gain support from additional sources of information while at the same time reducing the weight given to causal information. This would result in an increased amount of "noise" in the

comparison process[9]. This makes the decision process for sentence pairs connected via temporal connectives more difficult resulting in a decrease in the confidence levels and an increase in reading time compared to sentence pairs that use causal connectives. If the alternative interpretations are close to each other, it might also result in the choice of interpretation (i.e., disambiguation) being more variable.

2.2.2 Design

Experiment 2 had the form of a 2x2x2 factorial.

The main dependent variable was the proportion of object choices in the disambiguation of the tail sentence pronouns. Secondary dependent variables were the confidence participants have in their choice and their reading times (measured for the entire sentence pair and analyzed as reading time per syllable).

The independent variables are the type of head (subject-preferred vs. object-preferred for a *result* relation), type of connective (causal vs. temporal), and the temporal sequence of events (forward vs. backward). The different types of head are illustrated in Table 6 while the type of connective and temporal sequence are illustrated in Table 5 using the rows and columns respectively. Importantly, the temporal sequence of events depends on the coherence relation when a causal connective is used (*result* is a *forward* temporal relation while *explanation* is a

---

[9] These additional sources of information themselves are, of course, not noise. However, because they are likely to influence all of the temporal interpretations under considerations relatively equally, they might reduce the relative "mental distance" between the alternatives and consequently increase the influence of random noise in the stochastic decision function. Likewise, the reduced weighting of the causal information would also serve to increase the relative importance of existing noise.

backward temporal relation). The use of the terms neutral terms *forward* and *backward* serves to highlight the relationship between the relevant causal and temporal connectives. The connectives "and so" and "before" both indicate a *forward* temporal relation while "because" and "after" both indicate a *backward* temporal relation.

2.2.3 Materials

The materials for this experiment were similar to those used in Experiment 1.

In order to manipulate the implied coherence relation one of four connectives was used to connect the two clauses (see Table 5). Each participant received 4 sentence pairs using each connective (for a total of 16 sentence pairs). Two had a head clause for which the subject is the preferred interpretation for the *result* relation and the other 2 had a head clause for which the object is the preferred interpretation for the *result* relation[10].

The sentence pairs were organized in 8 groups of 2 heads and 2 tails each, such as the one described in Table 6. Appendix B includes a complete list of the materials used in this experiment. Each group had one head sentence of each of the two types of head. The two tail sentences were designed such that they could be paired with either head. The pairing of heads and tails within each group was counterbalanced across participants.

---

[10] This experiment is primarily concerned with the comparison between causal and temporal connectives for the same sentence pairs and interpretations. While different interpretations and temporal orders might represent sentence pairs with different plausibility, these variations are presumed to be independent of the manipulation of the connective.

2.2.4 Procedure

At the beginning of the experiment, the computer presented the participants with instructions. They informed participants that the experiment concerned the manner in which people determine the agents to which pronouns refer and that they would be asked to make such a determination for each of the sentence pairs they will be presented with. Participants were not provided with any explicit instructions on reading speed as the experiment presented participants with complete texts at a time and the assumption was that participants would read these texts at their normal pace.

Each sentence pair was presented at the top left of the computer screen. After reading the sentence pair, participants pressed the space bar and were presented with a screen that asked them to determine the referent of the ambiguous pronoun “he”. To indicate their choice of antecedent, participants pressed one of two keys – The ‘1’ key was used to indicate the subject and the ‘0’ key was used to indicate the object. These choices were presented directly below the sentence pair on the screen together with the names of the agents associated with them for that particular sentence pair. The name of the subject was on the left side and the object was presented on the right side. Following their choice, participants were also asked to indicate their confidence in their decision by typing a number between 1 (not confident at all) and 9 (extremely confident). The computer recorded participants’ choices as well as how long they took to read the sentence pair. The next sentence pair was presented immediately after they indicated their level of confidence.

2.2.5 Participants

Forty Northwestern university undergraduate students participated in this experiment.

2.2.6 Results

The proportion of object disambiguations, the mean confidence ratings, and the mean reading time per syllable for each of the conditions are given in Table 7. Prior to calculating mean reading times, outliers were removed from the data on a per-participant basis. Following Tukey (1977), an outlier was defined as a result that is further than 1.5 times the inter-quartile range from the appropriate quartile (below Q1 or above Q3). This resulted in the exclusion of 6.9% of the RTs. In addition, one participant did not provide confidence ratings in a condition and as a result their data was not used for the confidence analysis.

Overall, participants made the choices as predicted when a causal connective was used. When presented with a head sentence such as “John asked Mark for some study tips” and the tail sentence “he aced the test” they were likely to choose John as the referent for the ambiguous pronoun when the connective “and so” was used and Mark when the connective “because” was used. However, this effect was only slightly diminished when the temporal connective “before” was used and was not diminished at all when the temporal “after” was used. Nevertheless, participants’ confidence in their choice and their reading time did follow my prediction.

Participants were more confident in their choices when a causal connective was used and took less time to read sentence pairs involving causal connectives.

Table 7 – Disambiguation proportions and Mean Reading Times in Experiment 2 (standard deviations are given in parenthesis)

Connective Type	Temporal Sequence (Connective)	Expected Causal Disambiguation	% of Object Disambiguations	Mean Confidence	Mean Syllable Reading Time (ms)
Causal	Forward (and so)	Subject	0.35 (0.30)	7.19 (1.46)	348 (251)
		Object	0.76 (0.28)	6.99 (1.38)	327 (226)
	Backward (because)	Subject	0.31 (0.31)	7.14 (1.55)	333 (209)
		Object	0.78 (0.32)	6.99 (1.17)	339 (240)
Temporal	Forward (before)	Subject	0.39 (0.29)	6.28 (1.83)	410 (287)
		Object	0.63 (0.32)	6.18 (1.73)	421 (340)
	Backward (after)	Subject	0.30 (0.30)	6.54 (1.47)	363 (279)
		Object	0.78 (0.32)	6.89 (1.64)	354 (240)

As before, these results were tested using separate analyses of variance were conducted with participants (F_1) and sentence pairs (F_2) as random effects. A 3-way ANOVA between the connective type (causal or temporal), the temporal sequence (forward or backward), and the

expected referent (subject or object) was used to explore participants' proportion of object disambiguations. It confirmed that participants chose the object as the referent of the pronoun more often when the causal information provided suggested the object as the referent than when it pointed to the subject as the referent (F_1 (1, 39) = 12.8, *MSE* = .10, $p < .001$; F_2 (1, 15) = 13.11, *MSE* = .36, $p < .01$). Participants' responses did not differ due to the connective type or temporal sequence ($F < 1$). The subject analysis, but not the item analysis, also indicated that participants showed more sensitivity to the expected referent for backward temporal sequence relations than forward temporal sequence relations (F_1 (1, 39) = 5.53, *MSE* = .08, $p < .05$; F_2 (1, 15) < 1). The subject analysis also suggested a trend that the reduced sensitivity on the part of the participants in the forward temporal sequence was only evident when the temporal connective *before* was used but not when the causal connective *and so* was used (F_1 (1, 39) = 3.59, *MSE* = .06, $p = .066$; F_2 (1, 15) = 1.16, *MSE* = .06, $p = .30$). While these trends are weak, they might point to an interesting effect of the temporal order of events on reading that will be discussed later. Finally, no differences were found in participants sensitivity to the expected reference due to the type of connective (F_1 (1, 39) = 1.40, *MSE* = .08, $p = .24$; F_2 (1, 15) = 1.10, *MSE* = .044, $p = .31$).

While the type of connective used did not greatly affect the pattern of pronoun disambiguation it did show an effect on the confidence participants reported in their choices. An analysis of the confidence ratings showed that participants were more confident in their choices for sentence pairs involving a causal connective than those involving a temporal connective (F_1 (1, 38) = 29.85, *MSE* = 1.23, $p < .001$; F_2 (1, 15) = 27.25, *MSE* = .46, $p < .001$). There was also a trend that indicated this difference between connective types was somewhat larger for connectives indicating a forward temporal sequence ("and so" vs. "before") than those indicating

a backward temporal sequence ("because" vs. "after") (F_1 (1, 38) = 5.26, MSE = 1.02, $p < .05$; F_2 (1, 15) = 3.37, MSE = .65, $p = .09$). Similarly, the subject analysis, but not the item analysis also showed a trend where reported confidence was slightly higher for pairs using connectives indicating backward temporal sequence than a forward temporal sequence (F_1 (1, 38) = 4.75, MSE = 1.33, $p = .067$; F_2 (1, 15) = 1.75, MSE = 1.27, $p = .21$). These trends mirror similar trends observed in the analysis of reading times above and their possible implications will be discussed below. Participant confidence did not vary based on referent ($F < 1$) nor did referent affect participants confidence in conjunction with the type of connective (F_1 (1, 38) = 2.07, MSE = 1.17, $p = .16$; F_2 (1, 15) < 1) or temporal sequence (F_1 (1, 38) = 1.19, MSE = .73, $p = .28$; F_2 (1, 15) < 1). Likewise, no 3-way interaction was evident in participants' confidence ratings (F_1 (1, 38) = 1.00, MSE = 1.10, $p = .32$; F_2 (1, 15) < 1).

Finally, the hypothesized effect on reading times was also tested using a 3-way ANOVA. Participants read sentence pairs involving a causal connective faster than those involving a temporal connective (F_1 (1, 39) = 11.37, MSE = .02, $p < .01$; F_2 (1, 15) = 5.59, MSE = 0.014, $p < .05$). The subject analysis, but not the item analysis, also showed that participants read sentence pairs with the temporal connective "before" slower than those with the connective "after" but a similar difference did not exist between the causal connectives "and so" and "because" (F_1 (1, 39) = 12.8, MSE = .10, $p < .001$; F_2 (1, 15) < 1). This interaction also resulted in the subject analysis indicating that participants read sentence pairs with forward temporal sequence slower than those with a backward temporal sequence (F_1 (1, 39) = 4.57, MSE = .01, $p < .05$; F_2 (1, 15) < 1). The analysis found no effects of referent on participant reading times ($F < 1$). Likewise, the referent did not differentially influence participant reading times based on either connective type

or the temporal sequence ($F < 1$). Finally, no 3-way interaction was evident in participants' reading times (F_1 (1, 39) < 1; F_2 (1, 15) = 1.06, MSE = .014, p = .32).

Following the predictions, participants gave higher confidence ratings for their choices of referents when a causal connective was present than when a temporal connective was present. Furthermore, they spent more time reading sentences that employed temporal connectives. Nevertheless, these effects did not seem to significantly alter their choice of referent, suggesting that they were still willing to use causal information for the disambiguation even when an explicitly temporal connective was used. It is possible that if an alternative source of information was available, participants would use it when a temporal connective is presented rather than rely on the degree of causal relatedness as they did in this experiment.

Interestingly, the temporal sequence of the sentences might play a role in these decisions – in the more statistically powerful subject analysis it significantly interacted with the connective for both the confidence ratings and the reading times and with the expected referent in the case of the proportion of object disambiguations. In all three cases, this appears to be driven by a greater effect of the connective type when the connective indicates a forward temporal sequence of events. That is, there was a greater difference between sentence pairs that used the connectives "and so" and "before" than there was between those that used the connectives "because" and "after". One possible reason for this is that in a forward temporal sequence there is a parallel between the order in which the sentences are presented and the order in which events unfold. As a result, it is the most common type of temporal order found in narration. In contrast, a backward temporal sequence is often used to describe enabling conditions and therefore often involves some level of causal relatedness. Consequently, it is possible that participants view causal relatedness as more relevant for the temporal connective "after" than for the temporal connective

"before". If this is the case, participants would be more likely to look for evidence other than causal relatedness when the connective used is "before" than when it is "after". This will result in the longer reading times observed, as well as in the lower level of confidence participants report.

Alternatively, it is possible that the connective used to indicate a reversed temporal sequence ("after") is more closely associated with causal inferences than the connective used to indicate the normal temporal sequence ("before"). In that case, participants would be more likely to use causal information when the connective "after" was used simply because of this association.

The results of this experiment demonstrate that readers' use of causal information is sensitive to the coherence relation employed in the text. When the text specifies a relation that is explicitly causal in nature (e.g., through the use of the connectives "and so" or "because") readers feel confident in using causal relatedness as their guide to identifying the referent of an ambiguous pronoun. In contrast, when the text specifies a relation that is not explicitly causal (e.g., through the use of the temporal connectives "before" or "after") readers are less confident in their use of causal relatedness for the purposes of pronoun disambiguation. Nevertheless, when no other information is present, readers still deem causal information relevant and would use it as was evident by the pattern of disambiguation observed when temporal connectives were used.

