

The Time Course of the Influence of Implicit Causality Information on Resolving Anaphors

Andrew James Stewart

Department of Psychology
University of Glasgow.

Submitted for the Degree of PhD to the Higher Degrees Committee of
the Faculty of Social Sciences, University of Glasgow.

January 1998

(c) Andrew James Stewart

ProQuest Number: 13818590

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 13818590

Published by ProQuest LLC (2018). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

GLASGOW UNIVERSITY
LIBRARY
1107S (copy 1)

GLASGOW
UNIVERSITY
LIBRARY

Abstract

This thesis asks whether initial anaphor processing proceeds in a restricted manner with reference only to a well defined set of information or whether it is the case that all factors that are potentially relevant for resolving an anaphor exert a processing influence at the same time. In an attempt to adjudicate between these possibilities, we focus on the nature of the processing influence of implicit causality information on anaphor resolution.

Following a summary in Chapter 1 of issues concerning possible cognitive architectures and a review in Chapter 2 of previous work on anaphor resolution, we propose a two-stage model of anaphor resolution. We propose that the first stage involves co-indexation between anaphor and antecedent and is informed by low-level factors. We claim this stage behaves in a modular or restricted manner. The second stage involves integrative processing and behaves in a nonmodular or unrestricted manner. We suggest that it is at this second stage of processing that implicit causality influences anaphoric processing.

Implicit causality (Garvey and Caramazza, 1974) is a property associated with a particular set of verbs which, in sentence fragments such as (1) and (2) below, influences interpretation of the ambiguous pronoun.

(1) John fascinated Bill because he ...

(2) John blamed Bill because he ...

The verb 'fascinate' is classified as an NP1 biasing verb as it biases towards the character occupying the first Noun Phrase as the locus of cause. Similarly, the verb 'blame' is an NP2 biasing verb as it biases toward the character occupying the second Noun Phrase as the locus of cause. Readers prefer to interpret the pronoun as coreferential with the character predicted by the verb. Previous work has demonstrated the influence of implicit causality in both language production (e.g. Garvey and Caramazza, 1974) and comprehension

(e.g. Caramazza, Grober, Garvey and Yates, 1977). A reading penalty arises when the information in the subordinate clause conflicts with the verb bias, i.e. when there is a mismatch between implicit and explicit causes, as in example (3).

(3) John blamed Bill because he hated Bill.

However, several major methodological criticisms can be raised against previous work examining the influence on comprehension of implicit causality. Variations in factors such as sentence length, sentence plausibility and non-homogeneity of strength of verb biases may have confounded previous research. Experiments 1a and 1b in this thesis were used to create materials controlled for plausibility and strength of bias. Average length of the experimental sentences was equated across conditions. From an initial set of 50 verbs examined in Experiments 1a and 1b, we selected 24 that were strongly biasing and of equivalent plausibility for each cause.

An initial self-paced reading experiment (Experiment 2) demonstrated an implicit causality congruency effect with our materials on whole sentence reading times. Experiments (3) and (4) involved presenting the experimental materials in two halves, with the split occurring following the anaphor (see (4) and (5) below). We added an additional between experiment factor of question type which encouraged either deep or shallow processing.

(4) John fascinated Bill because he/John was full of interesting stories.

(5) John fascinated Bill because he/Bill was easily entertained.

Each sentence was presented as two fragments with the split following the anaphor. If implicit causality exerts an early influence on processing we would expect to find evidence of the congruency effect on reading times to the first fragment. The only effect we found on reading time to the first fragment was a repeat name penalty resulting from repetition of the first mentioned character's name. This did not interact with verb bias suggesting implicit causality does not influence interpretation of the anaphor when it is

first encountered. Reading times to fragment 2 showed an effect of implicit causality. Our between experiment manipulation led to a reduction in the strength of the implicit causality congruency effect under circumstances where shallow processing was encouraged. In other words, the relative difficulty associated with reading sentence continuations going against the bias of the verb was reduced when readers were engaged in shallower reading. The depth of processing manipulation had no effect on the magnitude of the repeat name penalty. This suggests that these phenomena may be arising from processing at different stages within the system.

In Experiment (5) we attempted to separate factors arising as a result of the information presented to the reader in the experimental sentences from those factors which may have arisen as a result of the manner of this presentation. Information previously presented in a main-subordinate clause sentence was presented as two separate sentences (see examples (6) and (7) below).

(6) John fascinated Bill. This was because he/John was full of interesting stories.

(7) John fascinated Bill. This was because he/Bill was easily entertained.

We found the implicit causality congruency effect in the pronoun conditions but not in the name conditions. We suggest this may be due to the reader interpreting the repeat name anaphor as a thematic shift signal (cf. Vonk, Hustinx and Simons, 1992) or perhaps as a result of increased informational load faced by the system. We did not find a repeat name penalty associated with repeating the first mentioned character's name.

Experiments 6 and 7 examine the influence of implicit causality under conditions where there is a gender differentiation between the two characters. Experiment 6 employed the same self-paced reading methodology as was used elsewhere in the thesis while Experiment 7 employed an eye-tracking methodology.

(8) John fascinated Mary because he was full of interesting stories.

(9) John fascinated Mary because she was easily entertained.

Gender information alone is sufficient to identify the pronominal referent. We found however that gender information is not used when a pronoun is first encountered.

Experiment 8 is an attempt to adjudicate between two conflicting positions in the literature concerning the time course of the influence of implicit causality. McDonald and MacWhinney (1995) propose that implicit causality influences processing as soon as a pronoun is encountered while Garnham, Traxler, Oakhill and Gernsbacher (1996) propose that the influence occurs during integration. A number of differences exist between the experimental structures of McDonald and MacWhinney and Garnham et al. We re-constructed our materials along the lines of those examined by McDonald and MacWhinney but still found no evidence for an early influence of implicit causality. We argue that their finding arises as a result of an experimental confound.

Experiments 8a and 8b provide us with an off-line measure of another type of verb bias we refer to as implicit consequentiality (see examples (10) and (11)). Paralleling Experiments 1a and 1b, we used Experiments 8a and 8b to create materials controlled for plausibility and strength of verb bias.

(10) Because Harold dreaded Justin, Justin/he / steadfastly refused to go back to school.

(11) Because Harold dreaded Justin, Harold/he / was told to try acting less aggressively.

In Experiment 10 we found a similar pattern of data to Experiment 3. Verb semantics in the form of implicit consequentiality influences anaphoric processing during integration. On reading time to fragment 1 we found a repeat name penalty of the same type as has already been reported.

When a gender contrast is present (cf. Experiments 6 and 7), we find evidence that gender information is used immediately if it can unambiguously identify the pronominal referent. In line with our examinations of implicit causality, we find evidence that implicit consequentiality influences anaphoric processing during integrative processing.

So then, verb semantics influences processing at the second stage of our proposed model. The first stage of anaphor resolution is informed by low level factors. Gender information only exerts an initial processing influence under conditions where the pronoun appears as the grammatical subject of a main clause (i.e. in implicit consequentiality type sentences) and where it is sufficient to identify the pronominal antecedent. Verb semantics influences anaphoric processing at no point earlier than integration.

Contents

Chapter 1

1.0 Overview	17
1.1 Introduction	18
1.2 Theoretical Positions	18
1.2.1 Processing Architectures	18
1.2.2 Modularity	19
1.2.3 Serial and Parallel Processing	23
1.2.4 Modularity applied to language processing	24
1.3 Parallels between Accounts of Parsing and Accounts of Anaphor Resolution	26

Chapter 2

2.1 Introduction	30
2.2 Anaphora	31
2.2.1 Differences between Anaphors	31
2.2.1 What is focus ?	32
2.2.2 What determines focus ?	32
2.2.3 How does focus influence selection and interpretation of anaphoric form ?	34
2.2.4 The Influence of Focus on Anaphoric Processing	35
2.2.5 How anaphors access levels within the Discourse Model	36
2.2.6 The mechanism of thematic shifting explained	38
2.3 The Time Course of Anaphoric Processing	39
2.3.1 The Immediacy Hypothesis	39
2.3.2 Differences in Processing of Pronouns and Noun Phrase Anaphora	41
2.4 The Probe Task	47
2.4.1 Problems with the Probe Task	48
2.4.2 Probe Task Evidence	50
2.4.3 Are Pronouns Always Resolved ?	55
2.5 The First Mention Privilege	55
2.5.1 A non-psycholinguistic account of the First Mention Privilege	58
2.6 Structurally informed accounts of pronoun resolution	59
2.6.1 Parallel Function Strategy	60

2.6.2 Subject Assignment Strategy	61
2.7 Structurally Informed Accounts as Processing Theories	62
2.7.1 Crawley, Stevenson and Kleinman (1990)	62
2.7.2 Smyth (1994)	63
2.7.3 Stevenson, Nelson and Stenning (1995)	66
2.8 Centering Theory	67
2.9 Explanation at the Level of Thematic Roles ?	69
2.10 The Influence of Non-Structural Factors on Anaphor Resolution	72
2.10.1 When does semantics influence anaphoric interpretation?	72
2.10.2 When does gender information influence anaphoric processing?	73
Summary	76
2.11 A Two-stage Model of Anaphor Resolution	77
Chapter 3	
3.1 Introduction	83
3.2 Implicit Causality	83
3.3 Implicit Causality from Non-processing Perspectives	84
3.3.1 The Whorfian Hypothesis - What language reveals about cognition.	85
3.3.2 Attribution Theory	87
3.3.3 Conclusion	88
3.4 Evidence for the Influence of Implicit Causality Information in Language Production	89
3.5 Evidence for the Influence of Implicit Causality Information in Language Comprehension	91
3.5.1 Conditions Necessary for the Bias to be Manifested	92
3.5.2 Conclusion	94
3.6 The Time Course of the Influence of Implicit Causality Information in Language Comprehension	94
Conclusion	98
Summary	100
Chapter 4	
4.1 Introduction and Rationale	101

4.2 Problems with previous Empirical Work on Implicit Causality	101
Off-line Experiments	103
4.2.1 Experiment 1a	103
Method	103
4.2.2 Results	104
4.2.3 Experiment 1b	105
Method	107
4.2.4 Results	108
Experiments examining On-line processing	110
4.3 Experiment 2	111
4.3.1 Introduction and Rationale	111
Method	112
4.3.2 Results	114
4.3.3 Discussion	121
4.4 Experiments 3 and 4	123
4.4.1 Introduction and Rationale	123
Method	127
4.4.2 Experiment 3 Results	128
4.4.3 Experiment 4 Results	136
4.4.4 Comparing Experiments 3 and 4	139
4.4.5 Discussion	140
4.5 General Discussion	141
Chapter 5	
5.1 Introduction	144
5.2 Experiment 5	145
5.2.1 Introduction and Rationale	145
Method	147
5.2.2 Results	148
5.2.3 Discussion	155
5.3 Experiment 6	159
5.3.1 Introduction and Rationale	159
Method	161
5.3.2 Results	162
5.3.3 Discussion	167
5.4 Experiment 8	168
5.4.1 Introduction and Rationale	168

Method	171
5.4.2 Results	173
5.4.3 Discussion	179
5.5 General Discussion	181
Chapter 6	
6.1 Introduction and Rationale	184
6.2 Experiment 8	191
6.2.1 Introduction and Rationale	191
Method	191
6.2.2 Results	193
6.2.3 Comparing Experiments 3 and 8	198
6.2.4 Discussion	198
Chapter 7	
7.1 Introduction and Rationale	202
7.2 Experiment 9a	204
Method	204
7.2.1 Results	205
7.3 Experiment 9b	206
Method	208
7.3.1 Results	208
7.4 Experiment 10	211
7.4.1 Introduction and Rationale	211
Method	213
7.4.2 Results	214
7.4.3 Discussion	221
7.5 Experiment 11	224
7.5.1 Introduction and Rationale	224
Method	226
7.5.2 Results	227
7.5.3 Discussion	232
7.6 General Discussion	233
Chapter 8	
8.1 When does verb bias influence anaphoric processing ?	236
8.1.1 Implicit Causality	238
8.1.2 Implicit Consequentiality	241

8.1.3 Between Experiment Comparisons	243
8.2 Our Two-stage Account	244
References	248
Appendices	
Appendix 1 (Experiment 2,3 and 4 materials)	258
Appendix 2 (Experiment 5 materials)	265
Appendix 3 (Experiment 6 materials)	272
Appendix 4 (Experiment 7 materials)	278
Appendix 5 (Experiment 8 materials)	283
Appendix 6 (Experiment 10 materials)	291
Appendix 7 (Experiment 11 materials)	295
Appendix 8 (Filler materials)	299

List of Figures

Chapter 4

Figure 1 - anaphor type x referent	116
Figure 2 - verb bias x referent	117
Figure 3 - verb bias x referent (Name conditions)	118
Figure 4 - verb bias x referent (Pronoun conditions)	119
Figure 5 - main effect of referent (repeat name anaphor)	130
Figure 6 - verb bias x referent	131
Figure 7 - verb bias x referent (pronoun conditions)	132
Figure 8 - verb bias x referent (name conditions)	133
Figure 9 - referent x anaphor	134
Figure 10 - main effect of referent (repeat name anaphor)	137
Figure 11 - verb bias x referent (pronoun conditions)	138

Chapter 5

Figure 1 - anaphor type x referent	150
Figure 2 - verb bias x referent	151
Figure 3 - verb bias x referent (Pronoun conditions)	152
Figure 4 - verb bias x referent	166
Figure 5 - verb bias x referent	175
Figure 6 - verb bias x referent	177
Figure 7 - verb bias x referent	178
Figure 8 - verb bias x referent	179

Chapter 6

Figure 1 -verb bias x referent interaction	195
Figure 2 - verb bias x referent interaction (Pronoun conditions)	196

Chapter 7

Figure 1 - plausibility x NP	211
Figure 2 - referent main effect (repeat name anaphor)	216
Figure 3 - anaphor type x referent	218
Figure 4 - verb bias x referent (Pronoun conditions)	219
Figure 5 - verb bias x referent	229
Figure 6 - verb bias x referent	230

List of Tables

Chapter 4

Table 1 - Experiment 1a continuation results	104
Table 2 - Experiment 1b plausibility results combined with Experiment 1a continuations results	109
Table 3 - Reading times	115
Table 4 - Question response times	120
Table 5 - Reading times for fragment 1	129
Table 6 - Reading times for fragment 2	130
Table 7 - Question response times	135
Table 8 - Reading time for fragment 1	137
Table 9 - Reading time for fragment 2	138

Chapter 5

Table 1 - Reading times for first sentence	149
Table 2 - Reading times for second sentence	150
Table 3 - Question response times	154
Table 4 - Reading times for first fragment	164
Table 5 - Reading times for second sentence	165
Table 6 - Question response times	166
Table 7 - First-pass reading times for region 1	174
Table 8 - First-pass reading times for region 3	174
Table 9 - Total reading times for region 1	176
Table 10 - Total Reading times for region 3	176

Chapter 6

Table 1 - Reading times for first fragment	194
Table 2 - Reading times for second fragment	195
Table 3 - Question response times	197

Chapter 7

Table 1 - Experiment 9a continuation results	205
Table 2 - Experiment 9b plausibility results combined with Experiment 9a results	210
Table 3 - Reading times for first fragment	216
Table 4 - Reading times for second fragment	217
Table 5 - Question response times	220
Table 6 - Reading times for first fragment	228
Table 7 - Reading times for second fragment	229

Acknowledgements

I wish to thank Dr. Martin Pickering and Prof. Tony Sanford for supervising this thesis.

I also wish to thank Eugene Dawydiak and Kirsty Dalrymple for assistance in helping to run some of the experiments. Thanks to the following for discussion, advice and support : Aude Oliva, Asifa Majid, Holly Branigan, Roger van Gompel, Simon Garrod, Arlina Import, Frederic Gosselin, Frances Duffy and everyone else in the Psychology Department Language Group.

Also thanks to Petrea, Petrana, Jacqui and Joanna. Special thanks to James Osterberg.

This research was supported by a BBSRC studentship.

Declaration

I declare that this thesis is my own work carried out under normal terms of supervision.

CHAPTER 1

1.0 Overview

In this thesis we set out to examine how different types of information are used to resolve pronominal and repeat name anaphors. The research is motivated both by theoretical work on anaphora and by general processing perspectives which have been adopted by those working on parsing. Chapter 1 sets out several theoretical positions which have been central in guiding parsing research. We focus on those that have arisen as a result of adopting Fodor's Modularity thesis. We summarise the main points of this proposal and then set out what follows from considering anaphoric processing as operating in a modular fashion. We draw explicit parallels between anaphoric processing and accounts of parsing. Broadly speaking, accounts of parsing can be classified as those which are modular and those which are non-modular in nature. We apply this method of categorisation to possible models of anaphor processing.

Chapter 2 contains a general overview of theoretical concerns and existing process models of anaphor resolution. This chapter contains an examination of theoretical positions including the factors that influence the choice of anaphoric form, the role played by focus information on interpretation of anaphors and how different anaphoric forms interact differently with the reader's discourse model. From a processing perspective we outline existing accounts of anaphor resolution which are informed by both low and high level factors. Chapter 2 concludes with the proposal of a 2-stage model of anaphor resolution.

The central question we are interested in is the time course associated with the influence on anaphoric processing of structural and non-structural information. We focus on one particular type of non-structural information : a type of verb semantic information known as implicit causality. Chapter 3 contains a summary of previous empirical work examining this phenomenon.

Our experimental chapters (4-7) examine the on-line processing influence of verb semantics. Apart from one experiment (Experiment 7) where we adopt an eye-tracking methodology, we measure reading time using a self-paced reading technique and attempt to temporally separate the processing influence of low and high level factors.

In Chapter 8 we summarise our experimental findings and re-evaluate our 2-stage model of anaphor resolution.

1.1 Introduction

This thesis is an attempt at applying some of the theoretical and experimental principles which have been prevalent in the parsing literature to those aspects of the language system associated with resolving anaphors. This chapter sets out two important theoretical positions which have guided research on sentence processing over the last 20 years. The first, Fodor's Modularity thesis (1983), concerns the overall conception of how cognitive structure may be organised. This framework provides us with a way of empirically examining how certain cognitive processes operate. The second theoretical position outlines possible computational differences in the way in which those processes might operate. Roughly they may behave in a serial or parallel fashion. Initially we shall focus on the general position outlined by Fodor before describing the consequences of processing operating in a serial or parallel manner.

1.2 Theoretical Positions

1.2.1 Processing Architectures

Perhaps the most influential theoretical account concerning itself with the nature of the cognitive architecture is Fodor's Modularity thesis (Fodor, 1983). Any account of cognitive functioning is faced with the world knowledge problem. Roughly, the world knowledge problem is that it is not possible to formally capture the complex nature of world knowledge and its influence on processing. Modularity manages to get round this by proposing that initial processing operates only over a restricted, well defined set of information and that only at some

later point does world knowledge exert an influence. As this is absent during initial processing, if we know what information is used within a module it should be possible to explicate processing at the modular level without having to formalise world knowledge.

Modularity is a restricted processing account. Stated simply it proposes that only a restricted set of information is used within a module. For any process, consider that there is a large set of information which ultimately is used by that process. A restricted account simply states that only a well defined subset of this larger set is initially used by the system. Only later does the rest of this information exert a processing influence. If we know precisely what type of information is used by the system at this initial stage we should be able to formally capture how that information influences processing.

Alternatively, an unrestricted account proposes that all of the information that can exert a processing influence does so as soon as it becomes available. In other words there isn't an initial stage of the system during which only a particular well defined subset of information is used. All of the information available to the processor at any point in time is utilised by the system.

The following section includes a general outline of what follows from interpreting the cognitive system within the modularity framework. The initial summary is later followed by an instantiation of the modularity thesis with respect to sentence processing. We then draw parallels between accounts of parsing and accounts of anaphor resolution.

1.2.2 Modularity

Although the view that the cognitive system could be decomposed into separate processing domains was implicitly assumed for some time by many researchers within cognitive psychology (e.g. Forster 1979), Fodor's Modularity thesis (1983) set forth a precise characterisation of what followed from viewing the cognitive system as consisting of a discrete number of specifiable processing

components. In doing so, it revived an older notion that considered the difference between certain dimensions of cognitive functioning within the faculty psychology framework. Fodor takes the faculty psychology position as 'the view that many fundamentally different kinds of psychological mechanisms must be postulated in order to explain the facts of mental life' (Fodor 1983, p. 1). One possible dissociation between faculties is roughly that some can be considered vertical while others horizontal. Horizontal faculties are those aspects of cognitive functioning which exert an influence over all aspects of mental life. Examples include memory and attention. Conversely, vertical faculties are best characterised with reference to their subject matter, or their domain of operation. Aspects of the language or the visual processing systems could most accurately be described as vertical faculties as they are only operational with respect to their processing domains. The classification of verticality was earlier set out by Gall. The position not only considers that an aptitude such as one for music, say, is distinct from an aptitude for mathematics, but also that the psychological mechanisms underlying these capacities are also distinct. In other words, with reference to these examples Gall's position is that there is a certain delineable aspect of mental functioning that is concerned with mathematics and another separate aspect concerned with musical performance.

Fodor takes Gall to task over an extreme vertical faculty position. Gall further proposed that there is no such thing as acuity per se, but only acuity with respect to some aspect of cognition (e.g. visual acuity, auditory acuity etc). This stance is extended to apply to other aspects of cognition such as memory. However, the fact that an individual displays better memory for mathematical phenomena than linguistic phenomena does not necessarily warrant the conclusion that the individual possesses separate memory systems for the different aspects of cognition associated with mathematical or linguistic performance. It may simply be the case that the processes concerned with unravelling the mathematical input do a better job of producing an output of a form easily represented by the memory system than do the analogous processes associated with analysing the linguistic input.

Horizontal faculties such as memory or attention seem best used as descriptions of the overall manner by which the system operates. They characterise some nature of the resources available to the system as a whole and aren't instantiated with respect to any particular type of processing task.

Fodor recasts Gall's position and distinguishes five central properties associated with construing a mental faculty as vertical. These properties define a vertical faculty as domain specific, genetically determined, associated with distinct neural structures and computationally autonomous. Fodor explains this last property as following from the position that the vertical faculties do not compete for horizontal resources; in other words they do not compete for resources such as attention and memory. The knowledge pool accessible by a given vertical faculty is not the pool of knowledge about the world, but rather knowledge pertinent to the functioning of that faculty. A reformulation of the diagnostic criteria associated with vertical faculties forms the basis of Fodor's proposal that the modular perspective more accurately captures the nature of the cognitive architecture. Those parts of cognition which Fodor considers to be modular are the input systems: aspects of the cognitive system which form the link to the environment external to the organism, including the language system. He proposes that central processes act on the output of these modules and it is only at this point that world knowledge influences processing. The input systems are concerned with processing the exogenous information into some form interpretable by the central processing aspects of the system. Modules can be viewed as deterministic transformation functions. In programming terms, functions take an input and produce an output. A function responsible for addition, say, will always produce the same output given the same input. We know that if we input the numbers 5 and 6, the output will be 11. We know how addition works and we can formalise this. Similarly, a particular module receives input A and produces output B. It can be considered as operating in a precisely defined way, always producing the same output given the same input in a manner analogous to our addition function.

In summary, modules transform an input in a deterministic fashion; with the determinism defined solely with respect to the internal workings of the module. This notion of encapsulation, or of computation without recourse to information outwith the module is considered by Fodor to be one of the central defining features of what it constitutes for a processing component to be modular. There are five such defining characteristics to which we'll now briefly turn.

1. Input systems are domain specific. Each system operates over a precisely defined type of input. They are tuned to process only information of a certain type (such as linguistic information).

2. The operation of input systems is mandatory. When an individual hears an utterance it can't be perceived as anything other than an utterance, i.e. it can't be perceived as an uninterpretable stream of noise. This automaticity has been described by Marslen-Wilson and Tyler (1981) with respect to word recognition. Subjects couldn't help but identify words in the auditory stream even when explicitly told not to focus on that aspect of the stream.

3. There is only limited central access to the mental representations that input systems compute. There are some levels of representation within a module that are not available for conscious reflection.

4. Input systems are fast. This notion is intrinsically related to, and perhaps follows from, characteristic (2) which states that the modules operate automatically. Indeed, this type of behaviour was taken by Posner and Snyder (1975) as a central defining characteristic of what they term 'automatic' processing. This speed may be a consequence of the small, well defined set of information that needs to be considered by the module before an output is reached. Simply because there is less information pertinent to processing within the module (see following point), the solution is arrived at in a rapid manner.

5. Input systems are informationally encapsulated. The computations carried out within a module occur without recourse to information not represented within that module. Only the information within a

module influences processing within that module. Information not contained within the module can only influence processing at some later point.

The above summarises what it means for particular processes to be viewed as modular in nature. A further question can be asked concerning the precise computational nature of these processes. They may operate in either a serial or a parallel manner.

1.2.3 Serial and Parallel Processing —

Ambiguity in language is widespread, covering levels of word meaning, syntactic analysis and reference. From a processing perspective some basic questions are common at each of these levels. When an ambiguity is encountered, how is it treated? Is one possible interpretation selected rather than another or is every interpretation selected, either to the same or differing degrees? Consider the following example where an ambiguous word is encountered before its disambiguating context.

(1) The bug was found by the insect hunter/ security team.

(adapted from Swinney, 1979)

The word 'bug' is ambiguous between an 'insect' reading and a 'spying device' reading. When it is encountered the reader does not know which interpretation will be correct. There are two possible ways in which processing of the word may proceed at this point. One meaning of the word could be selected, perhaps determined by the individual's prior experience with the word. Serial processing involves a single solution being proposed at any one time. In the case of the above example, if the 'insect' reading is initially proposed but turns out to be incorrect, the alternative interpretation is adopted.

A different account which we can term a parallel processing account would posit that both meanings are accessed and held in mind until disambiguating information is able to select between the two. Multiple solutions can co-exist.

In the same way that words can be ambiguous, so too can sentences. Consider Examples (2) and (3) below (taken from Tyler and Marslen-Wilson, 1977).

(2) Flying planes are dangerous.

(3) Flying planes is dangerous.