These results are in line with the model I proposed. However, because the model as outlined is geared towards causal coherence relations, it might need to be modified to account for other types of relations. The reduced confidence ratings observed in this experiment suggest that weighing the evidence based on a combination of the type of evidence and type of relation might

be a plausible way to extend the model. Importantly, because in this experiment there was no competing (non-causal) source of information that participants could use, such weights would have little effect on the outcome of the disambiguation process as the ordering of the options would remain unchanged. However, because weights are often represented as a multiplication of the original scores, a reduction in the weights would reduce the absolute difference between the options readers are considering when making their choice. This is unlikely to significantly affect the choice made by the participant as the choice is governed by the order of the options. In contrast, confidence is likely to be based on the absolute difference between the options and should therefore be lowered by a reduction in the weights. This could explain why the pattern of disambiguation was largely the same regardless of the type of the connective, while confidence was higher for causal connectives than temporal connectives.

The evidence provided by these 2 experiments focuses on how causal information is used for pronoun disambiguation. In my interpretation of the results I suggested that participants are first inferring a causal relation between the events and then using this relation to identify the relevant interpretation. However, the causal information used in both of these experiments was based on the probability of one event (e.g., studying for a test) resulting in another (e.g., acing the test). An alternative explanation for these results could therefore be that the reader simply selects the interpretation in which the second event is the one that is most likely to follow the first (without requiring the intermediate step of inferring a coherence relation). Experiment 3 will use negation to explore a situation in which the predictions of the two explanations diverge.

2.3 Experiment 3 – Causality and Negation

The previous 2 experiments provided evidence that the process of pronoun disambiguation is sensitive to the degree of causal relatedness between sentences. In Experiment 1 participants used causal relatedness to determine the referent of an ambiguous pronoun even though they were not provided with any explicit information regarding the nature of the relationship between the two sentences. Experiment 2 examined the use of causal information in causal relations and temporal relations. Its results demonstrated that readers deem causal information as more relevant to the disambiguation of pronouns when the relationship is explicitly marked as causal than when it is marked with a temporal connective.

In both of these experiments participants selected the referent that produced the most coherent discourse. However, the interpretation chosen was also the one in which the two events had the highest degree of causal relatedness. It is therefore possible that participants made their choice not on the basis of coherence but because they were looking for an interpretation that maximized the causal relatedness of the events described.

Experiment 3 aims to contrast these two explanations by exploring a contrast between two different types of causal relations, namely *result* and *violated expectation*. The difference between these two relations is one of negation – whereas a *result* relation directly follows the direction of causality (i.e., the event described in the first sentence is the cause of the event described in the second), a *violated expectation* relation relates two events that *fail* to follow this relation even though such a causal relation is expected. This failure of expectations is often marked with connectives that have a sense of negation, such as *but*.

For example, it is relatively straightforward to infer that in the sentence pair (15), the pronoun "he" refers to Mark. That is, that Mark is the one who called the defense attorney. If the logic of the *result* relation is applied to (16), a different result is obtained. Because an event can either happen or not happen, the causal likelihood of something *not* happening (1-*p*) is generally the complement of the causal likelihood that it will happen (*p*). If participants rated the likelihood that Mark will call a defense attorney in (15) as 7 out of 9, then the likelihood of Mark *not* calling a defense attorney as in (16) can be inferred to be 2 (9 minus 7)[11]. Following the logic of the *result* relation, the addition of negation should reverse the preferred referent of ambiguous pronouns – i.e., the preferred referent for the pronoun "he" in (16) would be John. Therefore, if participants are simply using the degree of causal relatedness to identify the most relevant referent to an ambiguous pronoun in sentence pairs such as (15) and (16), negation of an event should lead them to reverse their preferred interpretation.

A different prediction arises if readers identify the relation in (16) as *violated expectation*. In that case, the expected inference would remain that the pronoun refers to the accused because in a *violated expectation* relation the important cause-effect relationship is between the cause and the *negation* of its effect rather than between the cause and its effect. That is, negation is an internalized as part of a *violated expectation* relation as explicated by Kehler (2004). The relation

[11] In reality, because causality is not the same as conditional probability, the negation of an unlikely causal consequence (e.g., "He did not run a marathon" as a consequence of (14a)) is unlikely to be perceived as a likely causal consequence either. Therefore, it is expected that ratings for negation will be uniformly low. However, a negation of a likely consequence is probably less likely than any random event so the overall pattern should still be reversed.

is the same as the *result* relation except that whereas in result the first proposition (*P*) normally gives rise to the second proposition (*Q*), in a *violated expectation* relation the first proposition (*P*) normally gives rise to the negation of the second proposition (¬*Q*). This means that if a reader decides that the appropriate coherence relation for (16) is *violated expectation*, they will determine the appropriate referent by considering the degree of causal relatedness of "John accused Mark of steal a computer" and the negation of "He did not call a defense attorney". Because the negation of "He did not call a defense attorney" is largely equivalent to "He called a defense attorney", an inference of a *violated expectation* relation in (16) should result in the same choice of referent as an inference of *result* in (15).

(15)

a. John accused Mark of stealing a computer.
b. He called a defense attorney.

(16)

a. John accused Mark of stealing a computer.
b. He did not call a defense attorney.

This case is interesting because it suggests that readers would be using the *same* causal information they use to identify a *result* relation to identify its complement, a *violated expectation* relation. However, because of its nature the use of the information in the latter case is based on counterfactual reasoning – the reader needs to be able to identify the relation between the event in (16a) and the event that *did not occur* in (16b). This can be somewhat

counterintuitive because it implies that readers use *negative information* (i.e., what is not the case) in order to determine the relevant relation. For instance, in (16) the reader is told that a defense attorney was not called. However, following the reasoning of a *violated expectation* relation, a reader that intends to determine the referent of the pronoun (i.e., *who* did not call a defense attorney) is actually trying to identify who the most likely referent of "He called a defense attorney" is. Importantly, because in (16) no defense attorney was called, it is true that neither John nor Mark called a defense attorney. This is not the case for (15), where it is likely that only one of the agents involved (Mark) called a defense attorney and therefore only interpretations in which the pronoun refers to Mark are true.

The notion that readers might be able to separate the negation in (16b) from the rest of the content of the sentence is in line with prior research into the processing of negation. For instance, many studies have found that negation incurs a processing penalty (e.g., Arroyo, 1982; Wason, 1961; cf. Kaup et al., 2007 for a review of these findings). Wason (1965) suggested that these difficulties are due to the pragmatics of the situation. Specifically, he found that the identification of a contrast to a negated item prior to exposure to a sentence facilitated its processing compared to negated items that did not have a salient contrast. This suggests that, at least in some cases, participants might be representing the non-negated form of the sentence in some manner. It is possible that such a representation could then be used to assess the validity of a *violated expectation* coherence relation.

Experiment 3 aims to contrast these two predictions regarding the negation of the effect in a cause-effect relation. On the one hand, if readers infer a *violated expectation* relation as defined by Kehler (2004), they should disambiguate it similarly to how they disambiguated the corresponding affirmative sentence pairs in Experiment 1. For instance, whenever participants in

Experiment 1 showed a clear preference for the subject as the person to whom the pronoun refers, so should participants in Experiment 3. For instance, if participants show a preference for Mark as the referent in (15b), they should also show a preference for Mark as the referent in (16b).

On the other hand, if readers maintain the *result* relation (i.e., choose the more causally related interpretation based on the given sentence pairs whether it is affirmative or negative) the pattern of disambiguation should be the opposite of that used for the corresponding affirmative version used in Experiment 1. In this case, whenever participants in Experiment 1 showed a clear preference for the subject as the person to whom the pronoun refers, participants in Experiment 3 should show a preference for the object as the person to whom the pronoun refers. That is, if participants show a preference for Mark as the referent in (15b), they should show a preference for John as the referent in (16b).

2.3.1 Predictions

Both hypotheses predict an interaction between the causal relatedness of the sentences and the preferred interpretation of the pronoun. However, if participants measure the causal likelihood between the sentences based on the form in which they are presented (i.e., they use a *result* relation in both conditions), then the preferred interpretation of the pronoun should be reversed when negation is applied to the tail sentence – that is, whenever participants prefer the subject position in a non-negated tail, they should prefer the object position in the corresponding negated tail and vice versa.

In contrast, if participants measure the causal likelihood between the sentences based on their affirmative form even when they are negated (i.e., they use a *violated expectation* relation in this experiment) then the pattern of pronoun interpretations should be the same regardless of the application of negation.

A third possibility is that the negation would cause participants to not infer any causality at all between the sentences – i.e., negation would disrupt the causal relation participants identify between the head and tail sentences. In that case there should be no effect of causal relatedness for the negated tails. That is, when reading sentence pairs with a negated tail from the high causality condition participants should perform similarly to how they perform in the low causality condition – pronoun disambiguations should be at chance and confidence should be low.

2.3.2 Design

Experiment 3 had the form of a 2x2x2 factorial.

The dependent variable was the proportion of object choices in the disambiguation of the tail sentence pronouns. Secondary dependent variables were the confidence participants have in their choice and their reading times (measured for the entire sentence pair and analyzed as reading time per syllable)[12].

The independent variables are the polarity of the tail (positive vs. negative), the causal relatedness of the head and the positive polarity version of the tail (high vs. low), and the

[12] Even though reading times were collected, no specific predictions were made were made for this measure.

preferred interpretation of the pronoun in the high causal likelihood, positive polarity, sentence pair (subject vs. object).

2.3.3 Materials

The materials for this experiment were roughly based on those used in Experiment 1 and used the same structure. However, in order to create a *violated expectation* relation, each tail had an additional version in which negation was added to the verb phrase. For instance the object tail "he called a defense attorney" was complemented by "he did not call a defense attorney" and the subject tail "he asked for payment" was complemented by "he did not ask for payment".

As in Experiment 1, each participant was presented with 16 sentence pairs. However, in this experiment half of these pairs involved a negated tail. Consequently, participants were presented with 2 sentence pairs in each of the 8 experimental conditions. The assignment of sentence pairs to condition was counterbalanced between participants in a manner similar to that used in Experiment 1. Appendix C provides a complete list of the materials used and notes the changes made to the tail sentences to transform the positive versions into negative ones.

2.3.3.1 Causality Ratings

To verify that the addition of negation changed the causal likelihoods in accordance with the hypothesis, another group of participants was asked to provide ratings of causality in a manner identical to the collection of those ratings in experiment 1. Raters were presented with a booklet containing 128 questions like those used in experiment 1 based on all possible combinations of head and tail sentences with the pronoun "he" replaced by either the subject

name or the object name (e.g., “If John accused Mark of stealing a car, how likely is that to cause John to call a defense attorney?” and “If John accused Mark of stealing a car, how likely is that to cause Mark to call a defense attorney?” as well as their negative forms “If John accused Mark of stealing a car, how likely is that to cause John to not call a defense attorney?” and “If John accused Mark of stealing a car, how likely is that to cause Mark to not call a defense attorney?”). However, because of the large number of possible combinations (256), two sets of booklets were made each containing the combinations of half of the sentences. An equal number of positive and negated tails was used in each booklet and all of the possible combinations for each tail were included in the same booklet. As in experiment 1, participants rated the causal likelihood on a 9-point likert scale.

2.3.4 Procedure

At the beginning of the experiment, the computer presented the participants with instructions. They informed participants that the experiment concerned the manner in which people determine the agents to which pronouns refer and that they would be asked to make such a determination for each of the sentence pairs they will be presented with.

Each sentence pair was presented at the top left of the computer screen. After reading the sentence pair, participants pressed the space bar and were presented with a screen that asked them to determine the referent of the ambiguous pronoun “he”. To indicate their choice of antecedent, participants pressed one of two keys – The ‘1’ key was used to indicate the subject and the ‘0’ key was used to indicate the object. These choices were presented directly below the sentence pair on the screen together with the names of the agents associated with them for that particular sentence pair. The name of the subject was on the left side and the object was

presented on the right side. Following their choice, participants were also asked to indicate their confidence in their decision by typing a number between 1 (not confident at all) and 9 (extremely confident). The computer recorded participants' choices as well as how long they took to read the sentence pair. The next sentence pair was presented immediately after they indicated their level of confidence.

2.3.4.1

2.3.5 Participants

Forty-six Northwestern university undergraduate students participated in this experiment, 33 in the disambiguation task and 13 in the rating task.

2.3.6 Results

The proportion of object disambiguations, the mean confidence ratings, and the mean reading time per syllable for each of the conditions are given in Table 8. The mean causal likelihood ratings for each of the conditions are given in Table 9. As before, prior to calculating mean reading times, outliers were removed from the data on a per-participant basis, as this was the unit predicted to have a uniform distribution according to the null hypothesis. Following Tukey (1977), an outlier was defined as a result that is further than 1.5 times the inter-quartile range from the appropriate quartile (below Q1 or above Q3). This resulted in the exclusion of 7.4% of the RTs.

Table 8 – Disambiguation proportions and Mean Reading Times in Experiment 3 (standard deviations are given in parenthesis)

Causality Level	Polarity	Expected Positive Disambiguation	% of Object Disambiguations	Mean Confidence	Mean Syllable Reading Time (ms)
High	Positive	Subject	0.03 (0.12)	7.83 (1.33)	307 (230)
		Object	0.95 (0.15)	7.31 (1.82)	365 (332)
	Negative	Subject	0.05 (0.15)	7.89 (1.25)	288 (196)
		Object	0.86 (0.26)	7.26 (1.49)	324 (259)
Low	Positive	Subject	0.27 (0.28)	4.12 (2.18)	387 (345)
		Object	0.62 (0.31)	3.77 (1.93)	382 (311)
	Negative	Subject	0.23 (0.31)	4.80 (1.98)	364 (386)
		Object	0.65 (0.34)	4.68 (2.14)	403 (356)

Table 9 – Mean ratings of causal likelihood in Experiment 3 (standard deviations are given in parenthesis)

Causality Level	Polarity	Expected Positive Disambiguation	Mean Causal Likelihood Rating	
			Subject	Object
High	Positive	Subject	6.40 (1.02)	1.37 (1.16)
		Object	2.66 (1.37)	7.07 (0.99)
	Negative	Subject	2.47 (1.17)	4.53 (2.16)
		Object	3.41 (1.98)	2.32 (1.33)
Low	Positive	Subject	1.32 (1.16)	1.34 (1.18)
		Object	0.99 (1.00)	1.09 (1.02)
	Negative	Subject	2.92 (2.20)	3.08 (2.28)
		Object	2.36 (2.18)	2.68 (2.47)

The general pattern of disambiguation and participants' confidence in their responses do not appear to be affected by negation. In contrast, negation has a dramatic effect on the ratings of causal likelihood provided by participants. The ratings of causal likelihood for high causality negative polarity sentence pairs show a reversed pattern than their positive polarity counterpart. Moreover, the overall rating in these cases is greatly decreased. This reduction indicates that negation tends to disrupt the causal links that readers draw between the head and tail sentences. In conjunction, these results suggest that participants use a coherence relation of *violated expectation* when the tail sentence of a sentence pair is negated rather than a *result* relation. Specifically, it appears that participants are basing their choice of referent on the positive polarity versions of the tail sentences even for negative polarity tail sentences.

2.3.6.1 Ratings of Causal Likelihood

Before examining the disambiguation, it is important to ensure that the addition of negation affected the perceived causal likelihoods. Specifically, I hypothesized that the negation of the high causality sentence pairs should disrupt the causal relationship and result in reduced ratings of causal likelihood. This hypothesis was tested by separate analyses of variance that were conducted with participants (F_1) and sentence pairs (F_2) as random effects. To simplify the interpretation of the analysis, the high and low causal likelihood sentence pairs were analyzed using separate 3-way ANOVAs, with the independent variables being the polarity of the tail (positive vs. negative), the expected disambiguation of the pair (subject vs. object), and the pronoun interpretation rated (subject interpretation vs. object interpretation).

As predicted, participants rated the high causality positive polarity items higher on average than the negative polarity items (F_1 (1, 39) = 12.8, MSE = .10, $p < .001$; F_2 (1, 7) = 6.23, MSE = .10, $p < .05$). Interestingly, they rated the low causality negative polarity items higher than the positive polarity items (F_1 (1, 39) = 12.8, MSE = .10, $p < .001$; F_2 (1, 7) = 6.23, MSE = .10, $p < .05$). While this effect was not predicted, it also was opposite in its direction to the predicted effect observed in the high causal likelihood condition. Participants rated high causal likelihood pairs as less likely to be causal when negated but they rated low causal likelihood pairs as more likely to be causal when negated.

The expected disambiguation and the interpretation used did not, by themselves, affect the rating of causal likelihood in high causal likelihood pairs ($F < 1$ in all cases). However, for the low causal likelihood pairs, there was a trend that was significant in the subject analysis but

not the item analysis where items that were expected to have a subject disambiguation were rated higher than those that were expected to have an object disambiguation (F_1 (1, 39) = 12.8, *MSE* = .10, $p < .001$; F_2 (1, 7) = 6.23, *MSE* = .10, $p < .05$). No such effect was found for the specific interpretation (F_1 (1, 39) = 12.8, *MSE* = .10, $p < .001$; F_2 (1, 7) = 6.23, *MSE* = .10, $p < .05$). It is possible that the effect of the expected disambiguation on the ratings was due to some peculiarities of the materials. Regardless, this main effect was relatively small and did not intrude on the predictions for this experiment as they focused on the effect of negation on high causal likelihood sentence pairs.

Most importantly, participants rated the high causality pairs according to my predictions. For positive polarity pairs the ratings were higher when the interpretation used matched the pronoun disambiguation expectation while for the negative polarity pairs this pattern was reversed – ratings were higher when the interpretation used did not match the pronoun disambiguation expectation (F_1 (1, 12) = 90.27, *MSE* = 2.85, $p < .001$; F_2 (1, 15) = 112.64, *MSE* = 2.75, $p < .001$). No such effect was found for low causal likelihood pairs ($F < 1$).

In addition, there were several 2-way interactions. Firstly, in the case of high causal likelihood pairs, participants rated items with an expected disambiguation to the subject lower than those with an expected disambiguation for the object in the positive polarity case but not the negative polarity case (F_1 (1, 12) = 25.28, *MSE* = .67, $p < .001$; F_2 (1, 15) = 9.15, *MSE* = 2.27, $p < .01$). Because participants rated the expected disambiguation high causality positive polarity items much higher than they rated the rest of the items (averages of 6.74 and 2.32 respectively) this translated to an overall higher rating for interpretations that matched the expected positive disambiguation than those that did not match (F_1 (1, 12) = 29.45, *MSE* = 2.18, $p < .001$; F_2 (1, 15) = 33.53, *MSE* = 2.39, $p < .001$). Lastly, for high causal likelihood pairs, participants showed

a trend in the subject analysis only where they were more sensitive to the interpretation used in positive polarity items than the negative polarity items (F_1 (1, 12) = 5.32, *MSE* = .78, $p < .05$; F_2 (1, 15) = 2.60, *MSE* = 1.91, $p = .13$). No such effects were found for low causal likelihood pairs ($F < 1$).

These results support the success of the manipulation. When presented with a high causal likelihood pair involving a positive polarity tail with an expected subject interpretation, participants rated the subject interpretation higher than the object interpretation for positive polarity pairs. This pattern was reversed for tails with an expected object interpretation. Furthermore, as predicted, this overall pattern was reversed for negative polarity tails.

2.3.6.2 Pronoun Disambiguation

Experiment 3 was designed to explore Kehler's (2004) suggestion that readers compute a *violated expectation* relation when there is a salient causal expectation that is the negated by the text. If participants compute this relation, then they should base their choice of referent on the causal likelihoods between the head sentence and the positive version of the tail regardless of its polarity. In contrast, participants might be interpreting sentence pairs using the *result* relation for both positive and negative tails. In that case, their interpretation should follow the causal likelihood ratings provided by the raters for the relevant sentence pair and should therefore differ based on the polarity of the tail. As a result their interpretation of the pronouns in negative polarity tails should be the opposite of that used for the positive polarity pairs (e.g., if they interpret the pronoun "he" in "he called the police" as referring to the subject, they should interpret it as referring to the object when presented with the negative tail "he did not call the

police"). Finally, if the negation simply removes the causal likelihood, they should interpret all the negated sentences similarly to their interpretation of the low causality sentence pairs.

Therefore, if participants are primarily using the *violated expectation* relation for high causal likelihood negative polarity pairs, results for both positive and negative polarity sentence pairs should replicate the results found in experiment 1. There should be an interaction between the level of causal likelihood and the expected disambiguation, as well as a main effect of the expected disambiguation. However, there should be no interaction of the expected disambiguation with the polarity of the sentence pairs. In contrast, if participants are using the *result* relation for high causal likelihood pairs regardless of their polarity, then an interaction of the expected disambiguation with the polarity of the sentence pairs is expected. Finally, if participants are treating high causal likelihood negative polarity sentence pairs as low causal likelihood then there should be a 3-way interaction between the variables as participants are treating negative polarity sentence pairs and low causal likelihood sentence pairs similarly and only using the available causal information for high causal likelihood positive sentence pairs.

To test these predictions, the proportion of object disambiguation were analyzed using a 3-way ANOVA with the independent variables being the level of causal likelihood (high vs. low), the polarity of the tail (positive vs. negative), and the expected disambiguation of the pronoun (subject vs. object). As before, separate analyses of variance were conducted with participants (F_1) and sentence pairs (F_2) as random effects.

The results of this experiment replicated those of Experiment 1 with respect to the positive polarity items. Moreover, negation did not appear to affect participants' choice of referent. Participants were more likely to disambiguate the sentence pairs based on the type of continuation in the high causality condition than the low causality condition (F_1 (1, 32) = 62.17,

MSE = .06, $p < .001$; F_2 (1, 15) = 19.66, *MSE* = .10, $p < .001$). As in Experiment 1, participants were also likely to disambiguate the sentence pairs based on their type of continuation (F_1 (1, 32) = 347.16, *MSE* = .08, $p < .001$; F_2 (1, 15) = 146.87, *MSE* = .09, $p < .001$) but showed no overall effects of level of causality (F_1 (1, 32) = 1.48, *MSE* = .04, $p = 0.23$; $F_2 < 1$). With respect to the specific hypothesis of this experiment, participants' choice of referent were not affected by the polarity of the tail (3-way interaction: F_1 (1, 32) = 2.51, *MSE* = .05, $p = 0.12$; F_2 (1, 15) = 1.47, *MSE* = .04, $p = .24$; All other effects involving polarity: $F < 1$).

As mentioned above, in addition to the main measure of proportion of disambiguation, participants were asked to provide confidence ratings for their choices. As in experiment 1, participants were generally more confident in their choices of referent for high causality sentence pairs than low causality sentence pairs. More importantly, the pattern for negative polarity items was the same as that for positive polarity items. This suggests that participants identified causal relationships between the head and tail sentences in the high causality condition for the negative polarity item. Combined with their pattern of choices described above, this is consistent with the choice of a *violated expectation* relation for the negative polarity sentence pairs rather than a *result* relation.

These results were tested using 3-way ANOVAs similar to those used to analyze the proportion of disambiguation. As in Experiment 1, participants were more confident in their choices in the high causality condition than the low causality condition (F_1 (1, 32) = 221.02, *MSE* = 3.12, $p < .001$; F_2 (1, 15) = 98.8 *MSE* = 3.33, $p < .001$). Interestingly, in the low causality condition, participants were more confident in their responses for the negative polarity items than the positive polarity items (F_1 (1, 32) = 4.41, *MSE* = 2.37, $p < .05$; F_2 (1, 15) = 6.88, *MSE* = .74, $p < .05$). As a result, participants also reported more overall confidence for negative polarity

items than positive polarity items (F_1 (1, 32) = 4.73, *MSE* = 2.21, $p < .05$; F_2 (1, 15) = 5.39, *MSE* = .95, $p < .05$). The subject analysis, but not the item analysis, also showed that participants reported more confidence in their choices when the expected referent was the subject than when it was the object (F_1 (1, 32) = 4.87, *MSE* = 2.23, $p < .05$; F_2 (1, 15) = 3.43, *MSE* = 1.53, $p = .084$). No other factors reliably affected participants' confidence in their choices (3-way interaction: $F < 1$; Causal likelihood and Expected Disambiguation: F_1 (1, 32) = 1.24, *MSE* = 1.55, $p = .27$; $F_2 < 1$; Polarity and Expected Disambiguation: $F < 1$). The main effect the level of causal likelihood provides evidence that the manipulation of the level of causal likelihood was successful. In contrast, the effects of polarity and expected disambiguation were not predicted. The difference based on expected disambiguation was also found in Experiment 1 and is likely due to differences among the items. It is possible that the difference in confidence based on polarity that was found in the low causality items reflects the higher ratings of causal likelihood for the low causality items. However, these higher ratings were still relatively low and did not appear to affect the actual choices made by the participants.

Finally, a similar analysis of variance was conducted on the syllable reading times. There were no a-priori predictions regarding reading time. Nevertheless, it is likely that the same factors that underlie the effects of confidence ratings would also affect reading times. Existing research further supports this and suggests that sentence pairs with a higher degree of causal relatedness are read more quickly (Keenan et al., 1984; Myers et al., 1987). The removal of outliers resulted in the removal of both responses for the low causal likelihood, negative polarity expected object interpretation sentence pairs of 3 of the participants and their reading times were consequently not used in this analysis.