In Example (2) the phrase 'flying planes' should be interpreted as a complex Noun Phrase. In Example (3) however it should be interpreted as a Verb Phrase. When syntactic ambiguities of this sort are encountered, how are they treated? How does the parser decide which interpretation to select? Is one solution initially adopted or are multiple solutions held in mind? What information is used to decide between alternatives? Some strong predictions can be generated on the basis of whether we consider the parser as operating in a modular or nonmodular fashion.

1.2.4 Modularity applied to language processing

Broadly speaking, theories of parsing fall into one of two categories: modular and nonmodular accounts. Modular accounts propose that an initial stage of the parser behaves in a restricted manner, initially sensitive to only some of the information potentially relevant for sentence processing. One of the most influential parsing accounts, the Garden Path theory (Frazier, 1979), proposes that this information is purely syntactic. Stages following this level of processing may take advantage of non-syntactic information to evaluate what has been output by the syntactically determined preceding stage (Frazier, 1987; Rayner, Carlson and Frazier, 1983).

Within modular accounts of parsing, at points of syntactic ambiguity in the input, the parser can make a commitment to one of the potentially available syntactic analyses or may propose analyses in parallel (Gibson, 19??). Decisions at this level can be made by the system using only syntactic knowledge (perhaps by parsing principles reflecting the nature of the underlying syntactic structure of the

sentence). Consider example (4) below (taken from Trueswell, Tanenhaus and Garnsey, 1994) :

(4) The evidence examined by the lawyer turned out to be unreliable.

The verb 'examined' is ambiguous between a past tense and a passive participial interpretation (i.e. 'that was examined'). The animacy of the preceding Noun Phrase, 'the evidence', provides information which rules out the past tense interpretation. Modular accounts propose that such information isn't available initially within the system to help resolve ambiguity (Ferreira and Clifton, 1986). It is only at some point following that information of this type is used.

Alternatively, nonmodular accounts place no restriction on when information can exert an influence and propose that all information that is available to the parser is available at the same (early) processing stage. Different sorts of information are considered to act as constraints which restrict possible syntactic analyses. The constraints do not qualitatively differ from each other although there may be quantitative differences in their relative contributing weights. In the case of Example (4), those arguing for a constraint based position claim that information about animacy is available to the parser to influence initial parsing. In other words the reader will not be garden pathed and will correctly initially interpret the verb 'examined' as a passive participial.

We can also interpret the behaviour of the parser with respect to the serial and parallel processing positions outlined above. Parallel modular accounts allow for the parser to construct multiple analyses which are then decided between using non-syntactic information at a subsequent stage. In the case of a serial parser, where only one analysis is pursued at any one time, at points of ambiguity a decision must be made as to which analysis to adopt. Modular serial accounts permit this decision to be made through reference to only the syntactic structure of the possible analyses. These accounts propose that when an ambiguity is encountered the parser employs the principles of Minimal Attachment and Late Closure in order to select an analysis to pursue (Frazier, 1979, 1987). These principles make

reference to the phrase-structure of the analysis being constructed. Minimal Attachment proposes that the parser will prefer to construct the interpretation that requires postulation of the minimal number of tree nodes while Late Closure proposes that the parser will prefer to incorporate a new phrase with the one currently being processed if grammatically permissible. These principles cover initial parsing. Only later does non-syntactic information exert an influence.

Constraint based accounts are parallel processing accounts as they propose that a number of analyses are pursued on the basis of the available information, although ultimately only one is selected. Both syntactic and non-syntactic information is used by the system at the same time.

In summary then, modular accounts of parsing propose that only syntactic information is used immediately. Semantic information exerts an influence later. Non-modular or constraint based accounts propose that all information that is relevant exerts an influence at the same time.

So far we have only focused on ambiguity at the level of parsing. There is also ambiguity at other levels of processing language. The following section briefly outlines the level of referential ambiguity.

1.3 Parallels between Accounts of Parsing and Accounts of Anaphor Resolution

Anaphoric pronouns refer to some previously mentioned character in a text. In order to correctly interpret a pronoun, the reader must correctly identify to which character it refers. In the same way that there are constraints as to which syntactic analysis is possible in the case of parsing ambiguities, there are also constraints on which reference assignment is possible in the case of referential ambiguity.

(5) John fascinated Mary because he was interesting.

The gender constraint in the pronoun limits the pronoun's antecedent to one matching this characteristic. Gender is a strong constraint and

cannot easily be violated. Ultimately it limits which analysis can be adopted. If the anaphor resolution system behaves in a modular fashion this information may not be used immediately however. An initial interpretation by the system may proceed with reference to solely structural information.

In example (6) the pronoun can potentially refer to either character although it is preferentially interpreted as coreferential with 'John'.

(6) John fascinated Bill because he was interesting.

This is because of a property associated with the verb called implicit causality (Garvey and Caramazza, 1974). The verb 'fascinate' possesses a bias which, in a sentence such as (6), biases interpretation of the pronoun as referring to the first Noun Phrase. Verb bias is a weaker constraint than gender as it can be violated (as in example (7)).

(7) John fascinated Bill because he was easily interested.

The pronoun is now interpreted as coreferential with 'Bill', although this goes against the implicit causality bias of the verb. When is verb semantic information such as implicit causality used by the system ?

We can ask ourselves the same basic question about anaphor resolution as has been asked about parsing. Does the system behave in a restricted or nonrestricted fashion? If it behaves in a restricted fashion, an initial stage of the anaphor resolution mechanism will operate only over a particular subset of the information available in the input. That information may be solely structural and an initial stage may employ some principle such as Parallel Function Strategy (PFS). Stated simply PFS proposes that a pronoun will be interpreted as coreferential with the character occupying the same grammatical role in the preceding clause. In other words the pronoun in example (8) below will be preferentially interpreted as referring to the character 'John'.

(8) John saw Bill and he waved from across the street.

PFS can be considered as analogous to the strategies of Minimal Attachment or Late Closure in parsing. It is a purely structurally informed heuristic. If the anaphor resolution system behaves in a modular or restricted fashion it may be the case that strategies such as PFS initially inform the system. Only later does information such as verb semantics exert a processing influence.

The information not considered initially will exert an influence at a later point, perhaps to guide reanalysis following misassignment. If we adopt a nonrestricted, constraint based position, we propose that all information that ultimately influences anaphor resolution does so during the same stage in processing and that no information is accorded a privileged status other than in the sense that a particular type of information may be assigned a greater constraining weight.

Stated simply, modular accounts of parsing propose that initial processing is informed by a well defined set of information. In the case of the Garden Path model this is syntax, i.e. low level information. At the start of this chapter we described how we wanted to draw parallels between the processes of parsing and reference resolution. A structurally driven strategy for resolving anaphors such as PFS is equivalent to the parsing strategies of Minimal Attachment and Late Closure. If the anaphor resolution system is modular, an initial stage of processing should proceed with reference only to a well defined set of information. We claim that if this low level information is purely structural, as is the case with modular parsing theories, a strategy such as PFS will initially determine how a referential ambiguity will be treated without reference to higher level factors such as verb semantics. Only at some subsequent stage of processing will semantic information exert a processing influence. If resolving anaphoric reference behaves in a nonmodular fashion, all information that is relevant will exert a processing influence at the same point (cf. Truswell, Tanenhaus and Garnsey, 1994).

The following chapter provides a summary of the literature examining anaphors in general. At the end of that chapter we shall

set out the precise predictions that follow from the reference resolution system behaving in a modular or nonmodular fashion.

CHAPTER 2

2.1 Introduction

This chapter provides a summary of the issues associated with anaphora. It begins by sketching the functional differences between anaphors of different forms. It then examines the role of discourse focus on anaphoric resolution and the time course of resolution. There then follows a summary of those accounts of reference resolution which we can term restricted accounts (see Chapter 1 for definition) focusing as they do on the role of structural information in the process of resolution. We then examine how other types of information are used by the anaphor resolution system before proposing a two-stage model of anaphor resolution informed at different points in time by structural and non-structural factors.

In contrast to research on parsing, research examining the processing of anaphors forms a less homogenous body of work. Fodor's Modularity thesis has played an important part in guiding empirical work on parsing behaviour but has had little influence on the literature examining the level of anaphor resolution. It is certainly possible to examine whether that part of the cognitive system responsible for resolving anaphoric reference behaves in a modular or nonmodular fashion. That it hasn't yet been done is perhaps indicative of the large number of other issues researchers examining the behaviour of anaphora have deemed more worthy of examination. The following section provides a summary of some of these theoretical issues and the empirical work generated as a result. A number of dimensions can be extracted from the literature along which we can categorise previous research. These dimensions include the time course of the anaphor resolution process, the behaviour of different forms of anaphora and the influence of both structural and non-structural factors on anaphor processing. Initially we will focus on the functional role of anaphora within a text before spending some time summarising the work pertinent to each of these dimensions.

2.2 Anaphora

In order to construct and maintain a mental model standing in relation to the text being read, readers must integrate each unit of text with those preceding it. In other words, they must be aware of the cohesive relations between units. Cohesion may be achieved in a number of ways. For our purposes we shall concentrate on referential cohesion. This is accomplished through the application of referring expressions.

(1) John saw Mary in the park. He waved at her.

In example (1) above, the two sentences can only be properly comprehended if the reader successfully interprets the pronouns in the second sentence. The referential links must be established between anaphor and antecedent before integration of the semantic information in the second sentence describing the relationship between the two characters can be achieved.

In the case of example (2) below the pronoun in the second sentence cannot be successfully resolved initially as the disambiguating information doesn't occur until after the pronoun.

(2) John saw Bill in the park. He waved at John.

If we replace the pronouns in example (1) with repetitions of the characters names (see example (3) below), the text sounds awkward although the referential links are maintained.

(3) John saw Mary in the park. John waved at Mary.

What determines the type of anaphoric form that may be used ?

2.2.1 Differences between Anaphors

In this section we describe why people select one anaphoric form over another and the processing behaviour associated with particular

types of anaphors. Selection of form is intrinsically linked to the degree of focus of the intended referent.

2.2.1 What is focus ?

We can define focus operationally as the entity which readers prefer a text to continue with reference to. In example (4) below, there is a preference for the sentence to continue with reference to the character 'John'.

(4) John fascinated Bill because ...

A continuation may be of the form '... he was very interesting.' The pronoun and NP 'John' are coreferential so for this sentence we can say that the character referred to by 'John' is in focus.

The notion of focus is central to several psychological theories of language processing (e.g. Sanford and Garrod, 1981; Gordon, Grosz and Gilliom, 1993). Roughly, an entity can be considered to be in focus when it occupies a privileged role in the reader's centre of attention.

2.2.2 What determines focus ?

A number of devices can be used to place an entity in focus including recency of mention (Clark and Sengul, 1979) and prior topicalisation (Anderson, Garrod and Sanford, 1983). Sanford, Moar and Garrod (1988) demonstrate that the way in which a character is introduced in a text influences how strongly focused it will be. Specifically, characters introduced by a proper name are more highly focused than characters introduced through the use of a definite description.

There is a difference between local and global levels of focus. The global level can be interpreted as corresponding to entities relevant to the topic of the text as a whole while the local level can be seen as being composed of temporary shifts in focus between main and transient characters (Anderson et al, 1983). Several psychological theories restrict their examination of focus to focus at the local discourse level. Centering (see below) restricts its account to adjacent

utterances. There is also a level of focus at the more global level, roughly what can be considered the discourse topic. The topic can be considered to be what the discourse as a whole is about. Local discourse focus can shift throughout a section of text but the global focus will remain more or less constant.

Sanford and Garrod (1981) propose their Memory Focus model in an attempt to account for the linkage between focus that occurs over short discourse segments and a more general global level of focus. They emphasise the role that background knowledge plays in structuring within this model. Very roughly, this Memory Focus model is proposed to consist of two dimensions connected by mappings. The first dimension, Explicit Focus, consists of tokens standing for characters relevant to a particular stretch of discourse. The Implicit Focus dimension is considered to be that part of the discourse model which maps onto pre-existing knowledge structures associated with the situation described by the text (Minsky, 1975). Garrod, Freudenthal and Boyle (1994) suggest that it is the level of Implicit Focus that pronominal anaphora access. Pronouns provide a direct route into the conceptual level of the reader's discourse model.

More explicit forms of anaphor access their antecedents in a less direct way. Within the Memory Focus account they correspond to tokens in the Explicit Focus dimension of the model. For characters central to the discourse, there is a rich set of mappings between the Explicit and Implicit Focus aspects of the discourse model. Anderson, Garrod and Sanford (1983) demonstrate that even with shifts of scenario (i.e. where Implicit Focus changes) these elements remain central in the reader's discourse model.

Pronouns map directly onto implicit focus and therefore access the conceptual level of the discourse model directly. The tokens in the Explicit Focus part of the model are associated with a more superficial level of representation. The mapping between these tokens and their corresponding roles in Implicit Focus must be understood for Noun Phrase anaphors to access the same level as that accessed by pronouns. This is a less direct route to accessing the conceptual level than the use of pronominal anaphora.

2.2.3 How does focus influence selection and interpretation of anaphoric form ?

Ariel (1990) argues that the choice of anaphor is determined by the degree of focus possessed by its intended referent. A highly focused referent will be preferentially referred to using a pronoun, while a full name will be used to refer to an antecedent not in focus. There is a negative correlation between informational content of an anaphor and degree of focus of its referent.

Centering theory also proposes that reference to a focused entity will be preferentially realised through the use of a pronoun. Pronouns referring to antecedents not in focus take a relatively long time to read. Ehrlich and Rayner (1983) demonstrated longer fixation times, and by extension processing difficulty, on the region following a pronoun when the pronoun's antecedent was at a distant point in the text. This is supported by Ehrlich (1980), Clark and Sengul (1979) and Frederikson (1981).

Gernsbacher (1989) proposes that recency of mention is also an important determiner of level of explicitness of anaphor. She states 'the longer the distance between an anaphor and its antecedent, the more explicit the anaphor' (p. 138) where distance is taken to mean the physical distance between an anaphor and its antecedent.

The degree to which an antecedent is topical or part of global focus also seems to partly determine the level of explicitness of a following anaphoric expression. Antecedents which have fallen out of focus needed to be referred to using a relatively informationally rich anaphor (Ariel, 1990). Anderson, Garrod and Sanford (1983) examined the role of prior topicalisation in determining antecedent accessibility and found a preference for pronouns to refer to topicalised entities over non-topicalised ones. Similar effects were found by Chafe (1974), Givon (1983) and Marslen-Wilson, Levy and Tyler (1982) who report that the more topical the antecedent, the less explicit the anaphor used to refer to it.

2.2.4 The Influence of Focus on Anaphoric Processing

Regardless of how we precisely characterise focus, antecedents which no longer occupy a position within the focus of a reader's discourse model must be re-introduced using a referentially specific device. When an anaphor is employed that is more referentially specific than appropriate (i.e. if its antecedent is in focus), the result is awkward and unnatural sounding text. However, Vonk, Hustinx and Simons (1992) report that under conditions where a text contains a shift of theme, overspecification acts as an important signal to the reader that such a shift is occurring. Consider the following set of sentences :

1. Sally Jones got up early this morning.
2. She wanted to clean the house.
3. Her parents were coming to visit her.
4. She was looking forward to seeing them.
5. She weighs 80 kilograms.
6. She had to lose weight on her doctor's advice.
7. So she planned to cook a nice but sober meal.

Although the pronoun 'she' in sentence (5) unambiguously refers to the character 'Sally Jones', Vonk et al describe a preference to use an anaphor of increased specificity, i.e. the repeat name 'Sally'. This is because sentence (5) shifts theme from the visit of Sally Jones' parents to her weight problem. Without an overt thematic shift signal, it is difficult to integrate the content of sentence (5) with what has been read previously until the connection becomes clear in sentence (7). Vonk et al propose that readers interpret anaphors of a more specific form than is necessary as signals of shifts of theme. We can replace the pronominal anaphors in sentences (1) through (4) with the referentially unambiguous repeat name, 'Sally Jones'.

1. Sally Jones got up early this morning.
2. Sally Jones wanted to clean the house.
3. Sally Jones' parents were coming to visit her.
4. Sally Jones was looking forward to seeing them.

Adopting the Vonk et al line of argument, we argue that the difficulty in reading such text is due to the reader interpreting the overspecific anaphor in each sentence as a signal for a thematic shift. As no shift occurs, the conditions for using such over-specification are violated and processing disruption is encountered. The Vonk et al position is supported by both production and comprehension data. They report that thematic shifts are produced by subjects when they have to use an overspecific anaphor and that when required to produce a thematic shift, subjects produce overspecific anaphors. They also report two comprehension probe experiments which indicate that the presence of an overspecified anaphor reduces the accessibility of information contained in the preceding sentence. If the overspecified anaphor is treated by the reader as a cue that the next stretch of discourse will introduce a new theme rather than continue the existing one, a new foregrounded segment of the discourse model will be introduced and the segment containing the content of the previous theme will fall out of focus, at least temporarily.

The thematic shift signal documented by Vonk et al is a very specific example associated with overspecified anaphors. As such overspecification is relative, it is not possible to determine how particular anaphoric forms per se differ in their processing consequences. Indeed, the behaviour of a particular anaphor relative to other possible anaphors may be of a very different form from the absolute behaviour of that anaphor in general.

2.2.5 How anaphors access levels within the Discourse Model

Cloitre and Bever (1988) and Garrod, Freudenthal and Boyle (1994) propose a much more general framework in which they consider the operational distinctions between different anaphoric forms. It is possible to consider anaphora as pointers within the reader's discourse model. There seems to be evidence to suggest that different anaphoric forms tap into different levels of representation within this discourse model.

Using materials such as (1) below, Cloitre and Bever presented subjects with a pair of sentences followed by a probe word which had

previously modified the anaphor's antecedent. Subjects were required to engage in tasks of recognition, category decision or lexical decision following presentation of the probe word.

- (1) The gangly busboy spilled soup on the famous actress.
- (a) A waiter ran to help the busboy. (repeat noun anaphor)
- (b) A waiter ran to help him. (pronominal anaphor)
- (c) A waiter smothered a giggle. (baseline control)

Probe : gangly

In both recognition and category decision tasks, greater facilitation was achieved following the pronoun. In the case of lexical decision, performance was facilitated following the presence of a repeat noun anaphor. This suggests that following a pronoun a more conceptual level of the discourse representation is accessed. A repeat noun anaphor accesses a more superficial level of representation and so a task requiring a level of processing associated with superficial features (such as lexical decision) will be facilitated.

We shall describe the Garrod et al position in detail below. For the moment we shall simply describe one aspect of the experiments they report. Basically, they found a difference between reading pronouns and repeat name anaphors. Following reading of a pronoun, interpretation of that pronoun was quickly influenced by factors such as discourse focus and semantics. This was not found when repeat name anaphors are initially read. Garrod et al argue this is because the different forms of anaphor access different levels of the discourse representation. Although grounded in different theoretical frameworks, both Cloitre and Bever and Garrod et al suggest that pronominal anaphora access a conceptual level of the reader's discourse model. Garrod et al account for this at the level of Implicit Focus within the Sanford and Garrod (1981) Memory Focus model.

Conceptually the Memory Focus model is identical to the Cloitre and Bever account. The Memory Focus model proposes that repeat noun anaphors initially map onto a relatively superficial level of representation. For pronominal anaphora, a level of representation at

the conceptual level of the antecedent is accessed. Cloitre and Bever suggest that ultimately repeat noun anaphors also tap into this level but this level of access takes some time to be realised when repeat name anaphors are used.

2.2.6 The mechanism of thematic shifting explained

The position advocated by Vonk et al above is consistent with the general account of different levels of discourse representation access associated with different anaphoric forms. It is possible to explain why thematic shifts are induced by overspecific anaphors in light of the position outlined by Cloitre and Bever and by Garrod et al.

If a repeat name anaphor is overspecific relative to the form appropriate at a particular point, it effectively shifts the level of grain within the reader's discourse model from the conceptual to a more superficial level. Recall both Cloitre and Bever and Garrod et al proposed that repeat name anaphors access a more superficial level of representation. If the anaphor's antecedent is in focus, a pronoun will be the preferred form of referring device. In the case of the actual anaphor being a repeat name, it will shift the level of grain within the discourse model from the conceptual to the more superficial. We can define a thematic shift as a shift in grain, for example a shift in topic. In order to understand the thematic shift, the reader must start a new discourse segment to contain this topic shift.

As a shift in granularity of representation is required in order to integrate the information associated with the new theme, the overspecification facilitates integration of subsequent information contained within the sentence in which the overspecified anaphor occurs. In other words, a shift of grain is required to interpret a new theme. As this shift is induced by the use of an overspecified anaphor, integration of subsequent information is facilitated.

2.3 The Time Course of Anaphoric Processing

Up to this point we have concentrated on how discourse focus and other factors influence interpretation of different anaphoric forms. In

this section we focus on the temporal dimension associated with the processes responsible for resolving anaphoric reference. When is it that different types of information exert a processing influence ?

Just and Carpenter (1980) propose an account relating eye movements to the processing of linguistic stimuli based on the assumption of immediacy of processing. Their account makes a number of strong predictions regarding the temporal nature of language processing.

2.3.1 The Immediacy Hypothesis —

The immediacy assumption proposes that the reader tries to interpret each word as it is encountered, even at the expense of making guesses that may turn out to be wrong. This has direct consequences for how we consider the processing of anaphors to proceed as not all the information pertinent to resolving an anaphor is necessarily available at the point at which the anaphor is read. In the example below the antecedent of the pronoun cannot be identified until the disambiguating information following the pronoun has been read. With respect to the immediacy hypothesis, do readers delay interpretation of the pronoun or do they make an initial guess using, say, structural information ?

(5) John ran after Bill as he owed John some money.

Using an eye-tracking methodology, Ehrlich and Rayner (1983) demonstrated that readers did not necessarily resolve pronouns immediately. This occurs especially under conditions where the pronominal antecedent is present at some distant point in the text. It also occurs even under conditions where a gender cue unambiguously identifies a pronoun's referent. They examined three types of contexts where the antecedent of the pronoun in the final sentence appeared at near (6a), intermediate(6b) and far(6c) points in the text.

(6a) A group of people who shared an interest in photography had recently started writing a newsletter of their activities. In fact, in one

room Mark was mailing a copy of the paper to Susan. She was very involved in photography and spent every weekend taking pictures.

(6b) A group of people who shared an interest in photography had recently started writing a newsletter of their activities. In fact, in one room Mark was mailing a copy of the paper to Susan. He was very involved in photography and spent every weekend taking pictures.

(6c) A group of people who shared an interest in photography had recently started writing a newsletter of their activities. Mark wrote most of the copy but the other members did a lot of work as well. In fact, in one room Cathy was mailing a copy of the paper to Susan. He was very involved in photography and spent every weekend taking pictures.

Ehrlich and Rayner found an increase in reading time for the region following the pronoun in condition 6c where the pronominal antecedent was distant. Regardless of focus information, gender alone should be able to uniquely identify the appropriate referent. Although this finding appears at odds with Just and Carpenter's position, it may be that gender information is not available to that part of the system responsible for resolving anaphors. This would be the case if some initial processing of the anaphor behaved in a modular fashion and this module was not sensitive to gender information. Perhaps it should be more correct to interpret the immediacy hypothesis as meaning that the reader attempts to interpret each word as it is encountered despite the restrictions on the information available for processing at that point. So, it may be the case that an incorrect interpretation will be given to a particular word not just because information necessary for arriving at an unambiguous interpretation hasn't yet been read, but because although this information has been read, it may not be available to exert a processing influence within a particular module.

2.3.2 Differences in Processing of Pronouns and Noun Phrase Anaphora

Sanford and Garrod (1989) propose the importance of distinguishing between initiation and completion of processing. The immediacy issue of Just and Carpenter becomes a question of what constitutes immediacy. Is processing simply initiated when each word is encountered or does completion of processing (i.e. full understanding) also occur? Given the example below, it is possible that although the reader starts interpreting the pronoun when it is first read, the actual process by which an antecedent is identified only occurs some time later.

(7) John blamed Sue because he was in a bad mood.

The point at which termination of processing can occur is determined by how quickly different types of information necessary for successful interpretation of an anaphor exert a processing influence. For anaphors, information above and beyond what is contained in the anaphor itself plays a role in interpretation. In Chapter 1 we mentioned one factor, implicit causality, which might be one constraint used to inform the system. There are factors at other levels however. One of these levels is discourse structure. The way in which a character is introduced to a discourse and the manner in which they are referred to within a particular discourse affects to what degree they are treated as a topic character in the text. Recall section 2.2.3 where we asked how focus influences selection and interpretation of anaphoric form. We can now ask the extended question of how quickly does focus information influence comprehension?

Garrod, Freudenthal and Boyle (1994) examined the time course of the influence of 3 factors on interpreting an anaphor. In two eye-tracking experiments they looked at the influence of discourse focus, form of anaphor used and pragmatic information. They examined passages such as (A) and (B) below.

(A) A dangerous incident at the pool -

Alexander was an inexperienced swimmer and wouldn't have gone in if the male lifeguard hadn't been standing by the pool.

But as soon as he got out of his depth he started to panic and wave his hands about in a frenzy.

(C1) Within seconds he sank into the pool.

(C2) Within seconds he jumped into the pool.

(B) A dangerous incident at the pool -

Elizabeth was an inexperienced swimmer and wouldn't have gone in if the male lifeguard hadn't been standing by the pool. But as soon as she got out of her depth she started to panic and wave her hands about in a frenzy.

(C3) Within seconds she sank into the pool.

(C4) Within seconds she jumped into the pool.

(C5) Within seconds he jumped into the pool.

(C6) Within seconds he sank into the pool.

Each passage focuses on the first mentioned character, either Alexander or Elizabeth. This is achieved by introducing these characters by way of a proper name and through subsequent reference in the following sentence. In other words, the discourse focus biases towards the first rather than the second mentioned character. Conditions C3 and C4 contain a pronoun referring unambiguously to the focused character, C5 and C6 contain a pronoun referring unambiguously to the unfocused character. C4 and C6 are pragmatically incongruent as, although the verb biases toward one character, the pronoun selects the other.

In C1 and C2 as the pronoun can potentially refer to either character, it is only discourse focus that can influence interpretation as soon as the pronoun is encountered and before the subsequent disambiguating information has been read.