In accordance with the effects reported in the literature, participants read sentence pairs in the high causality condition faster than those in the low causality condition (F_1 (1, 29) = 4.86, *MSE* = .06, $p < .05$; F_2 (1, 15) = 3.04, *MSE* = .07, $p = .11$). No other factors reliably affected participants reading times (interaction of level of causal likelihood and expected disambiguation: F_1 (1, 29) = 1.34, *MSE* = .02, $p = .26$; F_2 (1, 15) = 1.40, *MSE* = .02, $p = .26$; expected disambiguation: F_1 (1, 29) = 2.33, *MSE* = .02, $p = .14$; $F_2 < 1$; All other combination of factors: $F < 1$).

The previous experiments demonstrated that causal information plays a role in the disambiguation of pronouns. However, a causal relationship between events also implies that the two events are likely to co-occur. Therefore, it is possible that participants based their choices of pronoun referent simply based on this pattern of co-occurrence and not necessarily on the causal relatedness of the events described. The present experiment sought to decide between these two explanations through the use of negation. Specifically, negation would affect the likelihood that the two events co-occur but was hypothesized to preserve the cause-effect relationship between them through the use of a *violated expectation* relation.

The results of this experiment supported the hypothesis that when statements were negated participants preferred the pronoun interpretation supported by a *violated expectation* relation to that supported by the likelihood of co-occurrence interpretation. While the negation affected the ratings of causal likelihood provided by participants, it did not affect their choice of referents for the pronoun, nor did it alter their confidence in their responses or their reading times of the sentence pairs.

The experiments presented so far support a model in which readers' identification of causal relationships is tied with their choice of referents for ambiguous pronouns. I suggested

that this process starts with the identification of the possible interpretations of the discourse presented. Readers then compare the plausibility of these interpretations, weighing factors such as causal relatedness, and choose the most plausible interpretation. These experiments focused on the outcome of this process – the interpretation participants chose. In the experiment that follows I will attempt to look into the influence that restricting the number of interpretations has on the process. Specifically, I will suggest that readers might use an unambiguous pronoun as a clue to the type of causal relationship that might exist between a pair of sentences.

2.4 Experiment 4 – Time Course of Causal Coherence

Experiment 1-3 examined the outcome of the processing of the *result* coherence relations in cases where an ambiguous pronoun is present. Experiment 4 examines how the ambiguity of a pronoun affects the on-line processing of a sentence pair involving a *result* relation.

As mentioned earlier, in order to identify the coherence relation that binds a pair of sentences together, readers must evaluate alternative interpretations of the sentence pair. Two factors combine to generate these alternative interpretations:

1. The coherence relation.
2. Any syntactic, semantic or pragmatic ambiguities that might exist in the sentences themselves.

In the sentences presented in the previous experiments, there was no real ambiguity involved when a sentence pair was considered as a whole. But even in these cases, some ambiguity often persists until enough of the sentence has been processed. For example, without knowing what the event described in the second sentence is (i.e., prior to reading the verb), readers might be hard pressed to clearly identify which entity an ambiguous pronoun refers to.

Rohde et al. (2006) noted that there is a tendency for specific coherence relations to employ pronouns to refer to entities that play a specific role in the previous sentence. For instance, participants who used an *explanation* relation in their completion were highly likely to use the ambiguous pronoun to refer to the subject (or *agent*) of the previous sentence (e.g., “Matt

passed a sandwich to David. He didn't want David to starve")[13]. This suggests that the referent of a pronoun might be a useful clue to the coherence relation that might obtain between two sentences.

This hypothesis is not unlike the conclusion Rudolph and Forsterling (1997) draw in their review of the implicit causality literature. They suggest that implicit causality can be explained, in part, as a result of a covariation model where readers are responding to the overall proportion of subject or object references they encountered in similar constructions in the past (e.g., sentences involving the same causal verb). However, whereas the implicit causality literature explores the factors that determine biases for pronoun interpretation for *explanation* relations, the present experiment suggests that the *type* of relation is also subject to prediction based on factors such as covariation statistics. Specifically, I hypothesize that the pattern of pronoun referents observed by Rohde et al. (2006) for each coherence relation would also predict readers' expectations regarding a specific coherence based on the referent of the pronoun encountered. For instance, if a reader encounters a pronoun referring to the subject (or *agent*) of a previous sentence, an *explanation* relation would be considered more likely than a *result* relation.

---

[13] Rohde et al. term these relations in terms of the thematic role of the subject and the object. However, the thematic roles and syntactic roles coincide in their experiments as well as in the present study. Because syntactic roles are more uniform they will be used here instead of the *source* and *goal* thematic roles used by Rohde et al. For the present study it is important to note that the important distinction is between *agent* and *patient* of which *source* and *goal* are a subset. In active sentences, the *agent* generally corresponds to the *subject* and the *patient* to the *object*.

Experiment 4 tests whether unambiguous pronouns are used as predictors for the type of coherence relation that might link the new sentence with the previous one. Specifically, Experiment 4 will contrast sentence pairs that are bound via a *result* relation with pairs that are bound via an *explanation* relation. According to the data in Rohde et al. (2006) *result* relations predominantly refer to the *goal* (or object) position of the previous sentence while *explanation* relations predominantly refer to the *source* (or subject) position.[14]

As demonstrated by Kehler et al. (2008), while various verbs might generate expectations regarding the referent in the next sentence through biases such as the implicit causality bias, those biases are tied to specific coherence relations and consequently require that the reader be aware of the relationship between the sentences. When the relationship between the sentences is disregarded (e.g., when examining all of the completions reported by Kehler et al., 2008, in Experiment 3) a subject bias is commonly found. Consequently, since Experiment 4 does not use connectives or otherwise provide readers with a-priori reasons for expecting one type of coherence relation over another, the biases of the verbs should not have a great impact on the results of the study.

If readers use the co-reference information to guide their interpretation of the coherence relation then they should find the rest of the sentence easier to read and correspondingly read it more quickly than when the pronoun is ambiguous. This effect may be reversed when the co-reference information provided by the pronoun contrasts with the actual coherence relation

---

[14] The terms *source* and *goal* used by Rohde et al. (2006) were used because the sentences used transfer-of-possession verbs.

between the two sentences. In such cases readers might find the rest of the sentence more difficult to read than when the pronoun is ambiguous because they were expecting a different coherence relation to obtain between the sentences (this effect might be similar to that of "garden path" sentences, where an early syntactic ambiguity results in misleading expectations). Because this effect depends on the ambiguity of the pronoun, it should only be expected to affect reading times of words that follow the pronoun (e.g., the verb phrase).

For instance in (17) the pronoun "he" is ambiguous. Therefore the reading time of the verb phrase "reported it to the police" needs to include the disambiguation of the pronoun (i.e., the reader needs to decide that "John reported it to the police" is more likely than "Mark reported it to the police". The unambiguous referent for the pronoun in (18) is likely to result in faster reading times for the verb phrase in that case than in (17).

(17)

a. John alleged that Mark stole an expensive car.

b. He reported it to the police.

(18)

a. John alleged that Mary stole an expensive car.

b. He reported it to the police.

This effect might be further complicated by the prediction that readers expect that a pronoun referencing the subject will be more likely to be found when the coherence relation is

explanation while a pronoun referencing the object is more likely when the coherence relation is *result* (as in (17) and (18)). Consequently, we might expect the presence of an unambiguous pronoun to provide a greater decrease in reading times when it is congruent with the expectation of the coherence relation as in (19) than when it is incongruent with the expectation as in (17).

(19)

a. Jane alleged that Mark stole an expensive car.

b. He sold it to a fence.

The prediction is therefore that in the case of a *result* relation, participants will read the verb phrase fastest when the pronoun is unambiguous and referring to the object[15]. Likewise, when the two sentences are related with an *explanation* relation, participants will read the verb phrase fastest when the pronoun is unambiguous and referring to the subject.

In addition, if participants form expectation of the coherence relation when presented with an unambiguous pronoun, it is possible that this will result in the need for "backtracking" if this expectation is violated and therefore will result in increased reading times compared to cases

---

[15] While in the examples above the object of the sentence is actually an embedded clause ("X stole an expensive car."), the pronoun refers to a person mentioned in this embedded clause. 2 of the experimental groups used this construction while the others did not employ an embedded clause. For consistency, I will use the term 'object' for pronouns referring to both a person in the actual object position and to a person in the subject position of an embedded clause in the object position. However, this difference is not likely to influence the predictions since in all of the experimental sentence pairs the *subject* position precipitated an event that influenced the person in the *object* position.

where they do not form such expectations (e.g., when the pronoun is ambiguous). Consequently, it is possible that when the referent of the pronoun is incongruent with the coherence relation, participants will show higher reading times for the verb phrase in cases where the pronoun is unambiguous than in cases where the pronoun is ambiguous. However, ambiguous pronouns are also likely to incur an additional cognitive load on participants compared to unambiguous pronouns. Therefore, it is expected that participants will read the verb phrase more slowly overall when an ambiguous pronoun is used. This makes specific predictions regarding verb phrase reading times in the incongruent case more difficult.

Finally, following results such as those reported in Kehler et al. (2008), it is possible that an overall subject bias would be observed when an ambiguous pronoun is used. Such a bias would result in faster reading times for verb phrases when the ambiguous pronoun refers to the subject than when it refers to the object.

2.4.1 Predictions

The prediction of this experiment is based on the hypothesis that the role an entity plays in the first sentence influences the likelihood it will participate as the subject in either a *result* or *explanation* relation in the second sentence. Specifically, subjects in the first sentence are more likely to participate in an *explanation* relation while objects are more likely to participate in a *result* relation. Consequently, when the unambiguous pronoun refers to the object of the first sentence, it will facilitate the processing of *result* relations but not *explanation* relations. The reverse prediction holds for the case where the pronoun unambiguously refers to the subject of the first sentence. When an ambiguous pronoun is presented, participants are less likely to draw

specific predictions and therefore the type of coherence relation should not have much of an effect (although overall biases, such as the subject bias mentioned above might still influence their performance). Nevertheless, because an ambiguous pronoun results in an increase in the number of interpretations participants must consider, it is expected that participants will take longer to read the verb phrases of sentences involving ambiguous pronouns. Statistically, this prediction results in a 3-way interaction between the ambiguity of the pronoun, the type of the relation, and whether the pronoun refers to the subject or the object of the first sentence. Participants are predicted to read the verb phrase fastest when the pronoun is unambiguous and refers to the subject in the case of an explanation relation or the object in the case of a result relation.

Another question of interest regards the point at which these expectations might affect readers' performance. Because the core of an event is contained in the verb that describes it, it is possible for expectations regarding the type of coherence relation will affect the reading time of the verb itself (e.g., the relationship between the verbs *allege* and *report* in (17) is sufficient to identify a *result* relation). Another possibility is that the selection of a coherence relation requires more than just the verbs. That is, readers might wait for more information before committing to an interpretation (e.g., readers will only make their decision after knowing that the report was made to the police). In this case, we might expect prior expectations regarding the coherence relation to affect reading time only after the verb is read. Consequently, it is of interest to examine the reading time of the verb and its complement separately.

2.4.2 Design

Experiment 4 has the form of a 2x2x2 factorial.

The dependent variables are the reading time of the verb and the verb complement (all of the words past the verb) in the tail sentence.

The independent variables are the type of relation (*result* vs. *explanation*), the ambiguity of the pronoun when it is read (ambiguous vs. unambiguous), and the actual referent of the pronoun after the entire sentence is read (subject vs. object).

2.4.3 Materials

The materials for this experiment were pairs of sentences similar in their construction to those used in Experiment 1. However, instead of using 2 possible tails in each group there were 4 possible tails (see Table 10). Two of the tails were possible results of the events described in the head sentences (a *result* relation). In one of the *result* tails the pronoun referred to the subject of the head sentence (e.g. “He reported it to the police.”) while in the other the pronoun referred to the object of the head sentence (e.g., “He sold it to a fence.”). The other two tails were possible explanations of the events described in the head sentences (an *explanation* relation). As before, in one of the explanation tails the pronoun referred to the subject of the head sentence (e.g., “He saw it happen.”) while in the other the pronoun referred to the object of the head sentence (e.g., “He needed the money.”)[16]. Additionally, there were 4 heads in each group paired with the 4 possible tails. These heads were similar to each other such that each could be used with any of

[16] In 2 of the experiment groups the object of the head sentence was an embedded clause. In these cases the pronoun referred to the subject of this embedded clause. Notably, in these cases the referent of the pronoun is still a part of the object of the main sentence.

the tails. Therefore, each pairing of a head sentence and a tail sentence resulted in a sentence pair belonging to one of the following conditions (based on the type of the tail sentence):

(1) Head + Result Subject Tail

e.g., “John alleged that Mark stole an expensive car. He reported it to the police.”

(2) Head + Result Object Tail

e.g., “John alleged that Mark stole an expensive car. He sold it to a fence.”

(3) Head + Explanation Subject Tail

e.g., “John alleged that Mark stole an expensive car. He saw it happen.”