Garrod et al used an eye-tracking methodology. What is important is whether the different types of information exert a processing influence as soon as the reader first encounters the pronoun. On

examining the first pass reading time for the region containing the pronoun, there was no evidence for an influence of discourse focus. In other words, the pronoun in conditions C5 and C6 was not fixated for a longer time than the pronoun in the other four conditions. For the following region however there was evidence in the first pass reading time data for an influence of discourse focus. Consistent with the data reported by Ehrlich and Rayner (1983), when a pronoun refers to an unfocused antecedent, there is an increase in reading time for the region following the pronoun.

For verb plausibility, Garrod et al report that it plays an early role in the resolution process where the pronoun refers to the focused entity, but not otherwise. This makes sense if we consider that the verb information is necessary for integration of the information following the pronoun with what has been read previously. Integration cannot proceed unless the pronoun's antecedent has been identified. In cases where this identification is problematic, i.e. where the pronoun refers to an unfocused entity, integration cannot occur immediately following the pronoun and so no effects of verb plausibility were observed. The first experiment of Garrod et al indicates that factors such as discourse focus and pragmatics play an early role in the anaphor resolution system. Or at least when the reader is required to resolve an anaphor that is a pronoun.

In their second experiment Garrod et al replaced the pronominal anaphors with repeat name and definite description anaphors. Otherwise the passages were identical to those in their first study. In line with their first experiment they found no evidence of discourse focus playing a role when an anaphor itself is encountered. In contrast to the first study there was no evidence in the following region that discourse focus was playing a role. Neither was there any evidence of verb plausibility exerting an influence in this region. The most obvious conclusion to draw in light of this set of data is that in the case of pronouns referring to a focused antecedent, discourse focus and pragmatic information exert an early influence on processing. This is not the case with fuller referring expressions or with pronouns referring to unfocused antecedents. In summary then, there is evidence suggesting that given the appropriate conditions are

met, higher level factors such as focus and pragmatic information informs the anaphor resolution system at an early stage.

Further evidence for an early influence of discourse factors comes from Marslen-Wilson, Tyler and Koster (1993) who examined the degree to which pragmatic knowledge comes to bear an influence on resolving anaphoric reference. Garrod et al employed an eye-tracking measure while Marslen-Wilson et al used a cross modal technique. Subjects heard a context (see below) followed by a sentence fragment and then had to name a visually presented word, *him* or *her*. The structure of the discourse focused on one character. The verb in the final sentence either also biased towards this character, was neutral with respect to this character, or actually biased towards the alternative character. (see examples 8-10 below)

Condition 1 : Discourse bias with congruent verb bias

(8) After the surgeon had examined the 12-year old girl with the badly broken leg, he decided he would have to take immediate action. He'd had a lot of experience with serious injuries. He knew what he had to do next.

- A. He quickly injected ...
- B. She quickly injected ... Him/Her
- C. Quickly injecting ...

Condition 2 : Discourse bias with neutral verb

(9) As Bill was buying popcorn at the movies, he saw an old girlfriend get in line for a ticket. He had arrived at the movies especially early. He wanted to be sure of getting a good seat.

- A. He waved at ...
- B. She waved at ... Him/Her
- C. Waving at ...

Condition 3 : Discourse bias with opposing verb bias

(10) Mary lost hope of winning the race to the ocean when she heard Andrew's footsteps approaching her from behind. The deep sand was slowing her down. She had trouble keeping her balance.

- A. She overtook ...
- B. He overtook ... Him/Her
- C. Overtaking ...

The pattern of data associated with subjects' naming times indicated that the influence of both pragmatic and discourse focus factors was of sufficient magnitude to result in early processing consequences. That subjects in Condition 3C demonstrated a naming facilitation associated with the probe *her* indicates the rapidity with which the verb semantics have been integrated with the discourse model. For the reader to know which character was doing the overtaking and which character was being overtaken they must have interpreted the verb against their discourse model. In line with the Garrod et al data examining interpretation of pronouns, the results of Marslen-Wilson et al demonstrate the speed with which such high level pragmatic factors can exert an influence on interpretation.

The system appears to be behaving in a highly incremental manner and allows for, in this case, pragmatic factors to exert an early, strong influence on the associated processing mechanisms. That discourse focus information is also behaving in a similar manner is evidenced by a facilitation to naming the probe *her* in Condition 2C. The verb is neutral with respect to which character it biases toward. The structure of the context however biases towards interpretation of the character 'Bill' as the thematic subject associated with the discourse. In order for subjects to respond more quickly to the word *her* they must be sensitive to the discourse focus information and interpret the focused character 'Bill' as the grammatical subject of the verb 'waving'.

The data of Marslen-Wilson et al indicate that the information necessary for the process of anaphoric resolution to operate is certainly available to the reader when a pronoun is encountered; at least with respect to the constructions they examined and using the measures they employed. This can be interpreted as evidence supporting the potential for the reader to interpret the pronoun when it is encountered. As it is cross-modal, the task used by Marslen-

Wilson et al however is certainly not one necessarily tapping into processes associated with normal reading and therefore it doesn't necessarily follow that readers resolve pronouns when they encounter them under normal circumstances. The eye-tracking methodology used by Garrod et al is much more likely to be tapping into normal reading. In light of this we can conclude that given the appropriate conditions, high level pragmatic does inform the anaphor resolution system at an early point, although perhaps only in the case of pronouns.

Garrod and Sanford (1985) examined the processing of fuller anaphoric phrases. They examined materials similar to those employed by Garrod et al (repeated below). Subjects were engaged in a spelling error detection task where they were required to respond as soon as a spelling error was detected.

(A) A dangerous incident at the pool -

Elizabeth was an inexperienced swimmer and wouldn't have gone in if the male lifeguard hadn't been standing by the pool. But as soon as she got out of her depth she started to panic and wave her hands about in a frenzy.

C1 Within seconds Elizabeth jumped/ jimped into the pool.

C2 Within seconds the lifeguard jumped/ jimped into the pool.

C3 Within seconds Elizabeth sank/ senk beneath the surface.

C4 Within second the lifeguard sank/ senk beneath the surface.

Following anaphoric reference to either character, response latencies was shorter for predictable than unpredictable verbs. In other words, consistent with the findings of Garrod et al, discourse focus appears not to be exerting an influence as soon as these anaphoric forms are encountered. Garrod and Sanford argue that the anaphor must have been interpreted with respect to the discourse model in order for the difference in response latency between predictable and unpredictable verbs to have arisen but this interpretation was not influenced by discourse focus. When a pronoun rather than a fuller anaphoric form was used an effect of discourse focus was found. Only for focused

entities was there a difference in response latency between predictable and unpredictable verbs. Again this is in line with the Garrod et al position and further supports the claim that interpretation of pronouns is sensitive to discourse focus factors. For fuller anaphoric forms such factors do not influence processing at this point.

As evidenced by the response latency difference between predictable and unpredictable verbs following definite anaphoric expressions, Garrod and Sanford conclude there is ample evidence to suggest that with full anaphoric noun phrases, the anaphor's antecedent is identified rapidly.

In the case of pronominal anaphora the evidence is more equivocal. Identification of a pronoun's antecedent may be delayed. Recall the evidence from Ehrlich and Rayner (1983) and Garrod et al (1994) demonstrating effects of their experimental manipulations arising on the region following the one containing the pronoun.

So, in light of the above it appears that pronouns and repeat name anaphors are treated differently by the language system. High level factors such as pragmatics can influence the interpretation of pronouns at an early point in processing but only provided the appropriate conditions are met. The antecedent of an anaphoric noun phrase may be identified immediately but this is not necessarily true in the case of pronominal anaphors.

2.4 The Probe Task

So far all the experiments reported in this chapter have used either some form of reading time measure or a naming task. One paradigm which is assumed to provide more of a direct insight into the nature of anaphoric processing is the probe task. The task has been used extensively within areas of the literature including syntactic processing and anaphoric processing. When employed to address questions of anaphoric processing, it can be used to monitor changes in the level of activation of potential antecedents. The basic rationale behind the task is that subjects will respond quickly to a probe word

if the activation of the word monitored by the probe is higher compared to some resting baseline. The word monitored may be the same word as the probe or an associate in the case of lexical priming. With respect to anaphoric processing it is assumed that the consequence of identifying an anaphoric antecedent is an increase in the activation of that antecedent. For pronominal anaphora this has been found in a probe naming task (Leiman, 1982), in a lexical decision task (Cloitre and Bever, 1988) and in tasks where subjects have to decide whether the probe word appeared earlier in the sentence (Chang, 1980; Cloitre and Bever, 1988; Gernsbacher, 1989; McDonald and MacWhinney, 1995; Garnham, Traxler, Oakhill and Gernsbacher, 1996). Comparable effects have also been found for noun phrase anaphora (Dell, McKoon and Ratcliff, 1983; Gernsbacher, 1989).

2.4.1 Problems with the Probe Task

Although it may be able to focus on aspects of processing possibly unmeasurable using a reading time technique, there are some very serious restrictions on what can be inferred from probe task data. The first is that when the probe task is employed to examine the effect of anaphoric processing, it is effectively measuring the consequences of processing rather than the nature of processing itself.

Another caveat to raise at this point concerns one of the most basic assumptions of the probe task : that of lexicalisation. Although it is generally assumed that activated concepts map onto their lexical counterparts, this is not necessarily the case. The discourse representation which the probe task aims to measure may consist solely of tokens mapping onto information about characters in a text. In example (11) below, there may be a token created in the reader's mental model corresponding to the character 'John' mapping onto the information 'went to the park'.

(11) John went to the park.

It may be possible to access the word 'John' from the token representing him but the lexical item 'John' is not necessarily part of

this representation. As the activation of lexical items is not intrinsically part of the discourse model, the assumption that measuring activating of those items reveals something of the underlying discourse model structure may be incorrect.

Finally, another fundamental assumption of the probe task is that a probe corresponding to a highly active lexical item will be responded to more quickly than one of lower activation. This is contradictory to the position adopted by the Centering theorists (Gordon, Grosz and Gilliom, 1993). We shall explore Centering Theory in more detail below but for the moment simply say that Centering predicts a repeat name penalty associated with a repetition of the name corresponding to the character most in focus. The second occurrence of the name 'John' in example (12) below will be read more slowly than the second occurrence of the name 'Bill' in example (13).

(12) John waved at Bill when John spotted him.

(13) John waved at Bill when Bill spotted him.

The reason that the repetition of the word 'John' will be read more slowly in example (12) than the word 'Bill' in example (13) is because the character 'John' is more focused than 'Bill' and should be realised using a pronoun. The character 'Bill' is not as focused as 'John' so there is no penalty associated with repeating that character's name. To put it simply, Centering Theory predicts a reading time penalty associated with reading a repetition of the name corresponding to a highly activated (focused) antecedent.

What this amounts to is the opposite prediction from the probe task. The repeat name penalty predicts that a repeat name will be read **slowly** if its antecedent is highly activated, while the probe task predicts that such a probe will be responded to more **quickly** if its antecedent is highly activated. Both positions are supported by their respective literatures so the only way in which they can be reconciled is through considering that the tasks of reading and participating in response judgements are non-overlapping in several important aspects. Of course the act of responding to a probe word is very

different from one of reading but the explanations for both these phenomena are grounded in terms of the activation of entities within the reader's discourse representation. Accepting that both positions are correct requires accepting that the representational arenas for these types of processing may be separate.

Aware of these caveats, we shall now examine what the probe task data seem to suggest as to the nature of anaphoric processing. The interesting issue is one of at what point in time the activation differential between control word and probe word occurs. As described above, the presence of a differential is assumed to indicate anaphoric resolution has taken place.

2.4.2 Probe Task Evidence

Chang (1980) presented subjects with sentences such as (14) and (15) below :

(14) John and Mary went to the grocery store and John/ he bought a quart of milk.

(15) John and Mary went to the grocery store and Mary/ she bought a quart of milk.

He found that after reading such sentences, subjects responded more quickly to the probe word 'John' following sentence (14) than following sentence (15). The version of sentence (14) containing the repeat name anaphor led to faster response times to the probe than the comparable pronominal anaphor version, but both led to faster probe response times than either version of sentence (15). The facilitation to responding to the probe 'John' following the version containing the repeat name 'John' anaphor over responding following the pronominal anaphor version may be due to form priming (Forster, 1979) and should not necessarily be taken as evidence that repeat name anaphors somehow increase their antecedents level of activation relative to pronominal anaphors. The facilitation in responding to the probe 'John' in the pronoun version of sentence (14) however does suggest that pronouns increase the activation of their

antecedents, at least by the end of the sentence. That no probe response facilitation was found following the pronominal version of sentence (15) can be considered as evidence that non-antecedents are not activated at the end of a sentence. However, some researchers are of the view that initially all potential antecedents do increase in activation. Corbett and Chang (1983) report that the activation of both antecedent and non-antecedent increases following a sentence containing a pronoun such as example (16) but not when that pronoun is replaced by a repeat name anaphor.

(16) Karen tried to beat Polly in chess but Polly / she always managed to win.

Dell, McKoon and Ratcliff (1983) employed a similar probe paradigm, but instead of the probe word simply appearing at the end of the sentence, the probe position was varied. Each sentence of the passages examined by Dell et al was presented at a rate of one word every 250 msec with each word appearing to the right of the preceding word which remained on the screen. The probe word, 'burglar', was presented in capitals underlined by a row of asterisks in the position where the next word in the sentence was due to appear. Upon presentation of the probe word, the rest of the sentence was erased. Subjects had to decide whether the probe word had been present in the preceding passage.

A burglar surveyed the garage set back from the street.

Several milk bottles were piled at the curb.

The banker and her husband were on vacation.

The₁ criminal₂ slipped₃ away₄ from the₅ streetlamp₆. (anaphor condition)

A₁ cat₂ slipped₃ away₄ from the₅ streetlamp₆. (control condition)

The version of the final sentence containing the phrase 'A cat' was used by Dell et al as their control against which to measure the activation level of the NP antecedent 'burglar' following the version containing the anaphoric noun phrase 'criminal'. Dell et al found evidence for activation of the antecedent at the last three probe points. It is possible however that the faster response time to the

probe 'burglar' may result from the semantic association between that phrase and the word 'criminal'. In their second experiment Dell et al introduced the probe word 'garage', an associate of 'burglar' in the first sentence in the passage. An activation differential between the probe word following the anaphor condition and following the control condition was found at probe point 3. This was found for both the probe word 'burglar' and associate 'garage'. At probe point 5, there was only an effect for the 'burglar' probe condition. When the related probe word 'garage' was replaced by 'bottles', a word which had simply appeared in the passage read by subjects but had no association with the antecedent 'burglar' within the context of the paragraph, no reactivation effect for 'bottles' was observed. On the basis of the data discussed by Dell et al it is possible to suggest that anaphoric noun phrases activate both their antecedents and concepts associated with those antecedents. The final control condition rules out the possibility that the anaphor is reactivating all words present in the passage read by subjects.

Shillcock (1982) used a cross-modal technique with a lexical associate probe and found an activation differential immediately following a pronominal anaphor. However, as can be seen from example (17) below, there was no competition between potential antecedents.

(17) The teacher did not board the train for the simple reason that it/ he was not going to the South Coast of England.

Probe : school/ street.

The antecedent of the pronoun 'he' is 'the teacher' while the antecedent of the pronoun 'it' is the phrase 'the train'. Each pronoun uniquely identifies its antecedent. Shillcock found evidence for suppression of the pronominal non-antecedent as opposed to activation of the antecedent. Using a similar task Marslen-Wilson and Tyler (1980) were unable to find a facilitation for a lexical associate of an antecedent following a pronoun compared to when the pronoun took a different antecedent. In their study Marslen-Wilson and Tyler incorporated a different baseline from the one used by Shillcock. They looked at the difference between reaction times to the probe

following different pronouns. MacDonald and MacWhinney (1990) stress the importance of choosing an appropriate baseline against which to measure probe response times. Using the Marslen-Wilson and Tyler baseline comparison, Shillcock reports that no activation differential was found suggesting that the comparison between related and unrelated probe which he utilised is more appropriate.

Gernsbacher (1989) proposes that the activation differential found in studies employing the probe task is a consequence of the combined result of two aspects of processing. She describes the processes associated with anaphoric resolution as triggering the mechanisms of enhancement and suppression. Enhancement increases the accessibility of an anaphor's antecedent while suppression decreases the availability of non-antecedents. Employing the probe task, she examined both pronominal and repeat name anaphors and concluded that the antecedents of pronominal anaphors increase in activation, although not immediately following the anaphor. The earliest probe point at which the activation differential between antecedent and non-antecedent was found was at the end of the sentence. This result replicates that reported by Corbett and Chang (1983). For repeat name anaphora, an activation differential was observed both at the end of the sentence but also at the probe point immediately following the anaphor. Whether this is the result of activation per se or simply the result of lexical repetition is unclear. Gernsbacher rejects the suggestion that it is simply an instance of lexical priming as she claims that the presentation of the probe word in uppercase instead of the lowercase presentation of the name embedded within the sentence obviates this criticism. She further suggests that even if this were the case, it is an explanation incapable of accounting for the effect of suppression found for the non-antecedent.

The strength of this defence is debatable. By abstracting to one level of representation, the position arguing for the importance of difference in case seems unsatisfactory. The words are identical at every level other than case so a simple lexical repetition argument can still be made to account for this result. This is the position put forward by Forster (1979) and is termed form priming. The argument focusing on the observation of suppression of availability of

the non-antecedent is perhaps more convincing. These two criticisms are only appropriate for the repeat name anaphor conditions. For pronominal anaphors, the finding of an activation differential between antecedent and non-antecedent at no point earlier than the end of the sentence is strong evidence for a delay associated with resolving pronominal anaphoric reference. If we maintain the view that the activation differential is a direct consequence of anaphoric processing, it seems we must accept that such processing is delayed, at least in the constructions examined by Gernsbacher. This is also consistent with the reading time data described by Garrod et al.

MacDonald and MacWhinney (1990) are sympathetic with the position taken by Gernsbacher (1989) that mechanisms of both antecedent enhancement and non-antecedent suppression result from anaphoric processing. Using a cross-model probe paradigm they also found evidence for enhancement and suppression of the antecedent and non-antecedent respectively. In line with the data reported in Gernsbacher, MacDonald and MacWhinney do not find that these mechanisms produce an observable effect immediately following the pronoun. An effect of antecedent facilitation was first observed at the 250 msec delay point following the pronoun. When the pronoun was referentially ambiguous (i.e. no gender cue), this facilitatory effect was first found at the 500 msec delay point. Suppression of the non-antecedent was found at the 250 msec delay point in their first study but at the 500 msec point in their second. It may be the case that the mechanisms of suppression and enhancement are not driven by the same basic process and so are not bound together in time as they would if they were due to a single process. Indeed, it is possible that what is referred to as the mechanism of non-antecedent suppression arises as a logical consequence following the process of antecedent activation.

2.4.3 Are Pronouns Always Resolved ?

McKoon and Ratcliff (1992) outline the controversial position they refer to as Minimalism (for replies see Garnham, 1992 and Glenberg, 1993). They propose that the only inferences that readers generate when reading a text are those necessary for the maintenance of local

cohesion, or those based on easily available information. Their theory proposes a restriction on the amount of cognitive effort expended by a reader. The position is taken to an extreme in Greene, McKoon and Ratcliff (1992) where they propose that pronouns aren't always resolved although noun phrase anaphors are. They argue that pronoun resolution is not an automatic process and is, to a large degree, under strategic control.

They claim that the processes associated with pronoun reference resolution do not always produce a unique antecedent. The position is based on evidence gathered using the probe task which failed to find any activation differential between the referent and non-referent at any point following a pronoun. This appears at odds with the vast bulk of the literature on pronoun resolution discussed above and, on the basis of our discussion of potential problems with the probe task, should perhaps best be interpreted as further evidence that the probe task isn't necessarily measuring what it is assumed to be measuring rather than evidence that readers aren't resolving pronouns.

2.5 The First Mention Privilege

One phenomenon that has consistently been found by those employing the probe task to examine anaphoric processing is that reaction times to probes corresponding to the first mentioned character in a particular sentence are faster than reaction times to probe words corresponding to other characters also mentioned within that sentence. This has been documented by Gernsbacher and Hargreaves (1988) as the first mention privilege. It is not the case that because the first mentioned character also normally occupies the grammatical subject position or the thematic role slot of agent or that it is the first word in the sentence that gives rise to the phenomenon, but simply that it is the first character encountered within a sentence. This general pattern of data parallels a basic serial position effect known for some time to exist in the memorisation of word lists (Murdock, 1962). This primacy effect is simply that words presented at the start of a list are recalled more accurately than words presented towards the middle. Gernsbacher, Hargreaves and Beeman

(1989) propose that the primacy effect found in probe response times is perhaps because :

'first-mentioned participants are more accessible both because they form the foundations for their sentence-level structures, and because it is through them that subsequent information is mapped onto the developing structure.'

(Gernsbacher, Hargreaves and Beeman, 1989, p. 737; cf. MacWhinney, 1977)

They cite evidence of initial words in sentences being read more slowly than other words in the sentence except the final word (Aaronson and Ferrer, 1983; Aaronson and Scarborough, 1976; Chang, 1980), slower identification of phonemes and words at the beginning of sentences than at the end (Cutler and Foss, 1977; Marslen-Wilson, Tyler and Seidenberg, 1978) and a larger N400 for the first open class word in a sentence than a later one (Kutas, Van Petten and Besson, 1988) in an attempt to support their claim that initial words are used to lay down a foundation onto which subsequent information in a text is mapped. They argue that these effects are due to initial words forming the foundation for comprehension of the sentence in which they occur. An alternative account however would argue that the increase in time spent processing the initial words when read effectively causes them to be more richly represented and so more amenable to subsequent recall or recognition. It is possible that at least part of the reader's task during reading is one of prediction. It should be noted that initial words in a sentence are less constrained and therefore less predictable than those occurring later. Even at the level of syntactic category the category of words later in a clause is more constrained than that of words occupying an initial position. It's possible for a sentence to begin with something other than a Noun Phrase so there is no requirement on a Noun Phrase appearing in sentence initial position. However, given a Noun Phrase at the start of a sentence, a Verb Phrase will usually follow at some later point although not necessarily immediately. Increased processing by subjects of the first few words in a sentence may reflect their low cloze probability relative to those appearing later.

Gernsbacher, Hargreaves and Beeman (1989) report an additional result which again has parallels in the serial position recall memory literature. At a 0 msec probe delay point they actually find evidence of a recency effect; that is, responding to the most recently mentioned character is faster than responding to the first mentioned. At later probe delay points (1400 msec) this recency effect disappears and the first mention privilege returns. They interpret this as consistent with their 'structure building' framework. Given their framework, clauses are processed separately and are integrated only once the processing associated with each clause is complete. At early probe points, readers are still processing the most recent clause and so respond more quickly to probe words which are repetitions of the name of the participant in this most recent clause. Following some delay, readers finish processing the most recent clause and attempt to integrate clauses. At this point, Gernsbacher et al argue, the first clause acts as a foundation for interpreting the second and so a primacy effect resulting in faster response times to the participant mentioned in the initial clause is found.

The position adopted by Gernsbacher et al is similar to the delayed-integration hypothesis proposed by Millis and Just (1994). Under this argument, interclausal relations are not computed until the end of the second clause. In other words it rests on a belief in non-incremental processing of interclausal relations. In essence it is the position adopted by Gernsbacher et al when they argue for a level of clausal processing preceding the computation of interclausal relations. However, this delayed integration argument has been questioned by Traxler, Bybee and Pickering (1997) examining the processing of causal and diagnostic statements (Traxler, Sanford, Aked and Moxey, 1997). Their results indicate that readers don't delay integration of clauses conjoined by 'because' until termination of reading of the second clause. It is possible however that the computation of referential descriptions and referential relations may indeed proceed to some degree in a non-incremental fashion.

2.5.1 A non-psycholinguistic account of the First Mention Privilege

Although Gernsbacher and Hargreaves propose that the findings of the primacy and recency effects in language comprehension are the result of what they term 'structure building', Neath (Neath 1993; Neath and Knoedler, 1994) takes the alternative view that the result is no more than an instance of a general serial position recognition phenomenon. Neath proposes a model that incorporates variables corresponding to the serial positions occupied by experimental items, the retention interval and the interpresentation interval. On the basis of a 0 sec retention interval, his model predicts a recency effect but no primacy effect, while at a retention interval of 2 seconds, there is a prediction of an increase in primacy. His account is driven by focus on the distinctiveness of an item to be recalled, where distinctiveness is defined as the temporal position occupied by an item in a list. To draw the appropriate parallel between the tasks traditionally used in serial order memorisation and the reading task used by Gernsbacher, Neath considers the sentences read by subjects as a list of words. His model certainly suggests that the mechanisms giving rise to the recency and primacy effects reported by Gernsbacher and colleagues are no more than those arising during normal processing of item lists. In other words the level of explanation for what Gernsbacher interprets as psycholinguistic phenomena is at a non-psycholinguistic level of processing. The effects are not necessarily due to the 'structure building' account which she argues for. Of course there may be consequences for language comprehension as a result of this general memory effect which influences the manner in which the discourse representation is laid down, but this would then be very much a secondary phenomenon.

Following the account proposed by Neath, our earlier concern that the probe task may be measuring processes other than those involved in language comprehension appears supported. It is possible that whatever processes the probe task measures correspond directly to the consequences they have on the structure of the readers discourse model, but it's equally possible that the task is measuring a level of processing functionally distinct but related to those levels associated with language processing. If temporally prior, those areas may be involved at a level lower than language comprehension, while if

subsequent, they may be a consequence of those processes associated with language comprehension.

The following sections examined how different types of information are used by the system to interpret anaphors. We initially focus on the role played by structural factors. We argued in Chapter 1 that this information may play some special role in initial interpretation if the anaphor resolution system behaves in a restricted manner.

2.6 Structurally informed accounts of pronoun resolution

In the case of a pronoun that is referentially ambiguous, does the system take advantage of what information it does have available to attempt to resolve the ambiguity when it is encountered ? The accounts summarised below accord some special status to the role of structural information. If we consider that these accounts describe an initial rather than ultimate reference resolution preference we can construe them as being restricted processing accounts with initial processing occurring through reference to only structural information.

One problem with the literature focusing on the role played by structural factors in pronoun resolution is that the distinction between an initial first guess and ultimate resolution is not normally made. We can return to our analogy between anaphor resolution and parsing to separate these aspects. These levels may be informed by different types of information. The first by low level structural information, the second by higher level factors.

In other words, if the anaphor resolution system is modular the following accounts are analogous to principles such as Late Closure and Minimal Attachment in parsing. They may then capture something of initial processing. Recall not only structural information but also semantics and gender information are ultimately utilised by the anaphor resolution system. An initial stage of processing may proceed with reference to only structural factors with these high level factors exerting an influence at a later point.

2.6.1 Parallel Function Strategy

Parallel Function Strategy (Sheldon, 1974) makes a very specific claim about preference for pronominal reference :

A pronoun will be interpreted as coreferential with the noun phrase occupying the same grammatical position in the preceding clause.

So, examples (18) and (21) below will be initially preferred over (19) and (20) as the pronoun refers to the character occupying the same grammatical position in the preceding clause. In Example (18), this is the grammatical object position while in example (21) it is the grammatical subject position.

(18) Mary hit Bill and then Sarah kicked him.

(19) Mary hit Bill and then Sarah kicked her.

(20) Mary hit Bill and then he kicked Sarah.

(21) Mary hit Bill and then she kicked Sarah.

Of course, PFS can be ultimately overridden in light of additional pragmatic information but when no additional disambiguating information is available when the pronoun is encountered, PFS can be employed to make an initial first guess. If the anaphor resolution system is modular, PFS may be equivalent to Late Closure or Minimal Attachment.

2.6.2 Subject Assignment Strategy

Another structurally informed heuristic is Subject Assignment Strategy (SAS). In sentences such as (22) below, SAS predicts that the pronoun will be interpreted as coreferential with the character occupying the preceding grammatical subject position; in this case 'John'.

(22) John saw Bill in the street and he waved.

In Example (13) PFS makes the same prediction as SAS as to which is the preferred referent of the pronoun in the second clause but differs from SAS in predicting reference in Example (23) below :

(23) John saw Bill and Mary waved at him.

In this case, SAS would still predict that 'him' and 'John' are coreferential, while PFS would interpret the pronoun 'him' and 'Bill' as coreferential. This is because both the pronoun and noun phrase occupy the grammatical object position in their respective clauses.

PFS and SAS produce the same results for pronouns occupying the grammatical subject role and consequently can be treated as reasonably accurate heuristics. Hobbs (1976) reports that in texts, 90 percent of pronominal antecedents occupy the grammatical subject position in the preceding clause. Therefore SAS will be correct 90 percent of the time. If the pronoun also occupies the grammatical subject position in the clause in which it occurs, PFS will have the same level of success. Frederikson (1981) provides evidence that sentences containing sentence initial pronouns are read more quickly if the antecedent is the noun phrase occupying the grammatical subject position in the preceding sentence. This is consistent with both SAS and PFS.

If we consider that only structural information is initially available to the system, if the pronoun resolution mechanism wants to make an initial 'best guess', it could do worse than adopt one of these strategies. This is consistent with a restricted account of processing where information other than that which is determined structurally is considered only at a later stage of analysis. When the structurally determined 'best guess' turns out to be incorrect, subsequent revision must occur which may, from a processing perspective, be quite costly.

What experimental support is there for SAS and PFS as processing theories?

2.7 Structurally Informed Accounts as Processing Theories

Stevenson and colleagues (Crawley, Stevenson and Kleinman, 1990; Stevenson, Nelson and Stenning, 1995) have conducted a number of experiments to investigate the relative strengths of contribution of Subject Assignment Strategy and Parallel Function Strategy to the overall mechanism by which pronominal anaphora are resolved.

2.7.1 Crawley, Stevenson and Kleinman (1990)

In general, Stevenson and associates assume that structurally driven heuristics act as weak constraints and can be overridden when other informative factors are present. The consequence of this is a position which assumes that gender, pragmatic and other non structural constraints are considered at an early stage. This is tantamount to the view that such structural information plays no privileged role compared to other factors at an initial stage of the mechanism of pronoun resolution. The only possible position compatible with this view is of a highly interactive constraint based nature. All information would therefore be used without restriction by the system and if non-structural information is sufficient to uniquely select an antecedent structural factors will exert no influence. Crawley et al cite an example from Broadbent (1973) as evidence that without full knowledge of the content of a sentence, readers adopt a structurally informed heuristic to resolve reference of ambiguous pronouns; see Example (24) below.

(24) The feedpipe lubricates the chain, and it should be adjusted to leave a gap half an inch between itself and the sprocket.

The pronoun, 'it', following the conjunction is typically interpreted as coreferential with the noun phrase, 'the feedpipe'. This is compatible with both SAS and PFS.

Crawley et al find evidence supporting the general conclusion that SAS plays some role in the mechanism of pronoun resolution but they find no support for PFS. It is not possible to draw any conclusions about when SAS may be used, i.e. whether it is used initially or at some later point if other factors to facilitate interpretation of the pronoun are absent. It is also not clear to what extent this conclusion

can be generalised to sentences containing alternative syntactic structures.

2.7.2 Smyth (1994)

Smyth (1994) defends the role that Parallel Function Strategy plays in pronoun resolution. He focuses on the structure of the materials examined by Crawley et al and suggests that only a small number display strict parallelism. Smyth considers sentences where the two conjoined clauses have the same underlying syntactic structure to be strictly parallel. In Example (25) below, the two conjoined verb phrases possess the same syntactic frame. However, an example item from the list examined by Crawley et al can be seen not to be strictly parallel (Example 26 below).

(25) William[i] hit Oliver and he[i] slapped Rod.

(26) Liz tried to catch Melanie[i] and Frank chased her[i].

The conjoined phrases in example (26) are not of the same syntactic form. Smyth argues that only when strict parallelism is adhered to, will PFS be used by the system. When this parallelism criterion is not met, alternative strategies such as SAS will be used by the reader. In four production experiments, Smyth provides evidence that PFS plays more of a role in the anaphor resolution mechanism as a whole than the data of Crawley et al would suggest. In his first experiment, Smyth shows that PFS exerts an observable influence on interpretation only when sentences are strictly parallel. He suggests the reason why this parallelism must be maintained between the conjoined clause as being due to a mechanism not unlike syntactic priming (e.g. Bock, 1986; Branigan, Pickering and Stewart, 1997). Smyth postulates an increased level of activation in the syntactic frame of the first clause will cause the second to be processed more easily if it possesses the same syntactic structure. In other words, if two syntactically identical conjoined clauses are read, the pronoun in the second will be interpreted as coreferential with the entity occupying the same grammatical position in the first. Although this is the general PFS account, Smyth reasons that the mechanism

underlying this is an example of syntactic priming. Whilst there may be some truth in this explanation for why parallelism should seem to exist as a strategy, there are serious theoretical problems associated with adopting it in this strong form. In the case of example (27) below, the processor can only know that the conjoined phrases exhibit syntactic parallelism after the second clause has been read.

(27) John hit Mary and he kicked Bill.

If the parallel nature of the sentence can only be determined at the end of the second clause, in this case at the end of the sentence, then either assignment of the pronoun's referent must be delayed until this second clause has been, at least, superficially processed, or some initial strategy must be used at the moment the pronoun is first encountered which may then be modified by a parallelism preference at a later point. This position is equivalent to a 3-stage theory of anaphor resolution. A first stage produces a decision using very low level information, a second assesses whether an interpretation using parallelism as a basis is appropriate, while a third considers higher level semantic and plausibility information.

Although Smyth (1994) provides evidence that PFS plays a role in pronoun resolution offline, it is not obvious whether this is the result of some relatively late strategy adopted by readers when required to make a decision given an impoverished input stream, or whether it genuinely reveals something about the structure of the underlying mechanism by which pronouns are resolved normally.

Another important issue to consider at this point is how exactly readers interpret clauses conjoined by the connective 'and'. In the above studies it has been assumed that the ambiguity doesn't interfere with reference resolution. There are a number of possible interpretations associated with this type of conjunction. In Example (28) below, 'and' could be interpreted as 'and then'. The reader may understand the connective as one operating to temporally relate the two described events in a serial, non-causal and non-consequential manner (Lascarides, Asher and Oberlander, 1992). Alternatively the reader may causally relate the two events if doing so is pragmatically

likely. Depending on what interpretation is given to the connective, there may be different consequences for resolving the anaphor.

(28) John hit Bill and he hit Bob.

As the mental model is constructed to represent the actions described in this sentence, something approximating a description similar to 'John first did x and then he did y' may occur; where x is hit Bill and y is hit Bob. Example (29) below could also be interpreted in an analogous fashion. However, it is more likely that it is the consequential nature of the relationship that becomes the means by which the contents of the two clauses are related.

(29) John hit Bill and he fell over.

In other words the mental model represents the sentence as being of the nature 'John first hit x and as a result, x did y'; where x is Bill and y is fell over. There's certainly no principled reason to rule out a reader constructing a discourse representation in such a way to arrive at an outcome of this nature. The manner in which the reader interprets the connective may have a differential effect on anaphoric processing depending on whether, in the above example, a simple temporally serial or consequential interpretation is made. Without knowing quite what relationship the reader is interpreting as linking the two clauses it is impossible to consider what aspect of general knowledge they are drawing upon. Under certain circumstances one interpretation (resolution preference) may simply be more plausible than another and lead to one type of interpretation of the clausal relationship. In other words, plausibility determined by the type of relationship construed as relating the two events rather than any structural heuristic may be driving readers' interpretation.

2.7.3 Stevenson, Nelson and Stenning (1995)

In a number of on-line studies, Stevenson, Nelson and Stenning (1995) modified the position taken by Crawley et al by accepting that both PFS and SAS may inform the pronoun resolution system depending on the structure of the sentence. The position outlined by

Smyth suggests that PFS plays an important role when the two conjoined clauses possess identical syntactic frames. When this criterion is not met, SAS may inform the system.

Again we return to our analogy with positions on parsing. If the anaphor resolution system behaves in a modular fashion, PFS and SAS can be interpreted as equivalent to Minimal Attachment and Late Closure informing an initial stage of processing. If the right conditions are met PFS will inform this stage, otherwise SAS.

Simply because they investigated on-line reading times rather than effects on production, Stevenson et al's experiments are perhaps more informative than those conducted by Smyth (1994) for the evidence that they provide supporting the existence of structurally motivated strategies used in comprehending pronouns. However, all we can really accept at this point is that some part of the mechanism is sensitive to these structural factors. No theoretically motivated processing account has yet been proposed which can accommodate the influence of structural information in conjunction with other informative factors in the input. Indeed, the results of Crawley et al (1990) would suggest that when reference can be determined using non-structural factors, e.g. when the preceding character proper names are disambiguated by gender, PFS and SAS have little or no influence. We earlier raised a caveat over what interpretation we may put on this finding. There's certainly no reason why a structurally informed heuristic may not be used initially by the system, the influence of which quickly diminishes as alternative sources of disambiguating information become available. This is equivalent to a two-stage model of anaphor resolution. The first stage is informed by low level structural factors while the second by higher level semantic information.

2.8 Centering Theory

Centering Theory (Gordon, Grosz and Gilliom, 1993) is perhaps the most fully developed of the structurally motivated attempts to account for the nature of the referential links between discourse segments. One of the aims of those working with Centering Theory is

'to show that structural features of discourse contribute to pronoun interpretation independently of knowledge-based processes.' (Gordon and Searce, 1995, p. 315). There is a parallel argument in certain parsing theories postulating a level of sentence processing independent of general knowledge (e.g. Frazier, 1979) and is equivalent to an initial stage of the anaphor resolution system informed solely by low level structural information. Centering Theory focuses on the role of local discourse structure on maintaining inter- and intra-sentential referential cohesion.

Each utterance within a discourse is proposed to possess 2 sets of centres consisting of entities corresponding to potential referents. The backward looking centre (Cb) contains entities referentially linking an utterance with that preceding it. The forward looking centre (Cf) contains a ranked list of entities to which reference can be made by those utterances following. In the following example the Cf of the first utterance consists of the characters {John, Bill} while the Cb of the second utterance also consists of the characters {John, Bill}.

(30) John waved at Bill. He had known Bill for many years.

Centering proposes that the backward looking centre in an utterance, if it is also the most highly ranked entity in the forward looking centre in the preceding utterance, must be realised as a pronoun for it to contribute to coherence. Gordon, Grosz and Gilliom (1993) report a repeat name penalty found when the Cb was realised as a repeat name rather than as a pronoun when that entity was also a member of the preceding forward looking Centre. In other words, the second utterance in example (31) below will be read more slowly than the second utterance in example (30) above.

(31) John waved at Bill. John had known Bill for many years.

This penalty was found only for the entity in an utterance occupying the grammatical subject role. The penalty may simply arise under circumstances when a character is highly accessible and referred to using a repeat name. There may be something special about the grammatical subject of a sentence giving rise to this finding. The

character in the grammatical subject position of a sentence may be highly accessible because of its grammatical prominence. This position has much in common with the first mention effect reported by Gernsbacher (see section 2.5 above). Of course, she reports that this wasn't restricted to the entity occupying the grammatical subject position but there may be a common mechanism responsible for both of these phenomena.

Although there is ample experimental support for the predictions made by Centering, the evidence can also be interpreted with respect to other accounts in the literature on referential cohesion. The repeat name penalty central to centering theory has been documented by Vonk, Hustinx and Simons (see section 2.2.4 above). What differs between the framework proposed by Centering Theorists and by Vonk et al is the explanation put forward to account for the finding. While Vonk et al focus on the functional role that over-specification appears to fulfil and provide an explanation at a level of processing within the language system, Centering Theorists merely describe the conditions (defined within their framework) necessary for such a penalty to be accrued. That the explanatory power of their framework extends no further than the structural properties associated with the local discourse is intrinsic to the account they wish to propose but it does leave Centering weak as a theory of how the anaphor resolution system ultimately behaves if there are important determinants relevant to referential cohesion such as semantics and pragmatics which operate at a level other than the local discourse structure. Centering may occupy a position during an initial stage of anaphor processing but it cannot account for how other higher level factors influence ultimate resolution.

The preceding section summaries those positions within the literature suggesting that some level of the pronoun resolution system is informed, at least in part, by structural information. Those working in the area have so far failed to outline where in the anaphor resolution system as a whole structural information exerts an influence. By adopting an analogy with theories of parsing we propose that if the anaphor resolution system behaves in a modular fashion, initial processing will be informed by solely structural

factors. The way in which these factors are used by the system is reflected by the way in which PFS, SAS and Centering operate.

There is a level of structure in text that stands between low level information and higher level semantic and pragmatic factors. The following section provides an overview of a position suggesting that this level of thematic structure also informs the pronoun resolution system.

2.9 Explanation at the Level of Thematic Roles ?

Thematic roles can be considered as semantic constructs forming the link between the syntactic and discourse level aspects of representation. Carlson and Tanenhaus (1988) suggest that thematic roles are represented in the reader's discourse model relating entities mentioned in a discourse to the roles they play at points within that discourse. Verbs possess thematic frames which contain the roles that characters and phrases play in a text. These frames include the roles Agent, Patient, Stimulus, Experiencer, Goal, Theme and Source. For example, the verb 'fascinate' possesses the roles of Stimulus and Experiencer. In example (32), the character 'John' occupies the Stimulus role while 'Bill' occupies the Experiencer role.

(32) John fascinated Bill.

Similarly, action verbs such as 'hit' possess role slots of Agent and Patient. In example (33) following, 'John' occupies the role of Agent while 'Bill' occupies the role of Patient.

(33) John hit Bill.

Stevenson, Crawley and Kleinman (1994) propose an account of pronoun resolution determined by the thematic roles associated with particular verbs. They do not state whether this amounts to the equivalence of a first stage of processing an anaphor or whether they consider it to correspond to how an anaphor is ultimately resolved. As higher level pragmatic factors can always override lower level

ones, it seems best to interpret their position as analogous to a first guess mechanism.

Stevenson et al examined State verbs, Action verbs and Goal-Source verbs. State verbs link participants using a psychological state relation, as in the case of *John fascinated Bill*. The character 'John' as well as being the grammatical subject of the sentence also occupies the thematic role Stimulus, with 'Bill', the grammatical object occupying the role of Experiencer. Action verbs link characters involved in an action. For action verbs such as hit in the sentence *John hit Bill*, John occupies the thematic role of Agent while Bill the role of Patient.

John gave the book to Bill is an example of a sentence containing a Goal-Source verb with 'John' occupying the role of Source, 'Bill' the role of Goal and 'book' acting as the Theme. Stevenson et al argue that in sentences of these types a reader's focus is on the character occupying a particular thematic role. This focus pattern is considered to influence pronoun assignment. In the absence of a connective Stevenson et al report that in a production task there was a pronoun continuation preference to the character occupying a particular thematic role. The thematic role Goal was preferred over Source, Patient preferred over Agent and Stimulus over Experiencer. In the presence of a connective, this pattern was modified. In the case of *so*, Goal again was preferred to Source, Patient preferred to Agent and Experiencer preferred to Stimulus. In the case where the connective was *because*, the Goal-Source effect disappeared to be replaced by a Source preference, there was a reduction in the Patient preference and in the case of State verbs, a preference for the Stimulus was found. Stevenson et al argue that this is due to the emphasis that the connective places on different aspects of the described event. They claim *because* focuses on the causal relationship between the two participants while *so* focuses on the consequences.

There are several problems with the Stevenson et al position. The first is the circularity in using thematic roles as an account of focusing preferences. The thematic roles are defined by higher level semantic factors based on properties associated with the individual verbs. The

nature of the roles proposed to be occupied by the noun phrases associated with these verbs may in part be driven by the same factors underlying the focus of the event. In other words, a factor related to focus inherent in an event will be used to assign thematic roles, the particulars of which are then used by Stevenson et al to predict focus.

Their examination of differences between connectives is interesting. The focus pattern associated with connectives of particular types seems to interact with the pattern associated with a particular verb. In the case of *because*, focus is placed on what could be considered to be the causer of the event, the Stimulus in the case of state verbs and the Agent in the case of action verbs. In the case of Action verbs Stevenson et al report a preference to continue with reference to the character occupying the role of Patient, although this may be restricted to the types of verbs they examined. This does suggest that other factors are contributing to subjects' continuation preferences.

For *so*, focus appears to be placed on the thematic role most likely to be affected by the consequences of the described event. However, *so* can be given both a consequential (John tackled Bill so he didn't score) and a motivational reading (John tackled Bill so he couldn't score) which does not seem to have been controlled in the Stevenson et al experiments. These two interpretations will have different associated patterns of focus so it is not clear on the basis of the data reported by Stevenson et al how these may individually interact with verb driven focus.

Although there is certainly a correlation between particular thematic roles and focus, this is not perfect. This correlation is inevitable if the definition of a verb's particular thematic structure is partly influenced by focus information as the correlation is between thematic roles (influenced by focus) and focus per se. We shall return to this in the following chapter as a similar account has been put forward to explain the nature of implicit causality verb biases.

The next section provides an overview of evidence for the influence of non-structural factors on anaphor resolution.

2.10 The Influence of Non-Structural Factors on Anaphor Resolution

In the same way that researchers in parsing have been interested in examining the influence on parsing of the context in which sentences are embedded, it is possible to examine the nature of the influence of equivalent contextual factors on anaphor resolution. In parsing, one of the central interests has been in the time course of the influence of non-syntactic information. As we mentioned above this interest was generated by the theoretical ground mapped out by modularity. If we find an initial influence of factors considered to lie outside the parsing module we can question the claim for an initial restricted stage of sentence processing. If we find an initial influence of non-structural factors considered to lie outside the anaphor resolution module we can similarly question the claim for an initial restricted stage of anaphoric processing. As the theoretical concerns for anaphor resolution haven't been specified in as clear a manner as those for parsing, the question of time course does not form one of the central issues within the area. With several notable exceptions (e.g. Garrod et al), the focus has been on what factors influence anaphor resolution rather than when.

2.10.1 When does semantics influence anaphoric interpretation ?

Consider the following example taken from Hirst and Brill (1980) :

(34) John stood watching while Henry fell down some stairs. He ran for a doctor.

(35) John stood watching while Henry fell down some stairs. He thought of the future.

Although readers may take advantage of one of the structurally determined heuristics such as SAS mentioned above to initially interpret the pronoun, the disambiguating information which indicates to which character the pronoun should be interpreted as

coreferential with does not appear until after the pronoun. In example (34) Hirst and Brill claim the information is clearly disambiguating on pragmatic grounds. It is unlikely that someone having falling down stairs would then stand up and run for a doctor. In example (35), Hirst and Brill claim the following contextual information is compatible with the pronoun being interpreted as referring to either character. Even under conditions where a gender contrast is present and the pronoun refers to its antecedent unambiguously, Hirst and Brill still found an effect of subsequent pragmatic information with faster response times to the pronominal antecedent following a continuation containing information that was more pragmatically consistent with that character compared to following information that wasn't more likely to be associated with one character than the other.

The question we can ask ourselves is when does this information influence interpretation of the pronoun? Is an initial first guess made on the basis of one of the structural heuristics and then compared against pragmatic information or is interpretation of the pronoun delayed until disambiguating information necessary for successful interpretation is encountered ?

2.10.2 When does gender information influence anaphoric processing ?

Information contained in the pronoun itself can also be used to facilitate interpretation.

(36) John waved at Mary as she passed in the street.

(37) John waved at Mary as he passed in the street.

As it is gender marked, regardless of whether the pronoun refers to the first or second mentioned character it does so unambiguously. The contextual information following the pronoun does not bias interpretation either way as the two are logically equivalent, i.e. John passing Mary in the street is equivalent to Mary passing John. When does gender information inform the system ?

Garnham, Oakhill and Cruttenden (1992) suggest that whether this information exerts an influence is under the control of the reader and that given certain reading strategies the reader appears insensitive to gender information. In several experiments they manipulated whether the readers' focus of attention was on pronoun resolution or on other aspects of a text. In cases where attention was diverted from the issue of pronoun resolution, gender cue did not seem to exert an influence on subjects' reading times suggesting that sensitivity to the cue was not automatic. This suggests that gender information does not exert an early influence in interpreting a pronoun.

Stevenson and Vitkovitch (1986) examined the influence of pragmatic factors on anaphor resolution under conditions where reference assignment was restricted by syntactic constraints. In the case of sentences containing elliptical reference such as (38) below, the referent is restricted to that which is occupying the grammatical subject position in the preceding clause (Hankamer and Sag, 1976). There is no such restriction where this reference is not subject to syntactic constraints as in example (39).

(38) Stan was much better at chess than Fred and won every single game.

(39) Stan was much better at chess than Fred and he won every single game.

Stevenson and Vitkovitch contrasted sentences such as the above, which they claim contains an informative verb, with ones such as example (40) below which they claim contains an uninformative verb. Their definition of informativeness seems to depend on whether or not given the first phrase, readers are likely to generate an inference comparable to the information contained in the second. From knowing that the character 'Stan' is much better at chess than 'Fred' it follows that 'Stan' is likely to beat 'Fred' if they play a game of chess. In example (40) below, however, knowing that 'Anna' lent 'Felicity'

the iron does not invite the inference that she forgot to give 'Felicity' the instructions.

(40) Anna lent Felicity the steam iron and (she) forgot to give the instructions.

Stevenson and Vitkovitch found that even in cases where there are clear syntactic constraints on reference (ellipsis condition), the informativeness of the verb was nevertheless important. Response times to the pronominal antecedent in a naming task were facilitated for those sentences containing informative verbs. By manipulating gender cue, Stevenson and Vitkovitch also report that the pragmatic information following the pronoun, which is unambiguous given this gender cue, affects the speed of assignment. In other words, there is an effect of pragmatics on response times even under conditions where a pronominal antecedent can be unambiguously identified using gender information. Stevenson and Vitkovitch take this as evidence that reference assignment is delayed.

Stevenson and Vitkovitch base their conclusion on both response time and reading time data. Their data cannot be interpreted as providing unequivocal support for delayed reference assignment. In fact, their pattern of results could primarily reflect the readers' difficulty associated with pragmatically less plausible sentences. In their uninformative verb conditions, the second clause, being unconnected in any obvious way with the first, may simply form a less plausible sentence when combined with the first thus leading to a general processing difficulty associated with comprehension of the sentence as a whole.

What Stevenson and Vitkovitch's data indicate is some manner of influence of pragmatic information on sentence reading time. As described above, a similar influence was found by Hirst and Brill (1980) for examples such as (41) and (42), repeated below.

(41) John stood watching while Henry fell down some stairs. He ran for a doctor.

(42) John stood watching while Henry fell down some stairs. He thought of the future.

Hirst and Brill employed a response time measure where subjects were required to press a button corresponding to which character they thought was the pronominal referent. Supporting Stevenson and Vitkovitch, Hirst and Brill found an influence of following pragmatic information on response time. They argue for the position that integration occurs in parallel with reference resolution. Again it is possible that a general cognitive difficulty follows as a consequence from the second sentence containing the less informative continuation. It certainly seems the case that readers are unable to ignore the pragmatic information contained in the text that they are reading. Response time is therefore effectively an index of both processes associated with identifying an anaphoric antecedent and those associated with subsequent integration.

What the results reported by Hirst and Brill and by Stevenson and Vitkovitch do indicate is that pragmatic factors influence integrative processing and that they do so automatically. This may occur either paralleling or following antecedent identification. The reading time and response time measures adopted are not sensitive enough to indicate at which point within anaphoric processing pragmatically driven influences are manifested.

Summary

The goal of this chapter has been twofold : to extract the dimensions from the experimental literature along which to categorise recent research perspectives and to sketch the framework for a restricted account of pronoun resolution placing special emphasis on structural or structurally determined information recoverable from the sentence in which a pronoun occurs. We started by examining the functional differences in discourse terms between anaphors of different types. Some time was then spent examining evidence that aspects of pronoun resolution are not completed immediately. The Gernsbacher (1989) probe data form perhaps some of the strongest evidence that pronouns are not resolved immediately. Recall she found an

activation differential between antecedent and non-antecedent at no point earlier than the end of the sentence.