(4) Head + Explanation Object Tail

e.g., “John alleged that Mark stole an expensive car. He needed the money.”

The ambiguity of the pronoun was manipulated via the gender of the entities described in the head sentence. While in the previous experiments the two entities in each sentence pair were of the same gender, here the entities had different genders in half of the sentence pairs presented to the participants. When both entities share the same gender, the pronoun used is ambiguous. In contrast, when two different genders are used, the pronoun is unambiguous.

Table 10 – Sample group of materials for Experiment 4

Head	Type of tail	Tail
John alleged that Mark stole an expensive car.	Result Subject continuation	He reported it to the police.
John claimed that Mark shoplifted a pricey suit	Result Object continuation	He sold it to a fence.
John stated that Mark forged a painting.	Explanation Subject continuation	He saw it happen.
John declared that Mark hijacked a motorcycle.	Explanation Object continuation	He needed the money.

There were 8 groups of materials such as that provided in Table 10 for a total of 32 sentence pairs. The complete list of materials used can be found in Appendix D. The order of presentation of the sentence pairs was randomized on a per-participant basis. In addition, each tail can be presented with 1 of 4 possible head sentences. Four different sets of 32 sentence pairs were generated such that each possible head and tail combination appeared in one of the sets. Additionally, there are 4 possible gender assignments for each head sentence (male or female subject and male or female object). Importantly, these gender assignments also control the ambiguity of the pronoun – The pronoun in the tail sentence is ambiguous if, and only if, the subject and the object in the head sentence share the same gender. These were also counterbalanced independently between participants. Of the 8 sentence pairs of a specific tail

type, 2 sentence pairs used each of the possible gender assignments. Consequently, there were 16 counterbalancing conditions, one for each combination of head/tail and gender assignments.

Each participant was presented with 32 sentence pairs. Of these 32, each of the 4 possible gender assignments was used in 8. These pairs were equally divided among the 4 different tail types and therefore each participant read 2 sentences for each combination of gender assignment and tail type. Because two of the gender assignments are ambiguous and two are unambiguous, each participant was presented with 4 sentence pairs of each of the 8 experimental conditions (2 with the pronoun referring to a male and 2 with the pronoun referring to a female).

Importantly, these sentences were all semantically reasonable. Moreover, since each sentence pair appeared in both the ambiguous and unambiguous conditions, differences between these conditions and interactions that include ambiguity as a factor should not be affected by factors intrinsic to the words used, such as word frequency.

2.4.4 Procedure

At the beginning of the experiment, the computer presented the participants with instructions. They informed participants that the experiment concerned the manner in which people determine the agents to which pronouns refer and that they would be asked to make such a determination for each of the sentence pairs they will be presented with. The instructions also mentioned that the sentences will be presented one word at a time and that the space bar would be used to proceed through the sentence.

Each sentence was presented on the top left of the screen one word at a time. Words not presented on the screen were replaced by blanks so the words participants view on the screen

progressed naturally.[17] After participants pressed the space bar to advance past the last word of the sentence they were presented with a screen that asked them to determine the referent of the ambiguous pronoun that was used in the pair (i.e., 'Which character did the pronoun in the second sentence refer to?'). To indicate their choice of antecedent, participants pressed one of two keys – The '1' key was used to indicate the subject and the '0' key was used to indicate the object. These choices were presented directly below the question, together with the names of the agents associated with them for that particular sentence pair. The name of the subject was on the left side and the object was presented on the right side. Following their choice, participants were also asked to indicate their confidence in their decision by typing a number between 1 (not confident at all) and 9 (extremely confident). The computer recorded participants' choices as well as how long they took to read each word. The next sentence pair was presented immediately after they indicated their level of confidence.

2.4.5 Participants

Sixty-five Northwestern University undergraduate students participated in this experiment.

2.4.6 Results

The proportion of object disambiguations, the mean confidence ratings, and the mean reading time per syllable for the verb and the verb complement of each of the conditions are

---

[17] The reason the non visible words are not replaced by dashes as is customary is because the critical pronouns 'he' and 'she' are of different length and might therefore provide clues to the participants prior to their presentation.

given in Table 11[18]. Prior to calculating mean reading times, responses that did not match the expected disambiguation[19] were removed from the data[20]. Additionally, outliers were removed on a per-participant basis, as this was the unit predicted to have a uniform distribution according to the null hypothesis. Following Tukey (1977), an outlier was defined as a result that is further than 1.5 times the inter-quartile range from the appropriate quartile (below Q1 or above Q3). This resulted in the exclusion of 4.4% of the RTs for the verb and 8.6% of the RTs for the verb complement. In addition, two participants made no correct responses in one of the conditions and as a result their data was not used for the reading time analysis. Finally, 3 items did not have a verb complement and were excluded from the verb complement analysis.

---

[18] The reading times for the pronoun were virtually identical in all conditions (within 30ms of each other).

[19] It is important to note that unlike Experiments 1 and 3, there were no "low causality" condition in this experiment. Consequently, when a complete sentence pair is read, the referent of the pronoun is clear even when the pronoun itself can refer to either the subject or the object. In this experiment, a pronoun is ambiguous only in so far as it matches the genders of both the subject and the object. This ambiguity is resolved once the entire sentence pair has been read. Therefore, each sentence pair has a "correct" and an "incorrect" response regardless of the pronoun ambiguity condition.

[20] The reading time effects described below remain largely the same when incorrect responses are included. However, when participants provide an incorrect response it is possible that they did not pay enough attention while reading and therefore it was considered preferable to analyze only correct responses.

Table 11 – Disambiguation proportions and Mean Reading Times in Experiment 4 (standard deviations are given in parenthesis)

Relation	Pronoun Referent	Pronoun Ambiguity	% of Object Responses	Mean Confidence	Mean Syllable Reading Time (ms)	
					Verb	Verb Complement
Result	Subject	Ambiguous	0.16 (0.19)	7.70 (1.16)	162 (150)	195 (165)
		Unambiguous	0.02 (0.09)	8.75 (0.42)	149 (132)	169 (145)
	Object	Ambiguous	0.95 (0.14)	7.85 (1.10)	177 (153)	223 (213)
		Unambiguous	0.96 (0.14)	8.82 (0.39)	175 (149)	150 (128)
Explanation	Subject	Ambiguous	0.07 (0.15)	7.76 (1.09)	159 (138)	208 (178)
		Unambiguous	0.02 (0.07)	8.71 (0.82)	159 (135)	170 (143)
	Object	Ambiguous	0.84 (0.22)	7.20 (1.41)	134 (114)	232 (209)
		Unambiguous	0.95 (0.15)	8.67 (0.66)	133 (121)	218 (185)

Reading times for the verbs show little difference based on the ambiguity of the pronoun. In contrast, the reading times for the verb complement are generally faster when the pronoun is unambiguous and show the predicted interaction of coherence relation and referent. Participants read the verb complement faster when an unambiguous pronoun referred to the object rather than the subject of a *result* relation and the subject rather than object of an *explanation* relation. No such pattern is seen when the pronoun is ambiguous. Instead, participants appear to have read the verb complement faster when the ambiguous pronoun referred to the subject than when it

referred to the object regardless of the coherence relation involved. This represents a subject bias that is congruent with the predictions drawn based on Kehler et al. (2008).

The proportion of object responses supports the results reported by Rohde et al. (2006). When the pronoun was ambiguous, participants showed a preference for the object in the case of *result* relations and the subject in the case of *explanation* relations. For instance, following the head sentence "Bart suggested Aaron was fibbing", readers chose Aaron as the referent of "He wanted to avoid punishment" only 75% of the time. An examination of the sentence pairs where participants were likely to make such unexpected responses suggests that these pairs have some ambiguity that warrants these interpretations. For instance, while it is easier to conceive that Bart was telling the truth and Aaron lied to avoid punishment, it is also possible that Bart's accusation was a misdirection because Bart wanted to avoid punishment. These interpretations were less likely than the primary explanations but can account for the observed results. Finally, as expected, when the pronoun was unambiguous there was very little variation in participants' choice of referent. Likewise, participants' ratings of confidence in their choices were higher when the pronoun was unambiguous than when it was ambiguous.

Before proceeding with the statistical analysis, it is important to note that the head sentences used two distinct classes of verbs – mental verbs (e.g., 'felt', 'desired') and transfer-of-possession verbs (e.g., 'gave', 'handed'). Since these verbs have been shown to have different implications and biases (e.g., mental verbs often exhibit implicit causality biases), it was possible that the class of verb influenced the observed effects. However, the class of verbs used in the head sentence did not affect the reading time of the verb phrases and the following analyses therefore collapses over the class of verb used in the head sentence.

As before, these results were tested using separate analyses of variance with participants (F1) and sentence pairs (F2) as random effects. Before examining the reading times of the verb and verb complement of the tail sentences, it might be instructive to explore how the type of relation affected participants' choice of referent. To that effect, a 3-way ANOVA between the type of relation (*result* vs. *explanation*), the expected referent (subject vs. object), and the ambiguity of the referent (unambiguous vs. ambiguous) on the proportion of object disambiguations was conducted. Not surprisingly, participants choices of referent followed the expected referent of the tail sentence (F_1 (1, 64) = 2189.83, *MSE* = .04, $p < .001$; F_2 (1, 28) = 1467.31, *MSE* = .008, $p < .01$). Participants also showed a preference for the object as the referent in the case of an *explanation* relation and the subject as the referent in the case of a *result* relation (F_1 (1, 64) = 20.08, *MSE* = .02, $p < .001$; F_2 (1, 28) = 6.35, *MSE* = .008, $p < .05$). Participants were more likely to show these preferences when the pronoun was ambiguous than when it was unambiguous (F_1 (1, 64) = 12.16, *MSE* = .02, $p < .01$; F_2 (1, 28) = 5.50, *MSE* = .007, $p < .05$). In addition, participants were more likely to choose the correct referent for unambiguous tails than for ambiguous tails (F_1 (1, 64) = 52.06, *MSE* = .02, $p < .001$; F_2 (1, 28) = 14.99, *MSE* = .007, $p < .01$). No other reliable differences were found in the analysis (Ambiguity: F_1 (1, 64) = 1.68, *MSE* = .02, $p = .2$; F_2 (1, 28) < 1; All other effects: $F < 1$).

The pattern of these results supports the observation that when the pronoun was unambiguously referring to the subject or the object, the proportions of object disambiguations were at floor or ceiling, respectively. However, when participants were presented with an ambiguous pronoun, they presented a slight bias that differed based on the coherence relation between the sentence – they showed a bias towards the object when the relation was *result*, and a bias towards the subject when the relations was *explanation*. Because the proportions in both the

ambiguous and unambiguous cases were equally divided between cases where the pronoun refers to the subject and those in which the pronoun refers to the object, the overall proportions for both the ambiguous and unambiguous cases are close to 50%. As a result, the ambiguity of the pronoun only significantly affected performance as part of the 2-way interactions with the type of relation and expected referent[21]. Overall, these results are in accordance with the results presented by Rohde et al. (2006) where participants showed an object bias in *result* relations and a subject bias in *explanation* relations.

I now turn to examining participants' reading times of the verb and verb complement of the tail sentences. ANOVAs similar to that used to explore the proportion of pronoun disambiguation were conducted on the reading times for both the verb and verb complement. In the case of the verb reading times, no effect was significant for both subject and items. Overall, all of the reading times for the verbs were within 50ms of each other and the lack of effects that are statistically significant in both analyses suggests that the type of relation might not affect the reading time of the verb to a significant degree.

In contrast with the reading times for the verbs, the reading time for the verb complements was in line with the predictions outlined earlier. Overall, participants were faster to read verb complements for an unambiguous pronoun than an ambiguous pronoun (F_1 (1, 62) = 31.80, MSE = .006, $p < .001$; F_2 (1, 25) = 34.94, MSE = .0006, $p < .001$) and this difference was more pronounced for the *result* relation than the *explanation* relation (F_1 (1, 62) = 7.36, MSE =

[21] The interaction with expected referent is simply a result of participants making more errors when the pronoun was ambiguous. In those cases participants made more subject choices when an object was expected and object choices when the subject was expected.

.002, $p < .01$; F_2 (1, 25) = 4.38, *MSE* = .0006, $p < .05$). However, they gained more of an advantage from the unambiguous pronoun when it referred to the object of a *result* relation and the subject of an *explanation* relation (F_1 (1, 62) = 9.57, *MSE* = .004, $p < .01$; F_2 (1, 25) = 7.56, *MSE* = .0006, $p < .05$).