We outlined the first mention privilege, documented by Gernsbacher, although we raised the possibility that it can be accounted for at a non-psycholinguistic level of explanation. Following that we summarised accounts of pronoun resolution which are informed by low level information. This information is generally proposed to be structural, although in the case of Stevenson et al's account, thematic information plays an important role.

We then drew on evidence from Hirst and Brill and from Stevenson and Vitkovitch indicating that non-structural factors do play some role in anaphor resolution although their influence may be limited to the level of integrative processing.

2.11 A Two-stage Model of Anaphor Resolution

In light of the positions outlined in Chapter 1 and the research summarised in this chapter, we now briefly sketch a two-stage model of anaphor resolution. The model is motivated by both the evidence gathered from research examining anaphoric resolution as discussed above and accounts of parsing which have proposed that an initial stage of the parser behaves in a restricted manner operating only over a well defined set of information.

We propose that two stages are involved in resolving anaphors. Taken together the overall goal of the system is to maintain referential cohesion between units of text. However, when taken individually the goals of the two stages differ somewhat.

Before information following an anaphor can be integrated with information previously read, the antecedent of that anaphor must be identified. We call the process by which this is accomplished the co-indexation stage and argue that it is the goal of the first stage of the system. At this point in the system nothing more than an associative link is established between anaphor and antecedent. We propose that this stage is informed by low level factors and behaves in a restricted

manner. Following this stage of co-indexation the information following the anaphor is integrated with what has been read up to that point. Integration cannot occur without prior identification of the anaphoric antecedent.

We turn first to the process of co-indexation. Analogous to the strategies of Minimal Attachment and Late Closure in parsing, we propose that initial processing within the anaphor resolution system (our stage of co-indexation) will be informed by equivalent structural heuristics. The result of this stage of processing is to make an initial first guess as to which potential antecedent is correct which may be revised in light of higher level factors such as semantics and pragmatics. We propose that these higher level factors do not inform the initial stage of anaphoric processing. Only low level information plays a role during co-indexation. In example (43) below, the reader will initially interpret the pronoun as coreferential with the character 'John' (as predicted by both PFS and SAS).

(43) John fascinated Bill because he was easily fascinated.

Only once the information following the pronoun is processed does it become apparent that the pronoun should be interpreted as coreferential with the character 'Bill'. In example (44) below, the initial first guess turns out to be correct in light of the information following the pronoun.

(44) John fascinated Bill because he was very interesting.

If a strategy such as SAS or PFS is being employed at an initial stage of processing, then in constructions such as (43) and (44) the pronoun will be initially interpreted as coreferential with the character 'John'. SAS and PFS both predict this assignment. In Example (43) there will be a processing difficulty when it becomes apparent that the pronoun refers to the character 'Bill'.

So then, we propose that co-indexation behaves in a restricted manner without reference to higher level semantic factors. We propose that it is during the second stage of processing that higher

level semantic and pragmatic factors come to influence processing. The decision arrived at by the co-indexation stage on the basis of structural information will be accepted or rejected depending on the result of this second stage. This two-stage account makes a very clear prediction. If an attempt is made to co-index a pronoun with its antecedent using only a restricted set of information, 'sentences containing reference which ultimately turns out to be to the character predicted by both SAS and PFS will always be read quickly regardless of other factors.

However, what would happen if it was the case that co-indexation proceeded in an unrestricted manner with reference to all information that was potentially informative ? As all constraints influence processing at the same time, any principles such as SAS or PFS which may be operating could be overridden if other constraints are strong enough. For sentences such as (43) and (44) above there should be no difficulty in reading a sentence containing a pronoun referring to either the first or second mentioned character provided that cues indicating that this is appropriate are present. These cues may be verb semantics or pragmatic information following the pronoun.

To summarise then, if pronouns are initially co-indexed in a restricted manner, this stage will be informed solely by structural information. If co-indexation is unrestricted, both low and high level factors will exert an influence on processing. In other words, low and high level factors act as multiple constraints (cf. Trueswell, Tanenhaus and Garnsey (1994).

With respect to the question of the time course over which these two stages of anaphoric processing occur, we can ask ourselves what happens under conditions where processing is delayed. We propose that the discreteness between the stages of co-indexation and integration of our 2-stage model is preserved. The stage of co-indexation may occur at some point following the reader encountering the pronoun but we suggest that the information which informs this stage doesn't change as a function of this delay. If we find evidence that co-indexation behaves in a restricted manner then we argue that

that level of processing will always proceed in a restricted manner regardless of any overall delay in processing which may potentially allow time for other informative factors to be encountered. This is central to the proposal of modularity as a mechanistically restricted account as opposed to modularity if it had been formalised as a temporally restricted account. The nature of co-indexation doesn't change under circumstances where processing of the anaphoric expression is delayed. In the same way that the type of information used to inform co-indexation doesn't vary as a function of delay, we also argue that the outcome of this stage doesn't vary either. This outcome is simply one of establishing an associative link between anaphor and antecedent.

We propose that the stage of co-indexation must always precede any stage of integration. The motivation behind this claim is that before any information can be related to the appropriate antecedent, that antecedent must be identified. Before this identification, integration cannot occur. As the ultimate goal of the co-indexation stage is one of identifying the appropriate antecedent we suggest it is reasonable to propose that this stage always occurs prior to semantic integration regardless of whether the system as a whole exhibits delayed processing. Integration cannot take place without co-indexation.

So far in this section we have only described processing of pronominal anaphors. Might processing of repeat name anaphors be any different? Given the data of Garrod et al, it certainly seems to be the case that there are important differences in how anaphors of different types are treated. Recall, Garrod et al proposed that while there may be a delay in interpreting pronouns, in line with Ehrlich and Rayner's (1983) position, it appears that repeat name anaphors are interpreted immediately. Consider examples (45) and (46) below.

(45) John blamed Bill because Bill broke the window.

(46) John blamed Bill because John hated Bill.

In both (45) and (46) the referent of the anaphoric expression can be identified as soon as the anaphor is read. The repeat names are

coreferential with the characters denoted by the preceding occurrence of the names. No further information is necessary for this stage of co-indexation. For repeat name anaphors we propose that co-indexation between anaphor and antecedent occurs as soon as the anaphor is encountered and is informed by low level factors. Only later does higher level information exert a processing influence as the information following the anaphor is integrated with what has been read previously. As the antecedents of repeat name anaphors appear to be identified as the anaphors are read, we propose that initial processing will operate in a restricted manner without reference to higher level semantic factors.

In summary, there are two dimensions to our two-stage account. The first concerns how low and high level factors inform processing of anaphors and focuses on whether it is best to construe aspects of the system as operating in a restricted or unrestricted manner. The second concerns differences between the treatment of pronominal and repeat name anaphors. As mentioned in the preceding paragraph, it is possible to identify a repeat name anaphor's antecedent using low level information. It may therefore be the case that processing of pronominal and repeat name anaphors is different with respect to how and when low and high level factors are employed by the system to facilitate interpretation of the anaphor. Low level information is more informative in the case of repeat name anaphors than when the anaphor is a pronoun.

Advances in parsing have been made by examining the point in time at which non-structural factors influence processing. A similar examination of how such high level information influences processing of pronominal and repeat name anaphors should allow us to distinguish between the possible positions outlined above. In other words, do aspects of the anaphor resolution system behave in a restricted or unrestricted manner ?

The following chapter contains a summary of the influence of one type of non-structural information, implicit causality, within the anaphor resolution system. During the summary we shall again raise the question of the time course of the influence of such high level

factors generally, and implicit causality information specifically, on anaphoric processing.

CHAPTER 3

3.1 Introduction

This chapter provides a summary of experimental work focusing on the processing influence of a property associated with a particular set of verbs known as implicit causality. Following the groundwork established in the preceding chapter, where appropriate we shall relate empirical investigations of the influence of implicit causality to the question of time course of processing. We shall also interpret the data with respect to the light they can shed on whether it is more accurate to construe the anaphor resolution system as operating in a restricted or unrestricted manner. At the end of Chapter 2 we proposed that if the system behaves in a restricted fashion we expect to find evidence for high level semantic factors at a relatively late stage of processing.

3.2 Implicit Causality

Implicit Causality, first documented by Abelson and Kanouse (1966; Kanouse, 1972) and Garvey and Caramazza (1974), is a property associated with a particular class of verbs which, provided the appropriate conditions are met, exerts an influence on anaphoric processing. Consider example (1) below :

(1) John blamed Bill[i] because he[i] broke the window.

(2) John[i] fascinated Bill because he[i] was very interesting.

On the basis of morphosyntactic information, the pronoun in the subordinate clause is referentially ambiguous. It is pragmatically disambiguated by the information following it. In example (1) the pronoun is coreferential with the second mentioned character 'Bill', while in example (2) it is coreferential with 'John', the first mentioned character. There are certainly no restrictions, say, in terms of number or gender which limit reference in such a way. The subordinate clause in Example (1) could be replaced with 'he was in a bad mood.' where the pronoun is now interpreted as coreferential with 'John' and

similarly the subordinate clause in Example (2) could be replaced with 'he was easily fascinated.' where the pronoun is now coreferential with 'Bill'. Garvey and Caramazza considered that verbs such as 'fascinate' and 'blame' possess a bias encoding something of the causal structure underlying the events which they describe. They consider this an implicit causal bias reflecting the probable direction of cause. In sentences such as the above, the verb 'blame' is considered to bias towards the second mentioned character and is classed as an NP2 biasing verb. Conversely, the verb 'fascinate' biases towards the first mentioned character as the causer and so is classed as an NP1 biasing verb.

Although an understanding of the causal link between the characters denoted by the noun phrases is necessary for ultimate understanding of the sentence, the question of interest is to what extent information about the causal structure exerts an influence on processing.

Before examining evidence for implicit causality exerting a processing influence within the human comprehension system we shall first summarise evidence that people are sensitive to such information in language production.

We will spend the bulk of this chapter focusing on the processing consequences associated with implicit causality verbs. Firstly though we shall present a brief overview of other work reported in the literature which has also adopted the implicit causality phenomenon as a topic of examination, although from a perspective other than that associated with processing.

3.3 Implicit Causality from Non-processing Perspectives

A certain amount of research examining mechanisms that may be responsible for the implicit causality effect has either adopted a social psychology perspective emphasising the attributive nature of verbs exhibiting the bias, or a perspective relating these verbs to what is claimed to be an associated linguistic phenomenon. This perspective has chiefly adopted the line of argument that there is some form of

(interesting) relationship between implicit causality verbs and adjectives which can be derived on the basis of those verbs.

Initially we shall focus on this second view before providing a general overview of the social psychology perspective. Although the experimental examinations of the relationship between implicit causality verbs and other linguistic properties are not centrally concerned with dimensions related to processing, the results of some of the experimental techniques employed are relevant for aspects of processing and provide critical insight into the nature of the effect. Of particular relevance are the data reported in Au (1986), summarised in the section following.

3.3.1 The Whorfian Hypothesis - What language reveals about cognition.

The Whorfian hypothesis, that language determines thought, has been examined by several researchers with respect to implicit causality. Brown and Fish (1983) examined the idea that adjectives derived from verbs exhibiting implicit causality could be used to predict what type of bias the verb possessed. The verb 'liked' in example (3) below biases towards the character occupying the second noun phrase (NP2).

(3) John liked Bill because he was very helpful

The adjective 'likeable' can be used to describe the character 'Bill', but in English there is no comparable adjective such as 'likeful' to describe the character 'John'. In the case of implicit causality, the Whorfian hypothesis can be construed as proposing that the sorts of adjectives we have in English to describe characters engaged in relationships such as in example (x) above influences how we interpret the causal structure of these relationships. In the words, the existence of the adjective 'likeable' leads us to look for a causal explanation for the described event at the level of the character to which we can apply this adjective.

This mode of explanation was ultimately rejected as an explanation by Brown and Fish (1983) who instead adopted a view similar to those emphasising the importance of the thematic roles associated with verbs (see Stevenson section 2.9. above). They argue that certain thematic roles are interpreted as more important than others for indicating the locus of the causal relationship. Consistent with Stevenson et al and Au (see below), Brown and Fish demonstrate that the character occupying the Stimulus role is considered to possess greater causal weight than that occupying the Experiencer role. However, in contrast to Stevenson et al but again consistent with Au, Brown and Fish suggest that for Action verbs, greater causal weight is given to the Agent than the Patient. Recall that Stevenson et al proposed that the character occupying the Patient role was preferentially interpreted as the pronominal referent. It could be argued that the determinants of the causal role and the preferred pronominal referent differ in some crucial way but as the entities occupying these roles are both at least partly determined by the same factors, it is not obvious where that difference might lie.

Brown and Fish initially show that to some degree the adjectives derived from the verbs they examined are usually associated with the thematic role possessed by the character considered as possessing the greater causal weight. Adjectives derived from action verbs are usually attributive to the Agent, while those derived from state verbs are usually attributive to the Stimulus. As stated above, the adjective 'likeable' can be derived from the verb 'like' to describe the Stimulus, but there is no comparable adjective 'likeful' with which to describe the Experiencer. Brown and Fish reject this as an instance of the Whorfian hypothesis, instead suggesting that both implicit causality thematic role preferences and the type of adjectives derived result from the way in which people structure and interpret events in the world.

Hoffman and Tchir (1990) also focus on the nature of adjectives derived from implicit causality verbs although they take the view that the finding is an example of the Whorfian hypothesis. They interpret their data, similar to that reported by Brown and Fish (1983) 'as showing that the way in which people think about

interpersonal causality is related to, and perhaps affected by, the content of the interpersonal lexicon.' Greene and McKoon (1995) empirically tested the claim that the nature of the derived adjectives plays an important role in influencing how people represent events described in such a way. They examined NP2 verbs which possessed either derived adjectives that could be applied to the character occupying the grammatical subject or the character occupying the grammatical object role. Using a probe task they found no evidence for an effect of the type of derived adjective on the accessibility of the character to which that adjective could be applied. Rather, they simply found an influence of whether the verb was NP1 or NP2 biasing. Although at some level there may be a relation between the direction of a verb's implicit causality bias and the adjectives derived from that verb, this does not appear to exert an on-line processing influence.

Au (1986) examined to what extent people are sensitive to implicit causality biases even under experimental conditions where attention is not explicitly focused on the dimension. She presented subjects with a list of verbs and asked them to sort them into categories using whatever criteria they felt were appropriate. Subjects seemed to take implicit causality into account when sorting the verbs, demonstrating the high degree of salience associated with the property. When subjects were asked to attribute cause to one of the two participants in sentences containing implicit causality verbs, they attributed cause to the Stimulus for state verbs and to the Patient for interpersonal or action verbs. This supports the finding of Brown and Fish (1983).

3.3.2 Attribution Theory

Several researchers propose an account of implicit causality bias driven by attribution theory (Heider, 1958; Kelley, 1967) suggesting a role for information that may covary with individuals. Knowing that 'John likes Bill' could either be interpreted as arising because of 'Bill' being likeable or as arising because 'John' likes people generally. If there is high consensus between other individuals, it is the case that many other people also like 'Bill'. If there is high distinctiveness associated with the individual 'Bill', it is the case that 'Bill' is one of

the few people that 'John' likes. Under conditions of high consensus and high distinctiveness, it is likely that the locus of cause lies with 'Bill'. Under circumstances where there is low consensus and low distinctiveness, it is likely that the locus of cause lies with 'John'. Rudolph and Forsterling (1997) propose that it may be the case that knowledge about the world will constrain the flexibility of the covariation information thus resulting in a bias associated with verbs (assuming that those verbs are presented in sentences divorced from additional informative context).

3.3.3 Conclusion

Questions addressing factors such as those outlined in the preceding section are certainly reasonable and interesting questions to ask. The nature of the mechanism responsible for giving rise to the verb biases undoubtedly requires addressing. What research which has been directed towards this goal has been predominantly from a social psychology perspective and it may be that an explanation within the more rigorously defined cognitive science framework is possibly more suitable. A situation based mechanism driven by prior experience in the world along the lines of a script based theory (Minsky, 1975) could account for the direction and strength of implicit causality bias as a function of the interaction between assumed and contextually specific information. It is not clear whether the mechanisms responsible for producing the bias do so 'on the fly' or whether it is the case that a previously generated bias preference is accessed when the appropriate cues are made available. Previous experience with both a situation and the way in which a situation is described may have resulted in statistically generated biases.

Although it should be possible to link the mechanisms giving rise to the implicit causality effect and the processing consequences associated with those mechanisms, in general the two have been kept separate. The rest of this chapter focuses on the processing influence of implicit causality.

3.4 Evidence for the Influence of Implicit Causality Information in Language Production —

Garvey and Caramazza (1974) followed Chafe (1972) who proposed that verbs contain in their conceptual level of representation 'causes, accompaniments and results.' They suggest a property encoded at the lexical level attributing cause to one of the two mentioned noun phrases. In a production task subjects were required to complete sentence fragments of the sort :

(4) The prisoner confessed to the guard because he ...

As mentioned above, Garvey and Caramazza report that some verbs clearly impute cause to the character occupying the first noun phrase (NP1) while others impute cause to the character occupying the second noun phrase (NP2). Similar results were obtained by Au (1986) for both adults and preschoolers, by Hoffman and Tchir (1990) where subjects had to simply indicate which character they thought was responsible for the described event and by Grober, Beardsley and Caramazza (1978).

On the basis of their production data Garvey and Caramazza make a number of speculative suggestions. The first is that transforming the sentences to the passive voice seems to preserve the implicit causality bias, but only for NP2 biasing verbs. Furthermore they make the direct prediction that sentences with the subordinate clause forming a continuation **against** the direction of implied bias will result in reading times longer than for those whose subordinate clause continuation is consistent with the direction of bias. This could be due to a conflict between the implicit and explicit causes and may reflect processing difficulty associated with integrating unexpected information with the rest of the sentence. It could also be due to readers initially interpreting the pronoun as consistent with the verb bias and then having to reassign reference once the rest of the subordinate clause is read. We can think of verb bias as a constraint, weak enough to be overridden but not weak enough so as to not lead to some degree of processing difficulty when it is overridden.

The effect of passivisation was experimentally examined by Garvey, Caramazza and Yates (1976) in a series of production studies.

Following a passive transformation, verbs exhibiting a strong NP1 bias in the active voice demonstrated a weak NP2 bias in the passive. Garvey et al interpret this as reflecting a general preference towards pronominalisation of the grammatical subject of the sentence. This is consistent with both Subject Assignment Strategy and, as the pronoun in the subordinate clause occupies the grammatical subject role, Parallel Function Strategy.

Following a manipulation of polarity Garvey et al found a general tendency for subjects to produce continuations focusing on the first noun phrase when the verb was negated. As this first mentioned noun phrase was also the grammatical subject of the sentence this preference may again reflect a general tendency to refer to the preceding grammatical subject .

Garvey and Caramazza (1974) and Garvey, Caramazza and Yates (1976) have demonstrated that subjects have sufficient sensitivity to implicit causality information for it to exert an influence on language production. However, the question now becomes one of whether this generalises to language comprehension. It is possible that the strength of the influence of implicit causality information is artificially magnified by the nature of the task itself. Subjects are required to produce a continuation following a relatively impoverished stimulus. It is possible that this results in an amplification of the influence of information available in the experimental stimulus pertinent to the task of providing a sentence continuation. As it is certainly a very salient property associated with the verbs examined (Au, 1986), the relative contribution of implicit causality information may increase as a function of task demands. Balota (1997) proposes that the relative contributory weights of the processing pathways corresponding to processes operating over different types of information may be affected by the nature of the experimental task employed to examine a particular phenomenon. It may be that the processing pathways corresponding to those parts of the cognitive system sensitive to implicit causality information are given unrepresentative weightings in the case of the production studies carried out by Garvey and colleagues. The extra reflective

time available to subjects when engaged in written language production tasks may further increase these weights.

3.5 Evidence for the Influence of Implicit Causality Information in Language Comprehension --

Caramazza, Grober, Garvey and Yates (1977) examined the online influence of implicit causality in language comprehension obtaining the general pattern of results predicted by Garvey and Caramazza (1974). They employed a naming task where subjects had to name the referent of a pronoun. They report that naming was faster following a sentence where the subordinate clause continued with reference to the character consistent with the direction of the implicit causality bias (see examples (5) and (6) below). This was found even in conditions where gender information was sufficient to determine reference.

(5) The mother punished her daughter because she admitted her guilt. (consistent with bias)

(6) The mother punished her daughter because she discovered her guilt. (inconsistent with bias)

Caramazza et al admit that the causal bias should be treated as nothing more than a bias rather than as a binary feature. It's more accurate to consider a scale of biases from NP1 to NP2 with very strongly biasing verbs lying at either extreme of this fuzzy scale.

Garnham and Oakhill (1985) report a study indicating a processing influence of implicit causality information only when supported by gender information. From a methodological perspective there are other possible reasons why no clear effect was found. Only a small number of verbs were examined (4). As the experimental measure was reading time it is possible that the small variation in the verbs used in the experimental sentences (of which there were 32, each verb appearing on 8 occasions) lead to a set effect throughout the course of the experiment. Combined with noise resulting from varying levels of plausibility of the experimental materials, this may

have clouded the pattern of data. In light of these criticisms and by the very fact that such congruency effects have been consistently found by others (see above and below), Garnham and Oakhill's failure to locate an effect of implicit causality should not be overinterpreted.

Garnham, Oakhill and Cruttenden (1992) provide further support for the influence of implicit causality in comprehension. They report faster reading times for sentences in which the subordinate clause continues in a manner consistent with the direction of implicit causality bias.

They also report evidence suggesting that sensitivity to gender cue appears to be somewhat under the strategic control of the reader. In one experiment subjects had to read passages followed by a question which could only be answered correctly if the pronoun in the preceding target sentence had been successfully resolved. Under this condition subjects were faster at reading the subordinate clause when there was a gender cue present than when it was absent. In the case where the questions didn't focus on the pronoun in the preceding passage, there was no comparable effect of gender cue. Garnham et al take this as evidence that sensitivity to gender cue is to some degree under the control of the reader. It can be accounted for in the framework proposed by Balota (1997) emphasising the influence of the structure of the experimental task on processing. Over the course of the experiment, the need to answer questions requiring resolution of the pronoun may have affected the relative contributory weight of those aspects of processing associated with resolving the pronoun so amplifying the influence of factors relevant for that task. So, the explanation for the difference between the experiments containing different forms of questions may lie at the level of structure in the experimental environment which may then have lead to readers employing some sort of particular reading strategy.

3.5.1 Conditions Necessary for the Bias to be Manifested

Although the studies reported above have found strong evidence for a processing influence of implicit causality information in both language production and comprehension, it is unclear to what extent the

processing influence of causality biases generalises to sentences of alternative structures. Up to this point the effect has only been examined within the context of sentences of the form 'NP1 verb NP2 because Pronoun'. The connective 'because' explicitly focuses the reader's attention on the causal nature of the relationship between the two noun phrases. In the presence of connectives focusing on other dimensions of the underlying event, aspects of the relationship other than causal ones may be more salient. As implicit causality information can only provide some indication of the nature of the causal relationship of the event, it seems reasonable to suggest that when a dimension other than that of cause is focused upon, implicit causality information will be of very little help and so exert either a greatly attenuated processing influence or no processing influence at all.

Ehrlich (1980) examined the importance of the type of relationship between the two characters denoted by the noun phrases as instantiated by the connective in licensing the processing influence of the implicit causality bias. In a naming task, she examined the behaviour of the implicit causality bias under conditions where the connective was 'and', 'but' and 'because'. Under one of its simplest interpretations, the connective 'and' relates the events it conjoins in a simple serial temporal manner (i.e. 'and then...'). This interpretation places no special emphasis on a causal relationship between the mentioned participants (but see section 2.7.2 above). The connective 'but' is usually used to deny some presupposition or expectation that could logically be inferred given the information contained in the text up to that point. Again it places no special emphasis on the causal nature of the underlying event.

Along with type of connective, Ehrlich also manipulated gender. In the case of gender differentiation between the two participants, antecedent identification is possible without reference to any other factor. Subjects were engaged in a task where they were visually presented with a sentence after which they had to release a button and name the pronominal referent.

Ehrlich found a main effect of connective, with naming following conditions containing the connective 'because' significantly faster than following the comparable 'and' and 'but' conditions. Subjects also responded more quickly when a gender contrast was present. Ehrlich's implicit causality effect was only found for sentences containing the connective 'because' and didn't generalise to other types of connectives.

Ehrlich states that overall it seemed the case that a sentence possessing a main-subordinate clause structure (i.e. the 'because' conditions) was easier to process than one containing a main-main clause order (i.e. the 'and' and 'but' conditions). Whether this is due to a general increase in processing load required to integrate two main clauses, or to a general facilitation of processing when the integrating factor (in this case causality) has already been provided is unclear. The data certainly indicate that only when the causal relationship between the two clauses is overtly indicated does implicit causality influence pronoun resolution.

3.5.2 Conclusion

The preceding section provides a summary of those studies which have demonstrated an influence of implicit causality information on language comprehension. This has been found on naming time, reading time and on probe task recognition response times.

3.6 The Time Course of the Influence of Implicit Causality Information in Language Comprehension

In Chapter 2 we suggested a possible restricted account of pronoun resolution in which the initial stage of co-indexation operates only over a restricted pool of information available within the input. Implicit causality information is certainly available at some point within the system, witnessed by the effects reported above. What is less clear is at which point this information exerts an influence. Greene and McKoon (1995) suggest that implicit causality doesn't exert an influence before the point at which the pronoun is encountered. The nature of this temporal dimension has been

examined more rigorously by McDonald and MacWhinney (1995) and by Garnham, Traxler, Oakhill and Gernsbacher (1996). Both sets of studies employed the probe task (see section 2.4 above for an overview of the task) but produced data supporting conflicting positions. Roughly, McDonald and MacWhinney (1995) have found evidence for an effect of such information at an early point as soon as the pronoun is encountered, while Garnham, Traxler, Oakhill and Gernsbacher (1996) have found evidence only for an influence at a much later point.

McDonald and MacWhinney (1995) have produced data which suggests that implicit causality information is used as soon as the reader encounters a pronoun, as in example (6) below. Indeed, following the connective 'because' is the earliest point at which implicit causality information can potentially exert a beneficial processing influence. It is only when the connective is read does it become clear that the focus of the sentence is on the causal relationship between the two characters mentioned in the main clause. The referential ambiguity encountered upon reading the pronoun following 'because' can only immediately be resolved using implicit causality information (under the assumption that the continuation following the pronoun will be consistent with verb bias).

(7) John blamed Bill because he...

McDonald and MacWhinney performed three similar experiments. In each study, a sentence such as (8) was presented auditorily to subjects.

(8) Beth disappointed Pam bitterly because she was so hard hearted at the anniversary party.

A probe word was presented visually at various points relative to the auditory stream. The subjects' task was to carry out a recognition response to the probe. In McDonald and MacWhinney's first study, the probe point positions were 100 msec following the second name, immediately following the pronoun, 200 msec following the pronoun and at the end of the sentence. For NP1 biasing verbs, a general first

mentioned character advantage was found independent of probe point. Subjects were always faster responding to the probe word corresponding to the first mentioned character. For NP2 biasing verbs, the same first mention preference was found at probe points 1 and 3, but at points 2 and 4 (after the pronoun and at the end of the sentence) reaction times to the second mentioned character showed facilitation. This facilitation lead to second character probes being responded to with the same ease as first mentioned character probes. At no probe point was reaction time to a probe corresponding to the second mentioned character faster than reaction time to one corresponding to the first mentioned. McDonald and MacWhinney found no evidence that when the two characters were differentiated by gender this information was used immediately by the reader although overall the presence of a gender contrast lead to an increase in the accessibility of the pronominal referent.

In their second study, McDonald and MacWhinney demonstrated the necessity for the presence of a causal connective and the beginning of a subsequent clause focusing on the cause of the described event for the activation differential to be found. Using sentences of the sort (9) below, they found no evidence for an implicit causality effect at any probe point.

(9) Gary amazed Alan time after time at the juggling competition.

McDonald and MacWhinney's third experiment consisted of presentation of the main and subordinate clauses as separate sentences and without an explicit causal connective (example (10)).

(10) Gary amazed Alan time after time. He was so talented.

A similar pattern of data to that obtained in their first study was found, although the pattern across time of the influence of implicit causality information was markedly different. Without the explicit causal connective the reader has to generate an inference instantiating the causal relationship between the information contained in the two sentences. It takes longer to apprehend the causal nature of the relationship when that has to be generated as an

inference. As other cognitive processing is occurring to generate the inference, the role that implicit causality plays is reduced. The implicit causality effect was on going and didn't show the sharp effect at the end of the sentence as it did in their first study.

We shall say no more about McDonald and MacWhinney's second and third experiments as the tension in the literature is dependent on the data from their first study and the studies described in Garnham et al (1995). Evidence from McDonald and MacWhinney's first experiment seems to indicate strongly that implicit causality information exerts a processing influence upon readers' encountering of the pronoun.

In contrast, the studies described by Garnham et al provide no evidence for an early influence of implicit causality information. Their experimental method also consisted of a probe task. Subjects were presented with sentences visually, word by word. The probes were again presented visually and subjects were required to make a recognition response. Although this precise instantiation of the probe task differs in several ways from the technique used by McDonald and MacWhinney, we argue in Chapter 7 when we return to the tension between the positions of McDonald and MacWhinney and Garnham et al that these methodological differences alone cannot account for the contradiction in their findings.

Over their four studies, Garnham et al manipulated probe point position, delay of probe onset and gender contrast. Consistent with the data of McDonald and MacWhinney, Garnham et al also found a general first mention effect (Gernsbacher and Hargreaves, 1988). In their second and fourth experiments, the effect was magnified following a probe at the end of the sentence. In none of their studies did Garnham et al find evidence that implicit causality information was affecting the relative activation levels of the antecedent and nonantecedent at the point of the pronoun. Not only was this effect not found on the pronoun, but there was no indication that a referent/ nonreferent activation differential occurred at *any* of the probe test points. The only reliable and consistently found effect measuring antecedent accessibility was one of first mention. Only when a gender contrast was also present did Garnham et al find

evidence of an increase in the pronominal antecedent's activation over that of the non-antecedent.

The data of Garnham et al are supported by Greene and McKoon (1995) who also failed to find an early influence of implicit causality information on activation using the probe task. Greene and McKoon (1995) employed a probe task to measure differences in the accessibility of the two participants when contained in a sentence possessing an implicit causality verb.

For sentences containing NP2 biasing verbs Greene and McKoon (1995) found an increase in response time to the probe word corresponding to the second mentioned character (NP2) over response time to the probe corresponding to the first mentioned (NP1). For NP1 biasing verbs they found no difference between response times to the two probes words. This is consistent with the recency effect reported by Gernsbacher, Hargreaves and Beeman (1989) (see section 2.5 above) following a 0 second probe delay. Because of this recency effect, the activation associated with NP2 will be relatively high compared to activation associated with NP1. This difference will increase following an NP2 biasing verb and will be reduced following an NP1 biasing verb. This decrease in difference will arise because the activation associated with NP1 will increase following a sentence containing an NP1 biasing verb but the increase will not be of sufficient magnitude to surpass the activation of NP2 which has occurred as a result of a recency effect. Increasing the delay between sentence offset and probe is one way of getting round the consequences of the recency effect.

We focus on this tension between the positions of McDonald and MacWhinney and Garnham et al below in Chapter 6. Our Experiment 8 directly addresses the contradiction in their findings.

Conclusion

The studies summarised above contain a number of methodological problems. The first concerns the experimental materials, while the second concerns the techniques adopted to measure the processing of

those materials. With respect to the experimental materials, no study has controlled either the plausibility or length of the sentences. These are not necessarily orthogonal dimensions. For sentences containing continuations going against the bias of the verb, in order to produce a natural sounding sentence it is usually necessary to add a few extra words which make the link between the subordinate and main clauses more obvious. When a continuation is consistent with the direction of verb bias, because of the strength of the bias much can be left implied. This results in sentences containing continuations inconsistent with the direction of verb bias being longer than their consistent counterparts.

Where there is greater homogeneity between the two in terms of length, this has the direct result that there is a reduction in the likelihood that the plausibility between the two is comparable. Both plausibility and length differences are certainly capable of producing the pattern of data traditionally found in studies examining the processing nature of implicit causality. A reduction in the acceptability of a sentence on purely plausibility grounds could lead to a general comprehension problem mirrored in both reading times and time to respond to a probe word. Although admittedly this should reduce the response times to probe words corresponding to either character but preserve the relative difference between the two, the time and resources required to comprehend a less plausible sentence may interact differently with different components of that sentence (i.e. first and second mentioned characters).

A more serious question concerns the extent to which the response measure can be interpreted as unambiguously revealing something of the underlying anaphoric processing. It is assumed that lexical access and syntactic processes will be common to both consistent and inconsistent sentences. The sentences should differ in terms of processing related to resolving anaphoric reference. As we proposed at the end of Chapter 2, anaphoric processing can be defined as those aspects of processing associated with locating the appropriate antecedent combined with those aspects associated with integration. Given the measures so far adopted, it is not clear that they are

capable of measuring anything other than the sum of these aspects of processing.

Summary --

Following Chapter 2, in which we asked ourselves at what point non-structural factors influence anaphoric processing, in this Chapter we have focused on one particular type of non-structural influence : Implicit Causality.

This chapter has provided a summary of both the literature focusing on the processing nature of implicit causality information and the literature considering what generative mechanism may lie at the heart of the phenomenon. We started by demonstrating evidence for an influence of implicit causality in language production before examining evidence for a comparable influence within the language comprehension system. Following discussion of the conditions necessary for the bias to be realised, we spent some time focusing on evidence pertinent to the temporal aspect of the processing influence of implicit causality. We shall return to the conflict between McDonald and MacWhinney and Garnham et al in Chapter 6 where we shall attempt to reconcile their contradictory positions. As summarised above, with the exception of McDonald and MacWhinney (1995) the empirical evidence suggests that implicit causality exerts the bulk of its influence at the integrative stage of processing.

CHAPTER 4

4.1 Introduction and Rationale

In the preceding chapters we focused on several theoretical issues. These included the way in which different sorts of information influence anaphoric processing and the time course associated with that processing. We can now extend the issues summarised in Chapters 2 and 3 with respect to the influence of implicit causality information. Central to the issue of whether aspects of the processes associated with resolving anaphoric reference should be construed as operating in a restricted or nonrestricted manner is the question of whether implicit causality information exerts an immediate or delayed influence on processing. Is initial processing guided by low level factors with implicit causality information exerting an influence only at a later stage? What studies reported in the literature that have attempted to address this question have produced conflicting data (McDonald and MacWhinney, 1995; Garnham, Traxler, Oakhill and Gernsbacher, 1996). In this chapter we report four experiments. The first provides us with off-line measures of the role played by implicit causality while the others examine the online influence of implicit causality.

4.2 Problems with previous Empirical Work on Implicit Causality

There are several problems with previous experimental examinations of the influence of implicit causality on reading. Recall that the implicit causality congruency effect is a relative effect. It is measured by comparing reading times to sentences containing subordinate clause continuations consistent with the bias of the verb in the main clause with those containing continuations inconsistent with the bias of the verb. It is possible that continuations going against the bias of the verb are simply less plausible than continuations consistent with the bias and it may be this that is leading to a processing difficulty. We want to be sure that it is only the mismatch between implicit and explicit causes accounting for the increased reading time and not a general processing difficulty arising from reading implausible

sentences in themselves. Both Examples (1) and (2) contain a subordinate clause going against the bias of the verb. The implicit causes and explicit causes contradict each other.

(1) John blamed Bill because John was in a bad mood.

(2) John blamed Bill because John was tall.

In Example (1), the fact that John was in a bad mood seems to be a good reason for him blaming Bill. It is possible that John is being generally grumpy towards everybody. In Example (2) however, the fact that John is tall does not immediately provide a sensible reason for him blaming Bill. Either the reader has to generate a number of extra inferences to link the main and subordinate clauses, or has to wait until further information is provided to make sense of the sentence. Although the implicit and explicit causes contradict each other in (1) and (2), while (1) is a relatively plausible sentence, (2) is not. We want to be sure that the implicit causality effect reported in the literature is more than a plausibility effect. For the experiments reported in this thesis, the plausibility of Example (1) should be equivalent to the case where the subordinate clause is consistent with the verb bias, as in Example (3).

(3) John blamed Bill because Bill broke the window.

The circumstances of Bill breaking the window and John being in a bad mood should both be plausible reasons for leading to the event of John blaming Bill.

Secondly, it may be the case that sentences containing subordinate clauses going against the bias of the verb simply contain more words. Much information is implied by the verb. When the causality bias is violated this implied information becomes uninformative and it may be necessary to explicitly introduce new information in the utterance in order to produce a reasonable sounding sentence. This will result in sentences with continuations going against the bias of the verb containing more words thus leading to longer reading times. In our experiments we control the length of our experimental materials.

The implicit causality bias should be treated as no more than a bias. In previous examinations of the influence of implicit causality on reading verbs falling into one of the two NP1 and NP2 biasing groups have generally been considered as more or less equal within those groups. It may be the case however that there is great variation in strength of verb bias with those verbs exhibiting a strong bias contributing disproportionately to the effects reported. No strength of bias controls have been employed by those examining the processing influence of implicit causality.

Off-line Experiments

In order to satisfy ourselves that the effects previously reported in the literature could not be explained either in terms of plausibility differences between experimental conditions or variation between verbs nominally of the same type of bias we first performed 2 off-line experiments.

4.2.1 Experiment 1a

The goal of our first off-line study was to select implicit causality verbs possessing strong biases.

Method

Subjects

24 English speaking subjects participated.

Stimuli

Experimental booklets were constructed containing 110 sentence fragments. We selected 50 verbs from the literature and embedded each verb in a sentence of the form :

John VERBED Bill because ...

Names were varied. Each booklet contained 60 filler materials of a similar structure to the above fragment. Each implicit causality verb appeared only once.

Procedure

One random order was used for half the booklets while the reverse of this was order used for the other half. Subjects were simply instructed to provide a sensible continuation to the sentence fragments.

4.2.2 Results

The endings for each experimental sentence were scored on the basis of whether the first word produced following 'because' was a pronoun or repeat name referring to the first or second noun phrase. Continuations for which it was unclear as to which character the ending referred were scored as 'Indeterminate' and sentences containing continuations not belonging to any other scoring category, such as 'John blamed Bill because of the broken window.', were scored as 'other'. Due to a typo in the sentence containing one of the verbs 'believed', that verb's data is missing. For each verb, the total number of continuations which fell into each category of continuation are reported below :

Table 1 - Experiment 1a continuation results

	NP1 pronoun	NP2 pronoun	NP1 name	NP2 name	Indeterminate	Other
lied to	23	0	0	0	0	1
accused	12	6	0	4	1	1
fascinated	21	0	1	0	0	2
infuriated	22	0	0	0	0	2
questioned	4	1	0	0	1	18
agitated	17	0	1	1	2	3
disappointed	21	0	0	0	0	3
confessed to	22	0	0	0	1	1
apologised to	21	1	0	0	0	2
angered	18	1	0	1	1	3
inspired	16	0	2	1	0	5
flattered	12	8	0	0	2	2
confided in	13	6	2	1	1	1
troubled	19	1	3	0	1	0
dumbfounded	17	1	0	0	2	4

telephoned	20	1	0	0	0	3
bored	17	1	2	0	0	3
amused	19	1	0	0	0	4
confused	17	0	0	1	3	3
approached	9	5	0	1	5	4
concerned	19	0	0	1	3	1
amazed	18	1	0	0	0	5
called	16	0	0	3	1	4
aggravated	16	0	0	2	0	6
annoyed	19	0	1	0	1	3
adored	0	17	0	0	3	4
honoured	0	19	0	2	1	2
despised	1	16	0	4	1	2
liked	1	18	0	1	2	2
criticised	4	14	0	1	1	4
reassured	5	10	0	7	2	0
complimented	4	18	0	0	0	2
derided	6	6	0	2	3	7
valued	1	17	0	2	1	3
ridiculed	4	10	0	3	1	6
appreciated	1	16	0	2	0	5
resented	0	19	0	3	1	1
congratulated	1	21	0	0	1	1
deplored	1	17	0	0	1	5
thanked	2	17	0	2	0	3
detested	1	19	0	3	1	0
dreaded	3	15	0	0	1	5
loathed	2	18	0	2	0	2
praised	1	17	0	0	0	6
consoled	1	6	1	1	0	15
noticed	0	17	0	0	0	7
scolded	1	17	1	1	0	4
punished	1	14	1	2	2	4
believed	x	x	x	x	x	x
admired	0	18	0	1	0	5

The data reported in the above table provides us with an indication of the relative strengths of the implicit causality bias associated with each verb. The maximum number of continuations of each category for each verb is 24. The bias strength for the verbs we examined is reported below along with an analysis of the plausibility rating.

4.2.3 Experiment 1b

The goal of the second off-line experiment was to ensure that the materials in each of the experimental conditions we intended to examine in our on-line task were of equivalent plausibility. We wanted to be sure that the version containing the ending consistent

with the verb bias and the version containing the ending inconsistent with the bias matched on a plausibility rating. We did this to exclude an explanation for the difference in reading time between the consistent and inconsistent conditions in terms of a simple difference in their levels of plausibility.

We used the same set of verbs as was used in our first off-line task. For each verb we constructed two endings, one continuing with reference to the first Noun Phrase and one continuing with reference to the second Noun Phrase (see examples (4) and (5) below).

- (4) John blamed Bill because Bill broke the window.
- (5) John blamed Bill because John was in a bad mood.

The verb 'blame' is NP2 biasing so Example (4) is consistent with the bias as the reference is to the second Noun Phrase. Example (5) contains a continuation inconsistent with the bias as reference is to the first Noun Phrase. We want the plausibility of each of these sentences to be equivalent. The continuation should provide unambiguous evidence as to which character is the appropriate pronominal referent when the repeat name anaphor is replaced with a pronoun. Example (6) and (7) below should be implausible interpretations.

- (6) John blamed Bill because John broke the window.
- (7) John blamed Bill because Bill was in a bad mood.

To measure the plausibility of these four possible interpretations (examples 4-7) we constructed two endings which for each verb produced a set of 4 items. We equated the average length for each of these 4 conditions.

To recap, we want Examples (4) and (5) to be rated as highly plausible while we want Examples (6) and (7) to be rated as highly implausible.

Method

Subjects

24 English speaking subjects participated.

Stimuli

Experimental booklets were constructed containing 160 sentences. We examined the 40 most strongly biasing verbs. 20 were NP1 biasing and 20 NP2. Each subject saw 4 instances of each verb, once in each of the 4 conditions. Randomisation of presentation order occurred in 4 blocks to prevent any sentence pair containing the same verb being in close proximity. Half the subjects saw the sentences in one random order, for the other half this order was reversed. Subjects were simply asked to rate the plausibility for each sentence. The instructions given at the start of the booklet are repeated below :

For the sentences below, rate how plausible you think they each are as a description of an event.

For example :

Jon and Mary got married because they were in love.

If you think this is plausible you might do :

|0-----1-----2-----3-----4-----5-----6-----7|

Implausible

Plausible

Hilary lent Margaret some money because Margaret was short of cash.

Jon and Mary got married because crocodiles have sharp teeth.

10-----1-----2-----3-----4-----5-----6-----7|

Implausible

Plausible

Try and use the entire scale.

You can take a break as often as you like but complete all the sentences.

4.2.4 Results

We calculated the average plausibility rating for the experimental materials. From these we selected 24 verbs which exhibited the plausibility pattern we wanted and also possessed strong implicit causality biases. Two criteria were applied to select our set of verbs. The first was that the number of continuations for a verb referring to one Noun Phrase had to equal 16 or more for that verb to be classified as strongly biasing. The other criterion was that the plausibility ratings were to be statistically equal overall for the versions corresponding to plausible continuations for the first and second mentioned character conditions and for the versions corresponding to implausible continuations for the first and second mentioned character conditions. For NP1 verbs the average of the NP1 bias is 19.5, the average of the NP2 bias 0.417. For NP2 verbs the average of the NP1 bias is 1.0 and the average of the NP2 bias is 17.333. These biases are out of a total of 24. The scoring category Plaus1 corresponds to what we want to be a plausible interpretation of the version of the sentence containing reference to the first mentioned character (example 5 above). Category Implaus1 corresponds to what we want to be an implausible interpretation of the version containing reference to the first mentioned character (example 6 above). The scoring category Plaus2 corresponds to what we want to be a plausible interpretation of the version of the sentence containing reference to the second mentioned character (example 4 above). Category Implaus2 corresponds to what we want

to be an implausible interpretation of the version containing reference to the second mentioned character (example 7 above).

We performed a separate ANOVA on the Plausibility ratings for the NP1 and the NP2 sets of verbs. For the NP1 verbs we found a main effect of plausibility ($F(1,11)=536$, $p<0.0001$). There was no interaction ($F<1$) between plausibility and referent. In other words we can treat the plausibility ratings for Plaus1 and Plaus2 as statistically equivalent and likewise for Implaus1 and Implaus2.

For the NP2 verbs we found a main effect of plausibility ($F(1,11)=870$ $p<0.0001$). There was no interaction ($F<1$) between plausibility and referent. Again we can treat the plausibility ratings for Plaus1 and Plaus2 as statistically equivalent and likewise for Implaus1 and Implaus2.

Those verbs are reported in the table below.

Table 2 - Experiment 1b plausibility results combined with Experiment 1a continuations results

NP1 Verbs

	NP1 continuation	NP2 continuation	Plaus1	Implaus2	Plaus2	Implaus1
fascinated	21.000	0.000	6.417	1.292	5.667	3.000
infuriated	22.000	0.000	6.583	1.625	5.375	2.625
disappointed	21.000	0.000	6.458	1.208	6.292	3.083
confessed to	22.000	0.000	6.542	1.167	6.304	1.042
apologised to	21.000	1.000	6.565	0.667	6.500	2.042
inspired	16.000	0.000	6.042	2.333	5.333	1.750
troubled	19.000	1.000	6.500	1.708	5.875	2.792
telephoned	20.000	1.000	6.542	2.522	6.292	0.500
amused	19.000	1.000	6.250	1.083	5.833	2.583
concerned	19.000	0.000	6.125	1.391	5.292	2.125
amazed	18.000	1.000	6.292	1.000	6.125	2.458
called	16.000	0.000	8.000	2.542	6.792	0.957
Mean	19.500	0.417	6.370	1.545	5.973	2.080

NP2 Verbs

	NP1 continuation	NP2 continuation	Plaus1	Implaus2	Plaus2	Implaus1
admired	0.000	18.000	6.125	1.583	6.333	1.913
appreciated	1.000	16.000	6.125	1.792	6.583	2.708
praised	1.000	17.000	6.000	1.875	6.750	2.261
despised	1.000	16.000	6.250	1.917	6.583	1.250
liked	1.000	18.000	6.000	1.958	6.417	1.208
noticed	0.000	17.000	6.042	2.000	6.625	1.375
congratulated	1.000	21.000	5.875	2.000	6.833	1.083
loathed	2.000	18.000	5.708	2.375	6.125	2.958
punished	1.000	14.000	5.583	2.458	5.958	1.583
scolded	1.000	17.000	5.958	2.625	6.500	1.250
detested	1.000	19.000	6.042	2.958	6.375	2.292
thanked	2.000	17.000	6.583	3.417	6.583	1.318
Mean	1	17.333	6.024	2.247	6.472	1.767

Recall the goal of the above off-line tasks was twofold :

1. To ensure that our verbs exhibited strong implicit causality biases.
2. To ensure that the endings were of equivalent plausibility, regardless of whether they were consistent with the implicit causality verb bias.

The above set of 24 experimental materials forms the basis for all the studies reported in this thesis examining the influence of implicit causality.

Experiments examining On-line processing

With the exception of one experiment (Experiment 7 in Chapter 5) all the studies reported in this thesis examining the on-line influence on processing of implicit causality information employ a self-paced reading methodology. This consists of words being presented on a computer screen with degree of the presentation rate being under the control of the subject. Subjects press a button once they have finished reading what is presented on the screen. The response latency (timed from onset of presentation of the experimental material to the point at which the Subject responds by pressing the

button) is assumed to be an index of processing associated with the stimulus. For more details, see below.

4.3 Experiment 2 --

4.3.1 Introduction and Rationale

The goal of our first study is to provide a basic replication of the finding reported in the literature and first predicted by Garvey and Caramazza (1974) that sentences containing subordinate clause continuations inconsistent with the bias of the verb in the main clause take longer to read than those containing subordinate clause continuations consistent with the bias of the verb. As outlined in the introductory section to this chapter, because pretests of the sort we carried out above had not previously been employed, we wanted to demonstrate that non-homogeneity of the verbs examined or a simple plausibility or length difference in the experimental materials could not account for what is considered to be the implicit causality congruency effect.

In our first on-line experiment we manipulated 3 factors. These were anaphor, verb bias and referent. Anaphor could either be pronoun or repeat name. Verb bias, a between experimental items manipulation, could either be NP1 or NP2 biasing. Referent could either be the first or second mentioned character. The interaction between verb bias and referent corresponds to the implicit causality congruency effect. In other words, an NP1 biasing verb followed by a continuation referring to the first mentioned character is congruent, as is an NP2 biasing verb followed by a continuation referring to the second mentioned character. An NP1 biasing verb followed by a continuation referring to the second mentioned character is incongruent, as is an NP2 biasing verb followed by a continuation referring to the first mentioned character. In this experiment the sentences were presented in a single display.

Method

Subjects

Thirty two English speaking subjects participated.

Stimuli

We manipulated anaphor, verb bias and referent. Anaphor could either be pronoun or repeat name. Verb bias could either be NP1 or NP2 biasing. Referent could either be the first or second mentioned character.

There were 48 sets of experimental materials (see appendix 1 for full set). Below is an example of the 4 experimental materials for the NP1 biasing verb 'fascinate'.

Name / NP1 verb / Referent character1

Barry fascinated Derek because Barry performed magic tricks.

Name / NP1 verb / Referent character2

Barry fascinated Derek because Derek was easily entertained.

Pronoun/ NP1 verb / Referent character1

Barry fascinated Derek because he performed magic tricks.

Pronoun / NP1 verb / Referent character2

Barry fascinated Derek because he was easily entertained.

There were also 96 filler items. These were of similar type to the experimental materials. A full list can be found in Appendix 8. The 24 verbs used to construct the experimental materials were selected from the set of verbs examined in the off-line studies (see list above). Each subject saw each verb twice. On the first occasion with the continuation referring to one character, on the second with the continuation referring to the other. The experiment was divided into two halves with a break halfway through. The first occurrence of

each verb was in the first half of the study and the second occurrence in the second half.

Procedure

The experiment was run on an Apple Macintosh computer using the Psyscope experimental software (Cohen, MacWhinney, Flatt and Provost, 1993). A button box was connected to the computer which recorded subjects' responses with millisecond accuracy.

Each subject participated in 10 practice trials similar in structure to the experimental items at the start of the experiment. The experiment lasted roughly 35 minutes. Before the experiment Subjects were provided with both verbal instructions and written instructions repeated below :

In this experiment you are required to read a number of sentences on the computer screen.

At the start of each trial the phrase '++ Press button box for next sentence ++' will appear in the middle of the screen. Press any key on the button box and a star will appear - look at this star and then press a button on the button box - a sentence will appear. Having read this, press a key again and you will be presented with a question.

With each question there are two alternative answers - one will be presented on the left hand side of the screen, the other on the right. If you think the left hand answer is correct, press the red (leftmost) button on the button box but if you think that the right hand answer is correct, then press the green (rightmost) button on the button box.

After you have answered the question the prompt '++ Press button box for next sentence ++' will appear again. There is a break halfway through the experiment. If you

need to pause at any other time please try and do so only when the prompt '++ Press button box for next sentence ++' is on the screen.

If you have any questions ask the experimenter now.