The subject analysis, but not the item analysis, also indicated that participants read the verb complement of *result* tails faster than *explanation* tails (F_1 (1, 62) = 18.34, *MSE* = .004, $p < .001$; F_2 (1, 25) = 2.78, *MSE* = .003, $p = .11$) and that verb complements of tails in which the pronoun referred to the subject were read faster than those in which the pronoun referred to the object (F_1 (1, 62) = 13.33, *MSE* = .004, $p < .01$; F_2 (1, 25) = 3.14, *MSE* = .003, $p = .09$). That sentences involving a *result* relation are read faster than those involving an *explanation* relation can be seen as an artifact of the expectations of narrative – readers expect a forward temporal progression (as in a *result* relation). As for the second effect, it is possible that the subject of the sentence is more accessible and that this increased accessibility facilitates the reading of verb complements that refer to the subject.

The subject analysis also indicated that participants read verb complements of *result* tails that referred to the object faster than those that referred to the subject and *explanation* tails that referred to the subject faster than those that referred to the object (F_1 (1, 62) = 14.90, *MSE* = .003, $p < .001$; F_2 (1, 25) = 1.22, *MSE* = .003, $p = .28$). This result mirrors the main prediction (except that both ambiguous and unambiguous sentence pairs are considered) and the resulting 3-way interaction and is therefore likely to be a side effect of that 3-way interaction.

Finally, no reliable interaction was found between the referent and the ambiguity of the pronoun (F_1 (1, 62) = 1.46, *MSE* = .003, $p = .23$; F_2 (1, 25) = 1.57, *MSE* = .0006, $p = .22$).

With respect to the original prediction, the most important results are that participants were faster when reading the verb complement of the tail when the pronoun was unambiguous and that this effect depended on the combination of the coherence relation and the referent. Because these differences were found for the verb complement but not the verb, it seems likely that readers do not make their choices regarding the coherence relation immediately after they have read the verb but wait for a more complete representation of the event described in the sentence.

Moreover, because the reading times of the verb complement for ambiguous pronouns showed a subject bias rather than a preference based on the relation, it appears that the interaction of the referent with the relation was specific to cases when the pronoun referent was unambiguous. This supports the hypothesis that the source of the interaction between the pronoun referent and the coherence relation is due to expectations generated prior to the verb complement. Otherwise, the same effects would likely be observed regardless of the ambiguity of the pronoun.

Interestingly, the effect of ambiguity was also greater overall for the *result* relation than the *explanation* relation. This result was not predicted, but suggests that participants might be more likely to rely on predictions in the case of the *result* relation than the *explanation* relation.

The previous 3 experiments demonstrated that causal relations play a role in the disambiguation of pronouns. This experiment focused on the role of pronoun reference in generating expectations of either *result* or *explanation* causal relations. I hypothesized that reference to the subject or object of the previous sentence induces readers to form specific expectations regarding the nature of the coherence relation between the previous sentence and the current one. The experiment supported this hypothesis – but only after the reader has already

processed the main verb of the sentence. Therefore, it is likely that these expectations do not come into play immediately after the pronoun is read but only after the reader has a firm idea regarding the event described in the sentence. This is in line with the model I proposed in which readers first generate a set of possible interpretations and then choose the most plausible interpretation in the set. While the set of interpretations can be built incrementally as the sentence is read, it is likely that judgments of plausibility can only be made after a significant portion of the sentence has been processed.

3 Discussion

This dissertation is concerned with the relationship between causal coherence relations and pronoun disambiguation. The first experiment demonstrated that this relationship exists by identifying a strong correspondence between the choice of referent for an ambiguous pronoun and ratings of causal likelihood. Participants were presented with sentence pairs in which the second sentence started with a pronoun that could refer to either the subject or the object of the first sentence. Their choice of pronoun referent was congruent with the ratings of causal likelihood. When the ratings of the interpretation using the subject as the referent of the pronoun were higher than those of those using the object as the referent, participants showed a strong preference for the subject as the referent of the pronoun and vice versa. However, when the second sentence was unlikely to be a consequence of the first regardless of the referent of the pronoun, the pronoun was ambiguous. That is, participants' choice of referent for such sentence pairs varied considerably.

The second experiment explored the role causality played in these choices. Specifically, I hypothesized that readers are more likely to use causal information for pronoun disambiguation when the relationship between the sentences was explicitly identified as causal rather than temporal. The type of relationship was explicated using connectives that identified either a causal relationship ("and so" and "because") or a temporal relationship ("before" and "after"). While this manipulation did not result in a reliable difference in the proportion of disambiguation, participants were less confident in their responses when a temporal connective was used than when a causal connective was used. Participants also read the sentence pairs more slowly when they involved a temporal connective than when they involved a causal connective. This indicates that participants were aware that the causal information was more relevant for causal

relationships than temporal relationships. However, because there was no alternative source of information present, participants had to rely on causal information for the purposes of pronoun disambiguation even when a temporal connective was used.

The third experiment contrasted two possible mechanisms that might be used to explain the patterns of results obtained in Experiments 1 and 2. Following Kehler (2002), I hypothesized that participants in Experiment 1 were trying to find how the two sentences relate. Because of the way the sentence pairs were constructed, these relationships were likely to be based on causal information. However, this causal relationship depends on the referent of the pronoun and consequently, participants were forced to make such a choice as part of the process of identifying the relevant relation. Likewise, in Experiment 2, participants were trying to identify the best interpretation of the pronoun given a particular type of relationship between the sentences (as identified by the connective used). However, there is a simpler mechanism that can account for these results – participants might be choosing the pronoun such that the resulting interpretation maximizes the likelihood that the second sentence is the consequence of the first sentence. This mechanism does not require readers to look for any specific type of relationship between the sentence pairs. Instead, readers are simply weighing the odds that one event is to follow another.

Experiment 3 contrasted these two explanations by looking at cases where the second event is negated. Negation accords with Kehler's relation of *violated expectation* which explicitly incorporates this negation. Consequently, if participants choose to use this relation for negated events their choices of referent should follow the causal likelihoods of the interpretations of the positive versions of the sentence pairs. However, if participants were basing their choice on the most likely causal consequence of the first event regardless of negation, then their choices of referent should follow the causal likelihood of the actual sentence pairs (i.e., they should make

different choices for the negative versions of the sentence pairs than they did for the positive versions of the sentence pairs). The results of Experiment 3 show that while the ratings of causal likelihood were affected by the addition of negation, the choices of referent were not. Likewise, participants' confidence and reading times were also unaffected by the presence of negation. Consequently, Experiment 3 provides evidence suggesting that participants are trying to identify the type of relation that exists between the sentences and not basing their choice directly on the most likely consequence of the previous sentence.

Finally, Experiment 4 investigated how the ambiguity of the pronoun might affect readers' expectations regarding the type of coherence relation between the two sentences. I hypothesized that an unambiguous pronoun might lead readers to expect specific types of relations in favor of others. Specifically, based on the base rates reported by Rohde et al. (2006) I hypothesized that a reference to the subject should lead to a preference for an *explanation* relation while a reference to the object should lead to a preference for a *result* relation. If readers make their choice of relation based on the verb alone, then these preferences should affect the reading time of the verb. Alternatively, if readers make their choice towards the end of the sentence, these preferences should affect the reading time of the final part of the sentence (i.e., the verb complement in the materials used in Experiment 4). The results supported the hypothesis that participants form preferences for specific relations based on the referent of unambiguous pronouns. These preferences affected the verb complement but not the verb, suggesting that participants make their choice of coherence relation only after reading a substantial portion of the second sentence. Moreover, when the pronoun was ambiguous, a general subject bias was observed rather than an effect that depended on the nature of the coherence relation. This is congruent with general pronoun interpretation biases found in the

literature and specifically with those reported by Kehler et al. (2008). This suggests that readers form expectations based on the general distribution of referents and coherence relations – when the referent is not known, readers follow the general distribution of referents whereas when the referent is known they form more specific expectations that are based on the distribution of relations for the specific referent.

While most of the results from these experiments followed my predictions, there were several notable results that were not predicted and merit further research. The most important deviation from my predictions occurred in Experiment 2, where the temporal connectives showed the same pattern of disambiguation as the causal connectives. I believe this is likely because no alternative source of information was available. It would therefore be interesting to explore cases in which two different sources of information suggest competing interpretations. For instance, it might be possible to contrast causal information about the causes and effects of events with information regarding the topic of the narrative (e.g., who the subject of a story is). To some extent, Wolfe et al. (2004) and Kehler et al. (2008, Experiment 1) have evidence that bear on this question. However, the cases they explore contrast semantic information with linguistic information (e.g., syntactic parallelism) whereas the sources of information that might distinguish different causal and temporal relations are more likely to be based in world knowledge.

Also in Experiment 2, participants' responses to sentences using the connective "after" were more similar to their responses to sentences using the connective "because" than the equivalent pair of "before" and "and so" connectives. This suggests that participants might have perceived the connective "after" as carrying more causal force than "before". While there is little direct evidence regarding the causal implications of "before" and "after", Clark (1971) suggests

that the connective “before” is linguistically simpler than “after” because simple sentences employing it (e.g., “He jumped the gate before he patted the dog.”) have the same order as the temporal order of the events they describe. Consequently, the simpler nature of “before” might set it apart from other connectives. Likewise, the contrast between the temporal order of events and their order in a sentence involving after (e.g., “He jumped the gate after he patted the dog”) might lead readers to consider more complex interpretations, such as ones that involve causal forces, are implied. Regardless, more evidence is needed to clarify the differences observed between the two connectives before any conclusions can be drawn.

Another unexpected result was found in Experiment 3. Participants rated the causal likelihood of the negative polarity sentence pairs in the low causal likelihood condition higher than their positive polarity counterparts. I believe that this is because the participants found it easier to connect negative consequences (i.e., nothing happened) than positive consequences. This difference was also mirrored by participants’ confidence in their choices of referent. Participants were also more confident in their choice of referent for the negative polarity low causal likelihood sentence pairs than their positive polarity counterparts.

Overall, the results from these experiments demonstrate that the relationship between causal coherence information and pronoun disambiguation is bi-directional. The first three experiments show that causal coherence information plays a role in the disambiguation of pronouns, while Experiment 4 shows that an unambiguous pronoun produces expectations regarding the nature of the causal information that will be presented.

These results therefore support Kehler’s (2002) view that pronoun disambiguation is sometimes a by-product of the process of coherence resolution. However, while Kehler is only concerned with the effect coherence resolution has on pronoun disambiguation Experiment 4

suggests that pronouns may also have an effect on the process of coherence resolution. This type of interaction is incompatible with processes that first identify a coherence relation and only then use this information for pronoun disambiguation. Instead, the results I reported argue for a process in which both coherence resolution and pronoun disambiguation occur simultaneously.

Earlier, I proposed one such process, based on Rips et al.'s (2006) *causal continuer model*. This process involves two stages – an evaluation stage and a decision stage. In the first stage, the evaluation stage, possible interpretation of the sentences are examined while in the decision stage the reader choose one of these interpretations over the others. Because both coherence relations and pronoun referents are determined when readers construct an interpretation of the sentences, this model can explain both why causal coherence information affects pronoun disambiguation and why an unambiguous pronoun can induce preferences regarding the coherence information.

Moreover, the results of Experiment 4 show that the unambiguous pronoun only affected the reading times of the latter part of the sentence. This suggests that the interpretations are most likely computed only after a substantial portion of the second sentence has been read rather than online as the reader progresses through the text. This is logical because the verb itself rarely provides sufficient information for generating a reasonable interpretation of the causal relations between the previous event and the current one. In fact, coherence relations and pronoun referents would frequently be ambiguous if not for the use of connectives. Experiment 2 provides some examples of sentence pairs in which such ambiguities are present.

In addition, there is prior evidence suggesting that causal coherence information plays a role only in the later stages of sentence processing. Stewart et al. (2000) explored the time course of the effect of implicit causal information on sentence processing. Their stimuli consisted of

sentences involving two clauses connected by an *explanation* relation (through the connective "because") that were similar in structure to those used in Experiment 2. They contrasted an account of online processing (the *Focusing* account) with one that suggested a later integration stage (the *Integration* account). Like the results of Experiment 4, their results show an effect of causal information only during the later stages of sentence processing.[22]

While I believe the model I outline, and its variations, provides the best fit for the results I described, it is important to consider the types of models these results suggest should be ruled out. Firstly, models that base pronoun disambiguation in narratives on the most likely event to occur will have difficulty explaining the results of Experiment 3. At the very least, these models will have to make allowances for deviations from these straightforward narrative continuations to account for the effect of negation. Such a model would also need to explain why participants were more confident in their responses when a causal connective was used than when a temporal connective was used.

Both of these difficulties can be resolved by models in which the process of pronoun disambiguation is based on the pragmatics of the text. However, in order for such a model to account for the results presented here it would generally need to assume that readers base their

---

[22] The integration account explored by Stewart et al. is relevant to the present study because the both reflect the integration of different information sources. The focus account explored by Stewart et al. requires that the verbs in question have implicit causality associated with them and that the relevant connective (i.e. "because") be present so that participants can use that connective to focus on the *cause* of the first sentence. Experiment 4 did not use such explicit connectives. Moreover, any *focal referent* would not vary when the pronoun was ambiguous, and would be overridden by an unambiguous pronoun. Consequently, the results of Experiment 4 bear little relevance on Stewart et al.'s focusing account.

disambiguation on a contrasting set of possible interpretations. The computation or identification of this set of possible interpretations corresponds to the evaluation stage of the model I proposed. Consequently, such models will likely be similar to the one I proposed.

The model I propose has implications to theories that propose various biases on sentence processing, such as theories of implicit causality and consequentiality. While it is possible that implicit causality information is processed separately from the coherence relation itself, I find it more plausible to assume that the role implicit causality information plays is in providing basic biases that influence the process of coherence resolution. Intuitively, implicit causality information can provide a baseline when evaluating the plausibility of a specific interpretation. For example, when the first sentence involves a verb with an implicit causal bias towards the subject (e.g., *apologize*) a connective that marks an *explanation* relation would bias the reader to give precedence (i.e., a higher baseline evaluation score) to interpretations where the subject of the second sentence would be the subject of the first sentence. In this way various biases that are based in the pragmatics or semantics of the first event can be integrated into the interpretation of the two sentences as a unit.

The expectations generated by the unambiguous pronouns in Experiment 4 also fall into this category of biases. They form a basic set of expectations that guide the evaluation of inferences. Importantly, these baselines provide a starting point which later (and often more precise and concrete) information can override. However, overriding these preferences carries a cost in processing time as was seen in Experiment 4.

Another body of research that intersects with the topic of this dissertation regards the computation of bridging inferences. As mentioned earlier, the concept of bridging inferences seems related to that of coherence relations. But whereas bridging inferences are context specific

(and often explored through the priming of specific words related to the content of the inference), coherence relations are generic relations that exist between many different pairs of sentences.

It is possible that the computation of bridging inferences is triggered by the process of coherence resolution when none of the immediate relations cross the plausibility threshold. For instance, I suspect that the longer reading times observed for the low causal likelihood sentence pairs in Experiment 3 is largely due to participants' attempts at generating reasonable bridging inferences to explain the gap in coherence between the first and second sentence.

This model also shares some features with Kintsch's (1988, 1998) Construction-Integration model. Most notably, both models involve two distinct phases. In the first phase of Construction-Integration the reader elaborates on the propositions in the text. Likewise, in my proposed model the reader generates fleshed-out interpretations of the sentence pairs. In the second phase of Construction-Integration the reader selects the most relevant elaborations and integrates them into their representation of the text. This second phase is similar to the selection of one of the interpretations.

However, there are important differences between the models. Firstly, Construction-Integration uses a spreading activation mechanism to generate bridging inferences from world knowledge. In contrast, the model I propose evaluates a limited set of possible interpretations by removing specific types of ambiguities inherent in the language, such as coherence relations and pronoun referents. Secondly, while Construction-Integration selects a small subset of the generated inference for inclusion in the representation of the text, the selection process of my model is much more stringent. Only a single interpretation is selected. These differences follow from the difference between the bridging inferences that are the focus of the Construction-Integration model and the coherence relations which are the basis for the model I described. As

described above, bridging inferences are context-based and their number is not necessarily constrained whereas coherence relations form a pre-existing set from which a single relation is presumably selected for a specific pair of discourse segments. Nevertheless, I believe the two theories are essentially compatible and complementary and it might be instructive to try and combine them into a unified model in the future.

Finally, my model borrows heavily from the literature on discourse structure and coherence relations. However, I limited my exploration to causal relations which form only one type of information that can be used to form coherence relations. Extending the model to account for other sources of information should be relatively straightforward. For instance, Kehler's *resemblance* relations involve parallel structures between the sentences. Therefore, the evaluation of this relation should focus on identifying such parallels.

One aspect of the model that will require careful consideration in such an extension is the comparison of disparate sources of information. It is not immediately clear how a reader might compare an interpretation involving cause effect information with one that involves parallel information. Consider that the pronoun "he" in (20) appears relatively unambiguous in referring to John and is similar to its interpretations in (21a). However, when the connective "because" is used, as in (21b), the referent of the pronoun shifts. It is obvious that readers can entertain both options and choose among them, but it is less obvious how to properly adjust for comparison across such different sources of information.

(20)

a. John met Mark.

b. He knows Barry.

(21)

a. John met Mark and he knows Barry.

b. John met Mark because he knows Barry.

The existence of different sources of information brings up another important question. Thus far I have argued that each interpretation that readers consider involves a single coherence relation. However, as the example above demonstrates, sometimes different sources of information advocating for different coherence relations are present in the same sentence. If these different sources of information are in harmony rather than conflict, it should be possible for readers to entertain an interpretation that draws its plausibility from more than one source of information simultaneously. This would suggest that it might be possible for more than a single coherence relation to be inferred between a pair of sentences. Nevertheless, this will not have a substantial effect on the model I outlined as it assumes that readers generate interpretations of the discourse rather than individual coherence relations between sentences.

In this work I have also made several important assumptions regarding the nature of coherence relations. One assumption in particular bears examining and that is the assumption that coherence relations form a small and finite set. Many theories of coherence relations, such as Rhetorical Structure Theory (Mann & Thompson, 1988) and Segmented Discourse Representation Theory (Asher & Lascarides, 2003) provide a large set of possible relations. In contrast, both Kehler (2002) and Sanders et al. (1992) argue for smaller sets. This matter is complicated further by the similarities between coherence relations and bridging inferences, both in function and representation.

However, I will argue that for coherence relations to be cognitively useful, the set of such relations contemplated by readers at each point in the reading process must be reasonably small, likely no more than a handful. Because this dissertation dealt with only a few, highly salient, coherence relation, this issue did not arise. However, if coherence relations are to become true members of the psycholinguistic arsenal, this issue needs to be resolved.

There are several possible avenues that can provide such a small set. The first, and most obvious, is that there are only a few coherence relations. This might be the case, for example, if coherence relations find their foundations in logic operators (e.g., Kehler, 2002). Another possibility is that different relations are ranked based on some metric of likelihood and only a handful are contemplated at a time. This method is related to the proposed non-monotonic logic Asher and Lascarides (2003) suggested for coherence relations. However, neither of these approaches can easily accommodate the flexibility of bridging inferences or suggests when bridging inferences might be called upon.

An intriguing alternative is that coherence relations are in a sense a subset of bridging inferences that are so frequently called upon that they have somehow separated themselves and appear to be an abstract class of their own. Such a separation can be qualitative or quantitative. In the first case, coherence relations are an abstracted representation over bridging inferences in which specific families of bridging inferences are identified using the same relation. One mechanism that might account for such a phenomenon might be conventionalization. This notion is supported by the existence of specific connectives that encapsulate the meaning of these relations (e.g., ‘because’, ‘therefore’).

In the second case, no special process or distinction needs to be made between bridging inferences and coherence relations, except that the bridging inferences that form coherence

relations are so frequently called upon that they become more fluent or better practiced in comparison with less frequent bridging inferences. The main advantage of the latter case is that it can easily be thought of as a natural result of exposure to language. Nevertheless, it is also possible to argue that such frequently occurring inferences would quickly result in conventionalization. I believe that this is an important and interesting question that needs to be addressed by future studies.

4 Summary

This dissertation deals with the role causal information plays in the identification of coherence relations and the disambiguation of pronouns. Central to the theme of this dissertation is Kehler's (2002) proposal that pronoun disambiguation is often a byproduct of identifying the relevant coherence relation. This interrelation makes it possible to examine the processing of coherence relations by looking at this byproduct.

Experiment 1 demonstrated that causal information plays a significant role in pronoun disambiguation for a specific causal relation – *result*. Experiment 2 extended this result by showing that causal information does not play as significant a role in temporal relations where the role of causal information is implicit rather than explicit. Experiment 3 then shifted focus and examined how this process might deal with counterfactual causality as exhibited by negation, showing that participants are just as likely to use causal information regarding when an effect fails to occur as they are for simple causes and effects. Finally, experiment 4 provided evidence

that readers use information regarding referents in the discourse to form expectations regarding the type of causal relation the text is likely to employ.

As a whole, the results presented here suggest that the processes underlying the determination of coherence relations and pronoun referents are intertwined. On the one hand, causal coherence affects pronoun disambiguation (Experiments 1-3). On the other hand, an unambiguous pronoun creates expectations regarding the type of coherence relation resulting in decreased reading times when these expectations are met (Experiment 4). This bi-directional relationship paints a picture that is somewhat more complex than that proposed by Kehler (2002). Pronoun disambiguation is not only an occasional by-product of establishing discourse coherence, it can also lead to expectations regarding the type of coherence relation used in the text. Nevertheless, the results of Experiment 4 demonstrated such expectations only when the pronoun was unambiguous. It is possible that such expectations are only generated when a pronoun is unambiguously tied to a referent. Consequently, it seems plausible that the effects observed in this dissertation require that *either* the coherence relation or the pronoun be determined before they might exert their influence on the other.

While there is an extensive literature on the causal implications of verbs, this literature generally focuses on cases where the coherence relation itself is unambiguous and is often explicitly conveyed by means of a connective. There is little research investigating the role causal information plays in *establishing* that two sentences are indeed causally related (rather than related in some other fashion). This is not entirely surprising because the psycholinguistic literature has only recently started to explore the notion that a small set of relationships might constrain the possible ways in which adjacent sentences might interact.

This dissertation aimed to extend the body of research that explores this relationship in the context of causal information. However, much work remains to be done in this respect. For one thing, the studies described deal only with the role causal information plays in coherence. As Kehler and others suggest, other coherence relations are likely to involve different types of reasoning. For instance, the role of parallelism and analogical relations may offer interesting insights of its own into the processing of discourse. Specifically, it is of interest whether readers actually decide to endorse one type of information over another (i.e., whether they might decide that a specific pair of sentences is connected via *result* but not *parallel* even when both might be applicable).

Secondly, while there are a few suggestions for computational models of coherence relations (e.g., Asher and Lascarides, 2003), there are as of yet no proposals for a psycholinguistic model of this process. Such a model should be designed so that it integrates well within the scope of general models of discourse comprehension, such as Construction-Integration (Kintsch, 1988, 1998) or Landscape (van den Broek, et al., 1999).

Finally, the relationship between bridging inferences and coherence relations needs to be studied. They seem to perform a similar role in the establishment of local coherence but, as noted earlier, there are some important differences between them. Moreover, because of the generality of coherence relations and the specificity of bridging inferences it might be that they are actually different levels of the same cognitive representation.

Overall, I believe that the notion of a small set of general relationships that might bind sentences together into a coherent whole is appealing and worth of further study. In this work I attempted to extend the small field that explores these relationships empirically and to propose

how psycholinguistic studies might be able to further explore the relationship between coherence relations and well studied cognitive processes.

References

Arroyo, F. V. (1982) Negatives in context. *Journal of Verbal Learning and Verbal Behavior*, 21: 118–126.

Asher N. and Lascarides A. (2003) *Logics of Conversation*. New York: Cambridge University Press.

Au, T. K. (1986) A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. *Journal of Memory and Language*, 25: 104-122.

Clark, E. (1971) On the Acquisition of the Meaning of Before and After. *Journal of Verbal Learning and Verbal Behavior*, 10: 266-275.

Crinean, M. & Garnham, A. (2006) Implicit causality, implicit consequentiality and semantic roles. *Language and Cognitive Processes*, 21: 636-648.

Brown, R. & Fish, D. (1983b) The psychological causality implicit in language. *Cognition*, 14: 237-273.

Fletcher, C. R. (1984) Markedness and topic continuity in discourse processing. *Journal of Verbal Learning & Verbal Behavior*, 23: 487-493.

Garvey, C. M., Caramazza, A. & Yates, J. (1976) Factors influencing assignments of pronoun antecedents. *Cognition*, 3: 227-243.

Gordon, P. C. & Scearce, K. A. (1995) Pronominalization and discourse coherence, discourse structure and pronoun interpretation. *Memory & Cognition*, 23: 313-323.

Grice, H. P. (1975) *Studies in the Way of Words*. Boston, MA: Harvard University Press.

Grosz, B. J., Joshi, A. K. & Weinstein, S. (1983) Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*. Cambridge, MA: ACL.

Grosz, B. J., Joshi, A. K. & Weinstein, S. (1995) Centering: A Framework for Modelling the Local Coherence of Discourse. *Computational Linguistics*, 21: 203-225.

Hobbs, J.R. (1978) Resolving pronoun references. *Lingua,* 44: 311-338.

Hobbs, J.R. (1979) Coherence and Coreference. *Cognitive Science,* 3: 67-90.

Hobbs, J.R. (1985) *On the coherence and structure of discourse*. Technical Report CSLI-85-37, Center for the Study of Language and Information, Leland Stanford Junior University, Stanford, California.