A fixation point appeared on the left hand side of the screen. Subjects pressed a button on the button box and the fixation point was followed by presentation of a sentence. Once subjects read the sentence they pressed the button again. A question then appeared. An example of the type of question is 'Who broke the window ?' with the alternative answers 'John' and 'Bill' presented below the question. One name was presented on the left hand side of the screen while the other was presented on the right. The left hand button on the button box corresponded to the name on the left hand side of the screen while the right hand button corresponded to the name on the right hand side of the screen. Subjects had to press either the left or right hand button corresponding to whichever answer they thought was correct. Subjects received no feedback. Following their response to each question they were prompted to press a button on the button box for the next trial. The fixation point then reappeared and the above procedure repeated. A break of a minimum duration of 30 seconds occurred halfway through each experiment.

4.3.2 Results

We report analyses for the experimental sentences themselves and also for the questions following those sentences.

Outlier Replacement

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the sentence data we excluded response times that were below 500 msec or above 25 seconds. This accounted for 1.3% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.5% of the data was replaced in this way. For the question response data we excluded response times

that were below 500 msec or above 25 seconds. This accounted for 1.3% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.9% of the data was replaced in this way.

Sentence Analysis

For the experimental sentences, means were calculated for each subject. The means corresponding to each condition are reported in Table 3.


Table 3 - Reading times, subject means for each condition with all times in msec.

Verb / Referent	Name	Pronoun
NP1 / Char1	4717	4753
NP1 / Char2	4732	5604
NP2 / Char1	4670	4851
NP2 / Char2	3948	4318

We performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char 2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias with sentences containing NP2 biasing verb read more quickly than those containing NP1 biasing verbs ($F(1,31)=19.01$, $p<0.0001$; $F(1,46)=17.68$, $p<0.0001$). We also found a main effect of anaphor with sentences containing a repeat name anaphor read more quickly than those containing a pronominal anaphor ($F(1,31)=10.59$, $p<0.005$; $F(1,46)=6.07$, $p<0.05$). We found an interaction between verb bias and referent ($F(1,31)=51.05$, $p<0.0001$; $F(1,46)=27.75$, $p<0.0001$) (see Figure 2). This corresponds to the implicit causality congruency effect. We also found a marginal interaction between anaphor type and referent significant by subjects only ($F(1,31)=5.51$, $p<0.05$, $F(1,46)=2.56$, $p<0.12$) (see Figure 1).

We first explored the nature of the interaction between anaphor type and referent.


Figure 1 - anaphor type x referent



In order to interpret this interaction we performed means comparisons on each possible pair of conditions. The points corresponding to a repeat name anaphor referring to the first mentioned character and a pronoun referring to the first mentioned character are statistically equivalent (both $F_s < 1$). All other conditions differ ($F_1: p < 0.05$) but only by subjects ($F_2s < 2$) with the exception of a name referring to the second and a pronoun referring to the first mentioned character ($F_2: p < 0.05$). The difference between name referring to second mentioned character and pronoun referring to second mentioned character is marginal by subjects ($F_1: p < 0.1$, $F_2: p < 0.01$). This means that sentences containing pronominal or repeat name reference to the first mentioned character are read with equivalent ease. There is a suggestion that sentences containing repeat name reference to the second mentioned character might be read quickest although this isn't clearly supported by both subjects and items analyses.

We now turn to the verb bias x referent interaction corresponding to the implicit causality congruency effect.

Figure 2 - verb bias x referent




We further explored the implicit causality congruency effect by examining the Name and Pronoun conditions separately.

Name Conditions

Looking at the Name conditions first, we found a main effect of verb bias ($F(1,31)=12.87$, $p<0.005$; $F(1,46)=6.10$, $p<0.05$) and a main effect of referent that was marginal by items ($F(1,31)=9.13$, $p<0.005$; $F(1,46)=2.98$, $p<0.1$). As reported above, sentences containing NP2 biasing verbs were read more quickly than those containing NP1 biasing verbs. The interaction between verb bias and referent, i.e. the implicit causality congruency effect was significant by subjects although marginal by items ($F(1,31)=8.89$, $p<0.01$; $F(1,46)=3.23$, $p<0.1$). Figure 3 below depicts this interaction. In order to interpret this interaction we performed means comparisons on each possible pair of conditions.

Figure 3 - verb bias x referent (Name conditions)




We performed 2-way ANOVAs for NP1 and NP2 biasing verbs separately. We also performed 2-way ANOVAs for Char1 and Char2 referents separately. For NP1 verbs we found that Char1 and Char2 referents are equivalent (both $F_s < 1$). For NP2 verbs we found that Char1 and Char2 referents differ ($F_1 = 21.09$, $p < 0.0001$; $F_2(1,46) = 7.54$, $p < 0.05$). For Char1 referents we found that NP1 and NP2 verbs are equivalent (both $F_s < 1$). For Char2 referents we found that NP1 and NP2 verbs differ ($F_1(1,31) = 23.99$, $p < 0.0001$; $F_2(1,46) = 8.40$, $p < 0.01$). This means that for repeat name anaphors, reading times to a sentence containing reference to the first mentioned character is equivalent regardless of whether that reference is supported by verb bias. Reading a sentence containing reference to the second mentioned character is fast when supported by verb bias compared to when there is no verb bias support. The condition corresponding to reading a sentence containing reference to the second mentioned character is fast compared to the other three conditions.

Pronoun Conditions

For the Pronoun conditions, we found a main effect of verb bias ($F_1(1,31) = 10.59$, $p < 0.005$; $F_2(1,46) = 7.96$, $p < 0.01$). The interaction between verb bias and referent, i.e. the implicit causality congruency

effect was significant ($F(1,31)=25.99$, $p<0.001$; $F(1,46)=16.08$, $p<0.0005$). Figure 4 below depicts this interaction.



We performed 2-way ANOVAs for NP1 and NP2 biasing verbs separately. For NP1 verbs we found that Char1 and Char2 referents differ ($F(1,31)=12.33$, $p<0.005$; $F(1,46)=9.92$, $p<0.005$). For NP2 verbs we found that Char1 and Char2 referents differ ($F(1,31)=7.68$, $p<0.01$; $F(1,46)=6.16$, $p<0.05$). We also performed 2-way ANOVAs for Char1 and Char2 referents separately. For Char1 referents we found that NP1 and NP2 verbs are equivalent (both $F_s<1$). For Char2 referents we found that NP1 and NP2 verbs differ ($F(1,31)=23.15$, $p<0.0001$; $F(1,46)=21.15$, $p<0.0001$). This means that for pronouns, reading times to a sentence containing reference to the first mentioned character are equivalent regardless of whether that reference is supported by verb bias. Reading a sentence containing reference to the second mentioned character is fast when supported by verb bias compared to when there is no verb bias support. The condition corresponding to reading a sentence containing reference to the second mentioned character is fast compared to the other three conditions.

Question Response Time Analysis

For the question response time data we again performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 referent vs Char2 referent) ANOVAs for both subjects and items as random factors. This indicated a persistence of the anaphor main effect ($F(1,31)=15.21$, $p<0.0005$; $F(1,46)=18.53$, $p<0.0001$) and marginal verb bias main effect ($F(1,31)=5.55$, $p<0.05$; $F(1,46)=3.25$, $p<0.1$). We found an interaction between verb bias and referent ($F(1,31)=19.64$, $p<0.0001$; $F(1,46)=15.13$, $p<0.0001$), i.e. an implicit causality congruency effect. We also found a 3-way interaction of anaphor type x referent x verb bias ($F(1,31)=14.17$, $p<0.001$; $F(1,46)=5.30$, $p<0.05$). This 3-way interaction corresponds to the finding of a congruency effect on the Pronoun conditions ($F(1,31)=28.64$, $p<0.0001$; $F(1,46)=18.80$, $p<0.0001$) but not in the Name conditions (both $F_s<2$).

Table 4 - Question response times, subject means for each condition with all times in msec. Numbers in parentheses correspond to response accuracy expressed as percentage correct.

Verb / Referent	Name	Pronoun
NP1 / Char1	1673 (98.9%)	1902 (87.4%)
NP1 / Char2	1775 (93.6%)	2317 (74.3%)
NP2 / Char1	1732 (96.3%)	2066 (91.1%)
NP2 / Char2	1677 (98.4%)	1748 (97.4%)

Question Response Accuracy Analysis

In order to determine whether there was some form of a speed/accuracy trade off for response times to the questions, we examined the proportion of correct and incorrect responses for each experimental condition.

We performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 referent vs Char2 referent) ANOVAs on the proportion of correct responses for both subjects and items as random factors revealed main effects of anaphor type ($F(1,31)=69.40$, $p<0.0001$; $F(1,46)=27.12$, $p<0.0001$) and verb bias ($F(1,31)=26.30$, $p<0.0001$; $F(1,46)=26.82$, $p<0.0001$). Responses were more accurate

following repeat name anaphors and following NP2 biasing verbs. We also found an interaction between verb bias and referent ($F(1,31)=42.62$, $p<0.0001$; $F(1,46)=17.76$, $p<0.0001$), i.e. an implicit causality congruency effect. This is the same pattern of data as revealed in response times to both the sentence and the following question. An interaction between anaphor type and verb bias was also found ($F(1,31)=20.50$, $p<0.0001$; $F(1,46)=20.86$, $p<0.0001$). The 3-way interaction of anaphor type x verb bias x referent was significant but only by subjects ($F(1,31)=6.07$, $p<0.05$; $F(1,46)=2.37$, $p<0.14$).

4.3.3 Discussion

The main effect of verb bias, that sentences containing NP2 biasing verbs were read more quickly than those containing NP1 biasing verbs, may be the result of differences in the thematic structures of the verbs. NP2 biasing verbs were predominantly action verbs while NP1 verbs were predominantly state verbs. It may simply be that forming a representation of an action is easier (and quicker) than forming a representation of an internal state. The effect was found consistently in other experiments reported below. However as we did not control for frequency or length of the verbs, we cannot be sure that these factors did not contribute to the effect.

Regardless of what is responsible of the main effect of verb bias, it demonstrates that the verb is certainly being processed as the first sentence is being read and it's not simply the case in our previous studies that we don't find early effects of implicit causality because the verb hasn't been properly read before the reader has finished reading the fragment in which it is contained.

The main effect of anaphor type, that sentences containing repeat name anaphors are read more quickly than those containing pronominal anaphora, may result from repeat name anaphors identifying their appropriate antecedents more quickly. A level of processing no deeper than a form match is all that is required in the case of repeat name anaphors. The antecedent of a pronoun cannot be identified on the basis of their visual similarity to the pronoun.

This main effect of anaphor is consistently found in the experiments reported in this thesis. The finding is consistent with the proposal we outlined earlier with respect to a two-stage model of anaphor resolution. When a repeat name anaphor is encountered, co-indexation between anaphor and antecedent can take place immediately. However for pronouns this co-indexation requires information following the pronoun. It may occur immediately but further disambiguating information is required to avoid misassignment. We return to evidence suggesting that misassignment does not occur in Experiment 6 in Chapter 5.

The main effects of verb bias and anaphor type persisted in both the question response time data and also in question response accuracy.

The interaction between anaphor type and referent (Figure 1) is the result of a repeat name penalty. The interaction is driven by the condition where reference to the second mentioned character is achieved through repetition of that character's name. This condition is read the fastest. There is a penalty associated with referring to the first mentioned character through the use of a repeat name as evidenced by our main effect of referent found for the Name conditions. There is no penalty associated with referring to the second mentioned character through the use of a repeat name.

We found an implicit causality congruency effect reflected in response times both to the sentences containing the implicit causality verbs but also to the following questions and in the level of accuracy in responding to those questions. The Name conditions data contains both the pattern of the implicit causality congruency effect, corresponding to a verb bias x referent interaction, and the repeat name penalty associated with reference to the first mentioned character, as indicated by the main effect of referent. The interaction in the Name conditions is a blend of these two effects.

Looking at Figure 2, it can be seen that reference to the first mentioned character is always easy, reference to the second is only easy when supported by verb bias. Although there was no evidence from the pattern of reading time data to the sentence that this

implicit causality congruency effect interacted with any other factor, there was evidence from the question response and question accuracy data that this did interact with anaphor type. The implicit causality congruency effect was stronger in the Pronoun than in the Name conditions. When reference is made through the use of a proper name, subjects are not required to integrate the information contained within the main and subordinate clauses to identify the referent. In the Pronoun conditions, implicit causality becomes an important indicator of reference and the integration of the information in the main and subordinate clauses is necessary in order for the pronoun's antecedent to be unambiguously identified.

With respect to the pattern of data associated with the question response times and question response accuracy, it is not the case that this is due to a speed-accuracy trade off. Response times for the condition corresponding to reference to the second mentioned character in the context of an NP1 biasing verb were slow and response accuracy poor.

In summary, our first examination of the on-line influence of implicit causality information has replicated the implicit causality congruency effect and has also demonstrated a repeat name penalty associated with reference to the first mentioned character through repetition of that character's name. Our pretesting and experimental controls allow to discount alternative explanations due to plausibility or length differences and non-homogeneity of verbs.

4.4 Experiments 3 and 4

4.4.1 Introduction and Rationale

Our first experiment replicated the basic implicit causality congruency finding. Our question of interest for the purpose of this thesis is the nature of the influence of implicit causality on anaphor resolution, i.e. the point in processing at which implicit causality exerts an influence. Also of interest is how resolution of different types of anaphors occurs. In the following experiments we altered the characteristics of the self-paced reading paradigm employed in Experiment 2 above in

order to provide us with an insight into how different information is used by those processes associated with anaphoric reference resolution.

The point at which the reader encounters an anaphoric pronoun is the earliest point at which we would expect implicit causality to influence processing. Consider Example (8) :

(8) John fascinated Bill because he ...

Recall the verb 'fascinated' biases towards interpreting the pronoun as coreferential with the first mentioned character. When the pronoun is encountered implicit causality information is the only non-structural cue available to the system which may be informative as to which antecedent is appropriate. In our first experiment examining the on-line nature of the influence of implicit causality we showed that the reader is sensitive to implicit causality information. Recall the following sentence (example (9)) took longer to read than its counterpart (example (10)) where example (10) contains an ending consistent with the verb bias (NP2 bias).

(9) Jack liked Tony because Tony was full of incredibly helpful advice.

(10) Jack liked Tony because Jack was made to feel quite at home.

The continuation goes against the bias of the verb. Indication that this is likely to be the case occurs when the reader encounters the repeat name 'Jack'. If implicit causality information is used by the reader early, we would expect a processing difficulty to occur whenever this repeat name is encountered. In other words we want to adopt a measure which will shed some light on processing at the point at which the anaphor is encountered. In Experiment 2 we demonstrated a repeat name penalty associated with reference to the first mentioned character using a repetition of that character's name. If implicit causality is used at the point at which the anaphor is encountered we expect there to be an interaction between verb bias and referent on the anaphor itself. Our whole sentence reading time measure in Experiment 2 would not have been sensitive enough to

pick up this effect. In Experiments 3 and 4 we present the sentences in two halves. The first half includes the anaphoric expression.

On the strength of the theoretical positions outlined in Chapter 2 above, what empirical predictions can we make about processing the anaphor ? Centering theory predicts a repeat name penalty associated with reference to the most focused entity when reference is achieved through the use of a repeat name rather than a pronoun. In Chapter 2 we defined focus operationally. Our first off-line study indicated that implicit causality exerts a focusing influence on language production. An entity can be considered to be in focus if it is the preferred referential link between the current and following utterance. In the case of NP1 biasing verbs there is a preference to interpret the pronoun as coreferential with the first mentioned Noun Phrase. Conversely, in the case of NP2 biasing verbs there is a preference to interpret the pronoun as coreferential with the second mentioned Noun Phrase. The first NP is in focus following a sentence containing an NP1 biasing verb and the second NP is in focus following a sentence containing an NP2 biasing verb.

The natural consequence of the above is that Centering predicts a penalty associated with reading a repetition of the first NP in the context of an NP1 biasing verb and a penalty associated with reading a repetition of the second NP in the context of an NP2 biasing verb. We would therefore expect to find a processing difficulty in the sentence fragment containing the repeat name. So, the first fragment in (11) and (12) below should be read more quickly than (13) and (14):

(11) John fascinated Bill because Bill / was easily interested.

(12) John blamed Bill because John / was in a bad mood.

(13) John fascinated Bill because John / was very interesting.

(14) John blamed Bill because Bill / broke the window.

The above prediction was not made for Experiment 2 as we do not believe that a whole sentence reading time measure is sensitive enough to detect effects that may arise locally on the repeat name. As Experiment 2 only utilised a whole sentence reading time measure

we did not propose that such effects could be detected. Only with a finer grained measure do we think the above Centering predictions can be examined.

Recall Vonk, Hustinx and Simons (1992) proposed that an anaphor of a form more referentially specific than necessary will be interpreted as signalling a thematic shift. It could be argued that repetition of a repeat name in cases where a sentence continues against the bias of the verb will be ultimately beneficial from a processing perspective as it effectively signals a shift in the causal locus from the expected character to the unexpected. The Vonk et al position also predicts that (11) and (12) will be read more quickly than (13) and (14) but this prediction is for reading time to the second fragment. If the repeat name correctly signals a shift, subsequent information consistent with a shift of theme of some description having taken place should be read more quickly than information not consistent with a shift.

Although the predictions made by Centering Theory and the Vonk et al proposal are the same, the point at which the predicted effect will be found differs.

Experiments 3 and 4 are identical to each other except in the type of question that followed the experimental and filler items. We were interested in the influence of depth of processing on the previously reported findings of implicit causality and repeat name penalty. A markedly different effect under different reading conditions would suggest a degree of strategic control over the application of information. Recall Garnham, Oakhill and Cruttenden (1992) suggested that the strength of influence of gender information may to some degree be under the control of the reader. They looked at the effect of embedding experimental materials in a set of fillers that either drew the subjects attention away from or focused it on the task of pronoun resolution. Greene, McKoon and Ratcliff (1992) argue that under many conditions pronouns are never resolved. To be sure that we are measuring the impact implicit causality has on anaphor resolution we introduced a factor which either required subjects to resolve the referent of the anaphor or didn't require them to do so.

We manipulated the type of question following the trials as a between Experiment factor. As in Experiment 2 reported above, in Experiment 3 a question of the type 'Who broke the window ?' followed every trial. This question can only be successfully answered if subjects have resolved the anaphor in the preceding sentence. In Experiment 4 the question was of the type 'Was a window broken ?'. In this case a question followed only a third of the sentences. It could be answered without subjects having resolved the anaphor in the preceding sentence.

Method

Subjects

Thirty two English speaking subjects participated.

Stimuli

As in Experiment 2 we manipulated anaphor, verb bias and referent. Anaphor could either be pronoun or repeat name. Verb bias could either be NP1 or NP2 biasing. Referent could either be the first or second character.

We examined the same experimental materials as were used in Experiment 2. There were 48 sets of experimental materials (see appendix 1 for full set).

Name / NP1 verb / Referent character1

Barry fascinated Derek because Barry / performed magic tricks.

Name / NP1 verb / Referent character2

Barry fascinated Derek because Derek / was easily entertained.

Pronoun/ NP1 verb / Referent character1

Barry fascinated Derek because he / performed magic tricks.

Pronoun / NP1 verb / Referent character2

Barry fascinated Derek because he / was easily entertained.

As in Experiment 3, there were also 96 filler items (appendix 8). The experimental sentences were identical to those used in Experiment 3. The experiment was divided into two halves. The first occurrence of the verb was in the first half of the study and the second in the second half.

Procedure

Experiments 3 and 4 were run in an identical manner to Experiment 2. Rather than the sentences being presented as wholes, presentation was in two halves. Each material was split following the anaphor (see example materials above).

Each experiment lasted roughly 35 minutes.

Results

We shall first report the results for Experiment 3 where the experimental questions required resolution of the anaphor.

4.4.2 Experiment 3 Results

Outlier Replacement

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. We report analyses for the two fragments of the experimental sentences themselves and also for the questions following those sentences. For fragment 1, we excluded response times that were below 200 msec or above 17 seconds. This accounted for 0.6% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 3.2% of the data was replaced in this way. For fragment 2, we excluded response times that were below 200 msec or above 15 seconds. This accounted for 0.1% of the

data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 3.0% of the data was replaced in this way. For the question data, times falling above or below two and a half standard deviations from the mean for each subject were replaced with that point. 3.3% of the data was replaced in this way.


First Fragment Analysis

For reading time to the first fragment, we performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors. We found a main effect of referent that was significant by subjects only ($F(1,31)=13.33$, $p<0.001$; $F(1,46)=2.52$, $p<0.12$) with reference to the first mentioned character taking longer to read than reference to the second. This was independent of verb bias (both $F_s<1$). No other main effects or interactions approached significance. As the Pronoun conditions are identical in fragment 1 we examined the Name conditions separately. The main effect of referent was closer to significance by items now ($F(1,31)=12.13$, $p<0.005$; $F(1,46)=3.77$, $p<0.0584$) (see Figure 5).

Table 5 - Reading times for fragment 1, subject means for each condition with all times in msecs.

Verb / Referent	Name	Pronoun
NP1 /Char1	3086	2799
NP1 / Char2	2793	2713
NP2 / Char1	2976	3714
NP2 / Char2	2648	2639

Figure 5 - main effect of referent (repeat name anaphor)




Second Fragment Analysis

For reading time to the second fragment we performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs treating subjects and items as random factors. We found a main effect of anaphor type ($F(1,31)=15.62, p<0.0005$; $F(1,46)=19.42, p<0.0001$). We found an interaction between verb bias and referent ($F(1,31)=43.95, p<0.0001$; $F(1,46)=13.88, p<0.0005$), i.e. an implicit causality congruency effect (Figure 6). We also found an interaction between anaphor type and referent ($F(1,31)=11.28, p<0.005$; $F(1,46)=6.76, p<0.05$).

Table 6 - Reading times for fragment 2, subject means for each condition with all times in msec.

Verb / Referent	Name	Pronoun
NP1 / Char1	2078	2238
NP1 / Char2	2299	2794
NP2 / Char1	2418	2386
NP2 / Char2	1915	2375

Figure 6 - verb bias x referent




Although our 3-way interaction between congruency (verb bias x referent) and anaphor type was not significant (both $F_s < 1$), we also examined the 2-way interactions corresponding to the implicit causality congruency effect individually for the Pronoun and Name conditions. The congruency effect for the Pronoun conditions is reported first.

Pronoun Conditions

We found a main effect of referent ($F_1(1,31)=11.90$, $p < 0.005$; $F_2(1,46)=5.01$, $p < 0.05$) with endings following pronominal reference to the first mentioned character read more quickly than endings following pronominal reference to second. Our verb bias x referent interaction was also significant ($F_1(1,31)=14.70$, $p < 0.001$; $F_2(1,46)=5.43$, $p < 0.05$) (see Figure 7).

Figure 7 - verb bias x referent (pronoun conditions)




We performed 2-way ANOVAs for NP1 and NP2 biasing verbs separately. For NP1 verbs we found that Char1 and Char2 referents differ ($F(1,31)=26.19$, $p<0.0001$; $F(1,46)=9.51$, $p<0.01$). For NP2 verbs we found that Char1 and Char2 referents are equivalent (both $F_s<1$). We also performed 2-way ANOVAs for Char1 and Char2 referents separately. For Char1 referents we found that NP1 and NP2 verbs are equivalent ($F(1,31)=2.69$, $p<0.12$; $F_2<1$). For Char2 referents we found that NP1 and NP2 verbs differ ($F(1,31)=9.89$, $p<0.005$; $F(1,46)=5.47$, $p<0.05$). This means that the condition corresponding to integration following reference to the second mentioned character is slow when not supported by verb bias. All other conditions are equally fast. Integration following reference to the first mentioned character regardless of verb bias and integration following reference to the second mentioned character in the context of an NP2 biasing verb is equivalently easy.

Name Conditions


We also applied the above set of analyses to the Name conditions. We found an interaction between verb bias and referent ($F(1,31)=25.82$, $p<0.0001$; $F(1,46)=10.50$, $p<0.005$) (see Figure 8).

Figure 8 - verb bias x referent (name conditions)



Recall when taking the Pronoun and Name conditions together, we found an interaction between referent and anaphor. We examined the nature of the interaction between referent and anaphor (see Figure 9) by carrying out means comparisons on all possible condition pair combinations.

Figure 9 - referent x anaphor



Means comparisons reveal the points corresponding to a pronoun referring to the first mentioned character and a name referring to the first mentioned character to be equivalent (both $F_s < 1$). The points corresponding to a name referring to the first mentioned character and a name referring to the second mentioned character are also equivalent ($F_1(1,31)=2.63$, ns; $F_2 < 2$). All other points differ significantly by both subjects and items ($F_1: p < 0.05$; $F_2: p < 0.05$) except for the point corresponding to a pronoun referring to the first mentioned character and a name referring to the second ($F_1: p < 0.05$; $F_2: p < 0.08$). This means that integration following names referring to either the first or second mentioned character is equally easy. Integration following a pronoun referring to the first mentioned character is easier than integration following a pronoun referring to the second mentioned character.

Question Response Time Analysis

For the question response time data we performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. The main effects of verb bias ($F_1(1,31)=16.44$, $p < 0.0003$; $F_2(1,46)=18.70$, $p < 0.0001$) and anaphor type ($F_1(1,31)=35.123$ $p < 0.0001$;

$F_2(1,46)=8.96, p<0.005$) persisted. We found an interaction between referent and verb bias ($F_1(1,31)=30.10, p<0.0001; F_2(1,46)=8.56, p<0.01$) indicating the persistence of the implicit causality congruency effect. We also found an interaction between anaphor type and verb bias ($F_1(1,31)=6.77, p<0.05; F_2(1,46)=5.87, p<0.05$).

Table 7 - Question response times, subject means for each condition with all times in msec. Numbers in parentheses correspond to response accuracy expressed as percentage correct.

Verb / Referent	Name	Pronoun
NP1 / Char1	1982 (96.9%)	2673 (73.4%)
NP1 / Char2	2205 (81.3%)	3194 (67.7%)
NP2 / Char1	2163 (95.8%)	2721 (80.2%)
NP2 / Char2	1747 (79.2%)	2158 (84.4%)

Question Response Accuracy Analysis

We also examined the proportion of correct and incorrect responses for each experimental condition.