Kaup, B., Zwaan, R. A., & Lüdtke, J. (2007). The experiential view of language comprehension: How is negated text information represented? In F. Schmalhofer & C.A. Perfetti (Eds.), *Higher level language processes in the brain: Inference and comprehension processes* (pp. 255-288). Mahwah, NJ: Erlbaum.

Keenan J. M., Baillet S. D., & Brown P. (1984) The Effects of Causal Cohesion on Comprehension and Memory. *Journal of Verbal Learning and Behavior,* 23: 115-126.

Kehler, A. (2002) *Coherence, Reference, and the Theory of Grammar*. Stanford, CA: CSLI Publications.

Kehler, A. (2004) Discourse Coherence. in Horn, L. R. & Ward, G. (eds.) *Handbook of Pragmatics*. Basil Blackwell, Oxford.

Kehler, A., Kertz, L., Rohde, H., & Elman J. L. (2008) Coherence and Coreference Revisited. *Journal of Semantics*, 25: 1-44.

Kintsch, W. (1988) The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, 95: 163-182.

Kintsch, W. (1998) *Comprehension: A paradigm for cognition.* Cambridge, England: Cambridge University Press.

Luce, R. D. (1977) The choice axiom after twenty years. *Journal of Mathematical Psychology*, 15: 215–233

Mann, W.C. & Thompson, S. (1988) Rhetorical Structure Theory: towards a functional theory of text organization. *Text*, 8: 243-81.

McKoon, G., & Ratcliff, R. (1992) Inference During Reading. *Psychological Review*, 3: 440-466.

McNamarra, D. S., Kintsch, E., Songer, N. B. & Kintsch, W. (1996) Are good texts always better? Interactions of text coherence, background knowledge, and levels of understanding in learning from text. *Cognition and Instruction*, 14: 1-43.

Myers J. L., Shinjo M. & Duffy S. A. (1987) Degree of Causal Relatedness and Memory. *Journal of Memory and Language*, 26: 453-465.

Pickering M. J. & Majid, A. (2007) What are implicit causality and consequentiality? *Language and Cognitive Processes*, 22: 780-788.

Polanyi, L. (1985) A Theory of Discourse Structure and Discourse Coherence in *Proceedings of the 21st Regional Meeting of the Chicago Linguistic Society*. Chicago: University of Chicago.

Polanyi, L., van den Berg M., & Ahn D. (2003) Discourse Structure and Sentential Information Structure. *Journal of Logic, Language, and Information, 12*, 337-350.

Rips L., Blok S. & Newman G. (2006) Tracing the Identity of Objects. *Psychological Review,* 113: 1-30.

Rohde H., Kehler A., & Elman J. (2006) Event Structure and Discourse Coherence Biases in Pronoun Interpretation. in R. Sun & N. Miyake, eds. *Proceedings of the 28th Annual Conference of the Cognitive Science Society.*

Rudolph, U. & Forsterling, F. (1997) The Psychological Causality Implicit in Verbs: A review. *Psychological Bulletin,* 121: 192-218.

Sagi, E. (2004) The priming of coherence relations: How one sentence can affect another. Unpublished Master's thesis.

Sagi, E. (2006) Context and the Processing of Discourse: Priming and Genre Effects on Discourse Comprehension. In R. Sun & N. Miyake, eds. *Proceedings of the 28th Annual Conference of the Cognitive Science Society*.

Sanders, T.J.M., Spooren, W.P.M., & Noordman, L.G.M. (1992) Towards a taxonomy of coherence relations. *Discourse Processes,* 15: 1-35.

Schmalhofer F., McDaniel, M. A., & Keefe D. (2002) A unified model for predictive and bridging inferences. *Discourse Processes,* 33: 105-132.

Singer, M., Halldorson M., Lear, J.C., & Andrusiak, P. (1992) Validation of causal bridging inferences in discourse understanding. *Journal of Memory and Language,* 31: 507-524.

Stewart, A. J., Pickering, M. J. & Sanford, A. J. (1998). Implicit consequentiality. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Stewart, A. J., Pickering, M. J. & Sanford, A. J. (2000) The time course of the influence of implicit causality information: Focusing versus integration accounts. *Journal of Memory and Language*, 42: 423-443.

Tukey J. W. (1977) *Exploratory Data Analysis*, Reading, Mass: Addison-Wesley.

van den Broek, P., Young, M., Tzeng, Y., & Linderholm, T. (1999). The landscape model of reading. In H. van Oostendorp & S. R. Goldman eds., *The construction of mental representations during reading*. Mahwah, NJ: Erlbaum.

Wason, P. C. (1961) Response to affirmative and negative binary statements. *British Journal of Psychology*, 52: 133–142.

Wason, P. C. (1965) The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior*, 4: 7–11.

Wolfe F., Gibson, E. & Desmet T. (2004) Discourse coherence and pronoun resolution. *Language and Cognitive Processes*, 19: 665-675.

Appendix A: Materials used in Experiment 1

Head Sentence	Subject Continuation	Object Continuation
Larry shot Andrew in front of a crowd.	He was sentenced to death.	He was buried in a casket.
Albert kidnapped Walter.	He collected $100,000 in unmarked bills.	He escaped.
Fred rescued Trevor from a burning building.	He was heralded as a hero on the 5 o'clock news.	He was taken directly to the nearest hospital.
Bart ran over Aaron with a motorcycle.	He quickly fled the scene.	He limped for weeks afterwards.
Mike stabbed Baxter with a sword.	He confessed to a priest later that week.	He recovered after a few weeks
Julian accused Glen of stealing a computer.	He notified the police.	He called a defense attorney.
Cliff pushed Abe into a pool.	He watched the scene unfold.	He swam to the nearest ladder.
Edgar fired at Paul with a handgun.	He was arrested by the police.	He was dead within minutes.
Hank punched Sam in the face.	He apologized profusely.	He dropped down.
Charlie beat Jay in a game of chess.	He won the tournament.	He came last in the tournament.
Dwayne visited Randy.	He rang the doorbell.	He opened the door.

Head Sentence	Subject Continuation	Object Continuation
Nick cut Bob with a knife.	He expressed regret.	He bled for a few minutes
Brandon took Gabriel on a tour of the city.	He asked for payment	He learned quite a bit
Adam called Kevin using a cellphone.	He got a busy signal.	He answered without delay.
Oscar sold Bradley a diamond ring.	He made a nice profit.	He charged the purchase to a credit card.
Mark waited on Miles at a restaurant.	He got a respectable tip.	He left an ok tip.

Appendix B: Materials used in Experiment 2

Group	Head Sentence	Tail Sentence
1	Larry asked Andrew for some study tips	he aced the test.
	Albert tutored Walter for several hours	he got an A on the midterm.
2	Fred followed a stock tip from Trevor	he made a bundle on Wall Street.
	Bart gave investment advice to Aaron	he doubled his money in less than a month.
3	Mike paid Baxter for a poker lesson	he came first in a tournament.
	Julian taught Glen how to prevail in checkers	he won a cash prize at a competition.
4	Cliff took employment advice from Abe	he was hired by an HR company.
	Edgar offered some job leads to Paul	he signed up with a temp agency.
5	Hank scared Sam	he said very little.
	Charlie disliked Jay	he kept his distance.
6	Dwayne blackmailed Randy	he kept quiet.
	Nick apologised to Bob	he didn't raise the issue.
7	Brandon detested Frank	he seemed agitated.
	Adam troubled Kevin	he didn't want to open up.
8	Oscar loathed Miles	he commited a faux pas.
	Mark infuriated Bradley	he caused problems.

Appendix C: Materials used in Experiment 3

Changes made for negative polarity items given in parentheses.

Head Sentence	Subject Continuation	Object Continuation
Larry shot Andrew in front of a crowd.	He was (not) sentenced to death.	He was (not) buried in a casket.
Albert kidnapped Walter.	He collected (did not collect) $100,000 in unmarked bills.	He escaped (did not escape).
Fred rescued Trevor from a burning building.	He was (not) heralded as a hero on the 5 o'clock news.	He was (not) taken directly to the nearest hospital.
Bart ran over Aaron with a motorcycle.	He fled (did not flee) the scene.	He developed (did not develop) a limp.
Mike stabbed Baxter with a sword.	He confessed (did not confess) to a priest later that week.	He required (did not require) a week to recover.
Julian accused Glen of stealing a computer.	He notified (did not notify) the police.	He called (did not call) a defense attorney.
Cliff pushed Abe into a pool.	He watched (did not watch) the scene unfold.	He swam (did not swim) to the nearest ladder.

Head Sentence	Subject Continuation	Object Continuation
Edgar fired at Paul with a handgun.	He was (not) arrested by the police.	He died (did not die).
Hank punched Sam in the face.	He apologized (did not apologize).	He dropped (did not drop) down.
Charlie beat Jay in a game of chess.	He won (did not win) the tournament.	He came (did not come) last in the tournament.
Dwayne visited Randy.	He rang (did not ring) the doorbell.	He opened (did not open) the door.
Nick cut Bob with a knife.	He expressed (did not express) regret.	He bled (did not bleed) out.
Brandon took Gabriel on a tour of the city.	He asked (did not ask) for payment	He learned (did not learn) a lot.
Adam called Kevin using a cellphone.	He got (did not get) a busy signal.	He answered (did not answer) forthwith.
Oscar sold Bradley a diamond ring.	He made (did not make) a nice profit.	He charged (did not charge) the purchase to a credit card.
Mark waited on Miles at a restaurant.	He got (did not get) a respectable tip.	He left (did not leave) an ok tip.

Appendix D: Materials used in Experiment 4

Different conditions in the experiment used different gender assignments for the subject and the object. The pronoun presented to participants matched the gender of its correct referent.

Group	Head Sentence	Type of Tail	Tail Sentence
1	Larry alleged that Donna stole an expensive car.	Result Subject Continuation	He reported it to the police.
	Albert claimed that Brenda shoplifted a pricey suit.	Result Object Continuation	She sold it to a fence.
	Fred stated that Deborah forged a painting.	Explanation Subject Continuation	He saw it happen.
	Bart declared that Carol hijacked a motorcycle.	Explanation Object Continuation	She needed the money.
2	Mike decided to call Angela.	Result Subject Continuation	He looked up the number in the phonebook.
	Julian wanted to phone Lisa.	Result Object Continuation	She answered the call promptly.
	Cliff desired to speak with Betty.	Explanation Subject Continuation	He needed a favor.
	Edgar felt like talking to Kimberly.	Explanation Object Continuation	She failed to show up for work.

Group	Head Sentence	Type of Tail	Tail Sentence
3	Ruth said Andrew was lying.	Result Subject Continuation	She told the true story.
	Michelle implied Walter was untruthful.	Result Object Continuation	He denied the accusation.
	Sarah stated Trevor was stretching the truth.	Explanation Subject Continuation	She knew what happened.
	Elizabeth suggested Aaron was fibbing.	Explanation Object Continuation	He wanted to avoid punishment.
4	Melissa gave an urgent memo to Baxter.	Result Subject Continuation	She waited for an immediate reply.
	Barbara handed an important note to Glen.	Result Object Continuation	He read it carefully.
	Patricia delivered a critical message to Abe.	Explanation Subject Continuation	She was asked to.
	Helen brought a last-minute contract to Paul.	Explanation Object Continuation	He requested it.

Group	Head Sentence	Type of Tail	Tail Sentence
5	Hank flung a dart at Sharon.	Result Subject Continuation	He laughed.
	Robert hurled a tomato at Cynthia.	Result Object Continuation	She ducked.
	Dwayne threw am egg at Karen.	Explanation Subject Continuation	He repaid an earlier incident.
	Nick tossed a water balloon at Mary.	Explanation Object Continuation	She offended the entire school earlier that day.
6	Brandon faxed a document to Shirley.	Result Subject Continuation	He received notification when it was delivered.
	Adam mailed a letter to Maria.	Result Object Continuation	She got it the next morning.
	Oscar shipped a package to Linda.	Explanation Subject Continuation	He promised to send it as soon as possible.
	Mark transmitted an article to Susan.	Explanation Object Continuation	She asked for it.

Group	Head Sentence	Type of Tail	Tail Sentence
7	Nancy roasted a chicken for Sam.	Result Subject Continuation	She served it as a surprise.
	Laura cooked some eggs for Jay.	Result Object Continuation	He enjoyed eating it.
	Amy grilled a fish for Randy.	Explanation Subject Continuation	She thought it would make a nice gesture.
	Sandra baked a cake for Bob.	Explanation Object Continuation	He asked for it.
8	Jessica startled Gabriel by buying a new car.	Result Subject Continuation	She apologized.
	Jennifer scared Kevin by running into traffic.	Result Object Continuation	He felt angry.
	Dorothy shocked Miles by throwing a loud party.	Explanation Subject Continuation	She acted impulsively.
	Margaret upset Bradley by quitting without notice.	Explanation Object Continuation	He never expected it.