We performed 2 (Name vs Pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type ($F_1(1,31)=27.35, p<0.0001; F_2(1,46)=63.01, p<0.0001$) with more correct answers following sentences containing repeat name anaphors and a main effect of verb bias ($F_1(1,31)=17.07, p<0.0005; F_2(1,46)=18.02, p<0.0001$) with responses more accurate following NP2 biasing verbs. We found a significant interaction of anaphor type x verb bias ($F_1(1,31)=32.96, p<0.0001; F_2(1,46)=16.96, p<0.0005$). We also found an interaction between verb bias and referent ($F_1(1,31)=4.86, p<0.05; F_2(1,46)=6.42, p<0.05$) reflecting the implicit causality congruency effect. A 3-way anaphor x verb bias x referent interaction was also significant but only by subjects ($F_1(1,31)=4.87, p<0.05; F_2<2$) suggesting a larger anaphor x verb bias interaction following second character reference than following first.

4.4.3 Experiment 4 Results

The following analyses are for Experiment 4. Recall the only difference between this Experiment and the previous one is the nature of the question subjects have to answer. Experiment 4 contained questions that could be answered without subjects having resolved the anaphor. These questions only appeared on a third of trials.

Outlier Replacement

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. We report analyses for the two fragments of the experimental sentences themselves and also for the questions following those sentences. For fragment 1, we excluded response times that were below 150 msec or above 14 seconds. This accounted for 1.0% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.7% of the data was replaced in this way. For fragment 2, we excluded response times that were below 250 msec. This accounted for 0.7% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.6% of the data was replaced in this way.

First Fragment Analysis


For reading time to the first fragment, we performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors. We found a main effect of referent ($F(1,31)=11.89, p<0.005$; $F(2,46)=5.16, p<0.05$) with reference to the first mentioned character taking longer to read than reference to the second. We also found an interaction between anaphor type and referent ($F(1,31)=10.96, p<0.005$; $F(2,46)=6.59, p<0.05$). With respect to the main effect of referent reported, as the Pronoun conditions are identical in fragment 1 we examined the Name conditions separately. We found a main

effect of referent ($F(1,31)=20.31, p<0.0001; F(1,46)=12.63, p<0.001$) (see Figure 10).

Table 8 - Reading time for fragment 1, subject means for each condition with all times in msecs.

Verb / Referent	Name	Pronoun
NP1 / Char1	2366	2041
NP1 / Char2	2069	2117
NP2 / Char1	2273	2028
NP2 / Char2	2012	2032

Figure 10 - main effect of referent (repeat name anaphor)



Second Fragment Analysis


For reading time to the second fragment we performed 2 (Name vs Pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias that was marginal by items ($F(1,31)=5.89, p<0.05; F(1,46)=3.14, p<0.1$). A main effect of referent was marginal by subjects and items ($F(1,31)=3.10, p<0.1; F(1,46)=2.85, p<0.1$). We found an interaction between verb bias and referent ($F(1,31)=5.49, p<0.05; F(1,46)=7.70, p<0.01$), i.e. an implicit causality congruency

effect (Figure 11). Although our 3-way interaction was not significant ($F(1,31)=2.64$, $p<0.12$; $F_2<1$) the congruency effect was present in the Pronoun conditions ($F(1,31)=8.57$, $p<0.01$; $F_2(1,46)=7.70$, $p<0.01$) but absent in the Name conditions (both $F_s<1$). We also found a marginal interaction between anaphor type and referent that was significant by subjects only ($F(1,31)=4.52$, $p<0.05$; $F_2<2$).

Table 9 - Reading time for fragment 2, subject means for each condition with all times in msecs.

Verb / Referent	Name	Pronoun
NP1 / Char1	1624	1617
NP1 / Char2	1668	1910
NP2 / Char1	1621	1618
NP2 / Char2	1562	1594

Figure 11 - verb bias x referent (pronoun conditions)



To examine the nature of the implicit causality congruency effect found in the pronoun conditions, we performed 2-way ANOVAs for NP1 and NP2 biasing verbs separately. NP1 verbs we found that Char1 and Char2 referents differ ($F(1,31)=11.77$, $p<0.005$; $F_2(1,46)=6.10$, $p<0.05$). For NP2 verbs we found that Char1 and Char2 referents are equivalent (both $F_s<1$). We also performed 2-way ANOVAs for Char1 and Char2 referents separately. For Char1

referents we found that NP1 and NP2 verbs are equivalent (both $F_s < 1$). For Char2 referents we found that NP1 and NP2 verbs differ ($F_1(1,31)=8.056$, $p < 0.01$; $F_2(1,46)=6.99$, $p < 0.05$). This means that integration following names referring to either the first or second mentioned character is equally easy. Integration following a pronoun referring to the first mentioned character is easier than integration following a pronoun referring to the second mentioned character.

4.4.4 Comparing Experiments 3 and 4

To examine what influence our depth of processing manipulation had on the strength of the repeat name penalty associated with repetition of the first mentioned character's name and on the implicit causality congruency effect we performed a separate analysis treating this manipulation as a between subjects factor.

For reading time to the first fragment we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) x 2 (Experiment 3 vs Experiment 4) ANOVAs for both subjects and items as random factors for repeat name anaphors. Recall the repeat name penalty corresponds to a main effect of referent. We found no interaction between referent and experiment (both $F_s < 1$) indicating that there was no difference in the nature of the repeat name penalty as a result of our depth of processing manipulation.

For reading time to the second fragment we also performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) x 2 (Experiment 3 vs Experiment 4) ANOVAs for both subjects and items as random factors. We found an interaction between the implicit causality congruency effect (verb bias x referent) and experiment ($F_1(1,62)=10.83$, $p < 0.005$; $F_2(1,92)=8.57$, $p < 0.0001$). This corresponds to the implicit causality congruency effect being weaker in Experiment 4 where the questions following the materials could be answered without the reader necessarily having resolved the pronoun. Our depth of processing manipulation did have a consequence for the nature of the implicit causality congruency effect.

4.4.5 Discussion --

In both Experiments 3 and 4 we found a repeat name penalty associated with a repetition of the first mentioned character's name. The strength of the penalty wasn't affected by our depth of processing manipulation suggesting that it may arise at a relatively low level of processing. Within each experiment, the repeat name penalty wasn't affected by verb bias suggesting that implicit causality information wasn't used as soon as the anaphor was encountered. This result is not consistent with the predictions made by Centering. Neither did we find evidence for the prediction derived from the Vonk et al position that following a repeat name a continuation going **against** the bias of the verb would be read quickly.

The earliest point at which an influence of implicit causality was found was on reading time to the second fragment. This suggests that it exerts its influence as the information contained within the subordinate clause is integrated with that contained in the main clause.

We also found an interaction between anaphor type and referent (although it was marginal by items in Experiment 4). The pattern suggests that following a name, ease of integration is equivalent regardless of whether the name corresponds to the first or second mentioned character. The repeat name penalty found in analysis of reading times to fragment 1 has disappeared in fragment 2. Following a pronoun referring to the second mentioned character, integration is difficult compared to the case where the pronoun refers to the first mentioned character. We argue that names uniquely (and quickly) identify their antecedents but that pronouns take some time to achieve this same level of referent identification. We propose that antecedent identification or co-indexing for repeat name anaphors occurs when that repeat name is read. This co-indexation occurs with reference to only low level information (leading to the repeat name penalty) and doesn't take higher level factors such as implicit causality into account.

Information following a pronoun is important for indicating to which character that pronoun refers so we argue that the stage of co-indexation is delayed in the case of pronominal anaphora. It takes longer to determine a pronoun's antecedent if the continuation in the subordinate clause goes against the direction of verb bias. This is because the two factors of verb bias and the semantics of the subordinate clause contradict each other. Integration of information following pronominal reference to the first mentioned character appears easy regardless of whether it is supported by verb bias. This may be due to the first mentioned character occupying some privileged position within the reader's discourse model (cf. Gernsbacher and Hargreaves, 1988).

Although we found a decrease in the strength of the implicit causality congruency effect from Experiment 3 to Experiment 4, the effect *per se* was present in both cases. It does suggest that some factors pertinent to reading are not completely under strategic control, contrary to what Garnham, Oakhill and Cruttenden (1992) and Greene, McKoon and Ratcliff (1992) suggest. Even when subjects do not have to resolve the pronoun, they nevertheless perform enough processing for implicit causality to exert an influence. This contradicts the specific claim made by Greene, McKoon and Ratcliff (1992) that pronouns are not resolved under conditions where such resolution is not necessary but also contradicts the broader claims made in their Minimalism proposal that readers never engage in more processing than is required (either because of demands of the text being read or the experimental demands themselves).

We propose that subjects do resolve pronouns even under experimental circumstances where this is not strictly necessary.

4.5 General Discussion

Experiments 1a and 1b provided a control which has so far been lacking in the experimental examination of not just the influence of implicit causality information on anaphoric processing but also in the examination of anaphor resolution in general. They provided us with a measure of the strength of the implicit causality biases associated

with the verbs we examined and also the relative plausibilities of the events described in the sentences we examined. Experiment 2, our first examination of the on-line influence of implicit causality, replicated the basic implicit causality congruency effect that has been found in the literature. In light of the data from Experiments 1a and 1b we can confidently interpret this as an implicit causality effect rather than the simple plausibility effect that we outlined. Experiment 2 also indicated a repeat name penalty associated with reference to the first mentioned character.

Experiments 3 and 4 provided an insight into the nature of anaphoric processing. If implicit causality information exerts an influence at the point at which an anaphor is encountered we would have expected a difficulty reflected in longer reading times to the first fragment in examples (15) and (16) below.

(15) John fascinated Bill because Bill / was easily interested.

(16) John blamed Bill because John / was in a bad mood.

The repeat name indicates that the continuation in the subordinate clause will go against the direction of the verb bias. The only effect we found on reading time to the first fragment was a repeat name penalty associated with a repetition of the first mentioned character's name. There was no interaction with verb bias suggesting that implicit causality information isn't used to guide anaphoric processing at this point. An interaction between implicit causality and referent was found on reading time to fragment 2 suggesting that implicit causality exerts its influence at the integrative stage of processing.

The finding of the repeat name penalty only with respect to the first mentioned character is inconsistent with the predictions made by Centering Theory. Experiment 1a indicated that focus was driven by the verb bias. In Centering terminology the forward looking centre contains whichever character is supported by verb bias. Either Centering Theory cannot account for shifts in focus within conjoined sentences, or focus as we are defining it is not focus as defined by Centering.

Initial anaphoric processing does not appear to use implicit causality information. The lack of an interaction between the repeat name penalty and verb bias on reading time to the first fragment suggests that the penalty is not due to a violation of the preferred referring device for the entity in focus (given our operational definition of focus where the character in focus is the preferred character for subjects to continue with reference to in a production task). It may instead be due to lower level structural factors (say, first mentioned character) or to phenomena unrelated to language processing. It is not our intention in this thesis to provide an exhaustive test of Centering Theory but where pertinent our results shall be related to the predictions made by it.

CHAPTER 5

5.1 Introduction --

In the preceding chapter we reported three on-line examinations of the time course of the influence of implicit causality information within the anaphor resolution system. We found two basic results. The first was that implicit causality appears to exert the bulk of its influence at the integrative stage of processing. The second finding was of a repeat name penalty associated with reference to the first mentioned character through the use of a repeat proper name. This penalty is independent of implicit causality bias. We argue that this supports our general two-stage account of anaphor resolution. For repeat name anaphors the repeat name penalty arises at the stage of co-indexation. This is informed by low level information. Semantic information exerts a processing influence later. For repeat name anaphors reading time to fragment 1 in Experiments 3 and 4 reflects co-indexation processes. Reading time to fragment 2 reflects the second stage of semantic integration as the information predicated by the anaphor is integrated with information contained in the main clause. For pronominal anaphors co-indexation is delayed until information necessary for unambiguous identification of the pronominal antecedent is available. For pronouns, reading time to fragment 2 reflects both stages of co-indexation and semantic integration.

Three experiments are reported in this chapter. The first (Experiment 5) examines the behaviour of the repeat name penalty and the influence of implicit causality information in cases where the information contained within the main and subordinate clauses is presented in two main clauses (i.e. as separate sentences). We want to separate effects that we may have found in our previous experiments as a result from the way in which the information was presented to the reader from those effects that arise as a result of the actual content of the information. It may be the case that within sentence focus behaves in a manner very different from focus between sentences. Consider Examples (1) and (2) below. The same information is presented but rather than subjects reading materials

such as Example (1), the clauses are separated and the materials are of the form as in Example (2).

(1) John blamed Bill because he/Bill broke the window.

(2) John blamed Bill. This was because he/Bill broke the window.

The second experiment in this chapter (Experiment 6) examines the influence of implicit causality information in the context of a gender marked pronoun which unambiguously identifies its referent; see example (3).

(3) John blamed Mary because she broke the window.

The pronominal antecedent can be identified on the basis of gender information alone. Implicit causality information is not necessary for this identification process to succeed. We want to examine at what stage of our proposed model gender information is utilised by the system. Experiment 7 is a replication of Experiment 6 but using an eye-tracking methodology. This allows us to obtain a much finer grained measurement of subjects' reading through monitoring eye-movements. The methodology will be discussed in more detail below when we focus on Experiment 7. Experiments 5 and 6 employ the same self-paced reading methodology as has been used up to this point.

5.2 Experiment 5 --

5.2.1 Introduction and Rationale

As briefly mentioned in the introduction, Experiment 5 differs in one important way from the experiments so far discussed in this thesis. The information previously presented in main and subordinate clauses is now presented in two main clauses in separate sentences (see example (4)).

(4) John fascinated Bill. This was because John was very interesting.

The overt causal link between the utterances is maintained but in this case it takes the form of the phrase 'This was because'. As this appears to be more of an overtly marked form we propose that it explicitly directs the readers' attention to the underlying cause of the described event. This manipulation should make no difference to the manner in which implicit causality information is employed by the reader. We wanted to separate out effects that may have arisen as a result of the manner in which the information was presented to the reader from effects that arise because of the content of the information presented, regardless of the manner of that presentation. The two consistent effects we have found up to this point are the repeat name penalty associated with repeated reference to the first mentioned character, found in the pattern of reading times to the first fragment, and the implicit causality congruency effect, found in the pattern of reading times to the second fragment.

With respect to the repeat name penalty, Centering Theory makes the same predictions regardless of whether two clauses are presented in a main-subordinate clause pair or as two separate main clauses. Recall, the circumstances necessary for the repeat name penalty to be found. The most highly ranked forward looking Centre will be preferentially realised using a pronoun if it is the following utterance's backward looking Centre. If it is realised using a repeat name anaphor, a reading time penalty will be accrued. Although we find a repeat name penalty associated with the first mentioned character, this is not the character occupying the most highly ranked position within the forward looking centre. It is possible that our definition of focus (i.e. defined operationally) differs from that definition adopted by the Centering theorists. However, regardless of how we precisely define focus, Centering does not predict that we should find any difference in terms of the repeat name penalty between this experiment and those previously reported experiments examining main-subordinate clause order sentences.

We have no reason to believe that the repeat name penalty will be in any different in nature from the penalty reported in Experiments 2, 3 and 4 in the previous chapter, i.e. associated with repetition of the

first mentioned character's name and independent of higher level semantic factors.

The only requirements reported in the literature necessary for the implicit causality effect to be found is that there is some form of causal link associating the event with its cause (Ehrlich, 1980). That is present in this study in the form of the phrase 'This was because'. We therefore expect to find the same pattern of data reflecting an implicit causality congruency effect as we found in the experiments reported in the previous chapter.

Method

Subjects

Thirty two English speaking subjects participated.

Stimuli

We manipulated anaphor, verb bias and referent. Anaphor could either be pronoun or repeat name. Verb bias could either be NP1 or NP2 biasing. Referent could either be the first or second character.

There were 48 sets of experimental materials (see appendix 2 for full set).

Name / NP1 verb / Referent character1

Barry fascinated Derek. This was because Barry performed magic tricks.

Name / NP1 verb / Referent character2

Barry fascinated Derek. This was because Derek was easily entertained.

Pronoun / NP1 verb / Referent character1

Barry fascinated Derek. This was because he performed magic tricks.

Pronoun / NP1 verb / Referent character2

Barry fascinated Derek. This was because he was easily entertained.

There were also 96 filler items (see appendix 8). The experimental materials were the same as those in Experiment 2 but with the information presented in separate sentences as in the above examples. Each subject saw each verb twice. The experiment was divided into two halves with a break halfway through. The first occurrence of each verb was in the first half of the study and the second occurrence in the second half.

Procedure

The procedure was identical to Experiment 2 but with two experimental sentences presented instead of one. The sentences were presented one at a time.

Each subject participated in 10 practice trials similar in structure to the experimental items at the start of the experiment. The experiment lasted roughly 35 minutes. Before the experiment they were provided with both verbal instructions and written instructions modified from those given to subjects in Experiment 2.

5.2.2 Results

We report analyses for the two experimental sentences themselves and also for the question responses following those sentences.

Outlier Replacement

For reading times to the first sentence we excluded response times that were below 250 msec or above 13 seconds. This accounted for 0.8% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 3.13% of the data was replaced in this way.

For the data corresponding to reading times to the second sentence we excluded response times that were below 400 msec or above 13 seconds. This accounted for 0.5% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.0% of the data was replaced in this way.

For the question response data we excluded response times that were below 400 msec or above 13 seconds. This accounted for 0.7% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 3.0% of the data was replaced in this way.

First Sentence Analysis

We performed a 1-way ANOVA (NP1 verb vs NP2 verb) for both subjects and items as random factors. We found an effect of verb bias ($F(1,31)=8.91, p<0.01$; $F(1,46)=6.70, p<0.05$). This corresponds to sentences containing NP2 verbs being read more quickly than sentences containing NP1 verbs.

Table 1 - Reading times for first sentence, subject means for each condition with all times in msec.

Verb	Reading Time (msec)
NP1	1984
NP2	1809

Second Sentence Analysis

We performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type with sentences containing repeat name anaphors being read more quickly than sentences containing pronominal anaphors ($F(1,31)=6.32, p<0.05$; $F(1,46)=6.03, p<0.05$). We also found a main effect of verb bias with sentences following NP2 verbs being read more quickly than those following NP1 verbs ($F(1,31)=15.17, p<0.0005$; $F(1,46)=4.69, p<0.05$). We found an interaction between


anaphor type and referent ($F(1,46)=11.31, p<0.005; F(1,46)=4.91, p<0.05$) (see Figure 1 below). We also found an interaction between verb bias and referent corresponding to the implicit causality congruency effect that was marginal by items ($F(1,31)=6.62, p<0.05; F(1,46)=3.47, p<0.07$) (see Figure 2 below). We also found an anaphor type x verb bias x referent interaction ($F(1,31)=6.44, p<0.05; F(1,46)=7.82, p<0.01$).

Table 2 - Reading times for second sentence, subject means for each condition with all times in msec.

Verb / Referent	Name	Pronoun
NP1 / Char1	2678	2589
NP1 / Char2	2594	3081
NP2 / Char1	2540	2664
NP2 / Char2	2455	2521


We first explored the nature of the anaphor type x referent interaction.

Figure 1 - anaphor type x referent



In order to interpret this interaction we performed means comparisons on each possible pair of conditions. The comparisons reveal that the conditions corresponding to pronominal and repeat

name reference to the first mentioned character are statistically equivalent (both $F_s < 1$). The conditions corresponding to a repeat name anaphor referring to either the first or second mentioned character are also equivalent ($F_1(1,31)=2.408$, $p < 0.14$; $F_2 < 2$). The difference between the conditions corresponding to a pronoun referring to the first mentioned character and a name referring to the second mentioned character is marginal ($F_1(1,31)=3.499$, $p < 0.07$; $F_2 < 2$). All other condition pairs differ significantly from each other (at least $p < 0.05$ for F_1 and F_2).



We further explored the nature of this 2-way interaction by examining the Name and Pronoun conditions separately. This was qualified by our finding of a 3-way interaction of anaphor x verb bias x referent.


Name Conditions

Looking at the Name conditions first, we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias that was marginal by subjects and non-significant by items suggesting sentences following NP2 verbs being read more quickly than

sentences following NP1 verbs ($F(1,31)=4.13$, $p<0.06$; $F_2<2$). The verb bias x referent interaction was not significant (both $F_s<1$).

Pronoun Conditions

For the Pronoun conditions, we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias with sentences following NP2 verbs being read more quickly than sentences following NP1 verbs ($F(1,31)=13.03$, $p<0.005$; $F_2(1,46)=4.85$, $p<0.05$). We also found a marginal main effect of referent with reference to the first mentioned character read more quickly than reference to the second that was only significant by subjects ($F(1,31)=6.85$, $p<0.05$; $F_2(1,46)=2.788$, $p<0.11$). We found a verb bias x referent interaction corresponding to the implicit causality congruency effect ($F(1,31)=11.68$, $p<0.005$; $F_2(1,46)=9.55$, $p<0.005$) (see Figure 3 below).



We performed 2-way ANOVAs for NP1 and NP2 biasing verbs separately. For NP1 verbs we found that Char1 and Char2 referents differ ($F(1,31)=14.36$, $p<0.001$; $F_2(1,46)=11.05$, $p<0.005$). For NP2 verbs we found that Char1 and Char2 referents are equivalent

($F(1,31)=2.19$, $p<0.15$; $F_2<2$). We also performed 2-way ANOVAs for Char1 and Char2 referents separately. For Char1 referents we found that NP1 and NP2 verbs are equivalent (both $F_s<1$). For Char2 referents we found that NP1 and NP2 verbs differ ($F(1,31)=31.48$, $p<0.0001$; $F_2(1,46)=12.05$, $p<0.005$). This means that the condition corresponding to a sentence reference to the second mentioned character is slow when not supported by verb bias. All other conditions are equally fast. Sentences containing reference to the first mentioned character regardless of verb bias and sentences containing reference to the second mentioned character in the context of an NP2 biasing verb are read equivalently quickly.

In light of our focusing on the Pronoun and Name conditions separately, the 3-way interaction of anaphor type x verb bias x referent can be interpreted as corresponding to a difference with respect to the implicit causality congruency effect in the Pronoun conditions compared to the Name conditions. Basically, we found an implicit causality congruency effect (verb bias x referent interaction) in the Pronoun conditions, but not in the Name conditions. Reasons for this shall be discussed in detail below.

Question Response Time Analysis -

For the question response time data we again performed 2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a persistence of the main effect of anaphor type ($F(1,31)=36.92$, $p<0.0001$; $F_2(1,46)=66.01$, $p<0.0001$) and verb bias ($F(1,31)=17.27$, $p<0.0005$; $F_2(1,46)=5.97$, $p<0.05$). We also found a main effect of referent that was marginal by items ($F(1,31)=5.95$, $p<0.05$; $F_2(1,46)=3.72$, $p<0.06$) with response time following an ending referring to the first mentioned character being faster than following an ending referring to the second mentioned character. The only interaction that reached significance was that corresponding to the implicit causality congruency effect (verb bias x referent) ($F(1,31)=14.54$, $p<0.001$; $F_2(1,46)=13.61$, $p<0.001$).

Table 3 - Question response times, subject means for each condition with all times in msec. Numbers in parentheses correspond to response accuracy expressed as percentage correct.

Verb / Referent	Name	Pronoun
NP1 / Char1	1694 (97.4%)	2054 (87.9%)
NP1 / Char2	1927 (95.8%)	2666 (81.3%)
NP2 / Char1	1789 (96.3%)	2154 (91.6%)
NP2 / Char2	1609 (94.7%)	2070 (96.3%)

Question Response Accuracy Analysis

In order to determine whether there was some form of a speed/accuracy trade off for response times to the questions, we examined the proportion of correct and incorrect responses for each experimental condition.

2 (Name anaphor vs Pronoun anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs on the proportion of correct responses for both subjects and items as random factors revealed main effects of anaphor type ($F(1,31)=16.98$, $p<0.0005$; $F(1,46)=18.01$, $p<0.0001$) and verb bias ($F(1,31)=8.31$, $p<0.01$; $F(1,46)=5.69$, $p<0.05$). Subjects responded more accurately following sentences containing repeat name anaphors. They also responded more accurately following sentences containing NP2 biasing verbs. Two interactions were also highly reliable. The first was anaphor type x verb bias ($F(1,31)=15.93$, $p<0.0005$; $F(1,46)=13.95$, $p<0.0005$) and the second the implicit causality congruency effect (verb bias x referent) ($F(1,31)=5.11$, $p<0.05$; $F(1,46)=5.91$, $p<0.05$). The 3-way interaction of anaphor type x verb bias x referent was marginal by subjects and by items ($F(1,31)=3.96$, $p<0.06$; $F(1,46)=3.34$, $p<0.08$).

The 3-way interaction (anaphor type x verb bias x referent) is of the same sort as that found in the pattern of reading time data to the second sentence. The implicit causality congruency effect is present in the Pronoun conditions ($F(1,31)=5.60$, $p<0.05$; $F(1,46)=7.09$, $p<0.05$) but is absent in the Name conditions (both $F_s<1$).

5.2.3 Discussion

Our main finding in this experiment is with respect to subjects' reading times for the second sentence (see example (5) below).

(5) This was because Bill/he broke the window / John/he was in a bad mood.

Our lack of evidence for the occurrence of a repeat name penalty in this experiment is curious. In light of our previous findings we would expect a repeat name penalty associated with reading the sentence containing the repetition of the name 'John', i.e. the first mentioned character's name. Recall in Experiment 2 we found a repeat name penalty although the finding was marginal by items. In the current study, no hint of a penalty was observed in any of the analyses we performed. The predictions made by Centering do not differ in such a dramatic fashion as a function of whether intra- or inter-sentential reference is appropriate for linking utterances.

There are two possible ways in which Centering can explain this lack of an effect. The first is due to the nature of the referential link between the two sentences, while the second is due to a specific behaviour associated with over-specification of reference that we earlier described in section 2.2.4 above. The first account that Centering theorists may propose to explain the lack of a repeat name penalty is that the referential link between the two utterances is the cause of the event, as signalled by the phrase 'This was because'. This conceptual link occupies the backward looking Centre rather than either of the participants mentioned in the preceding utterance. The conditions necessary for the repeat name penalty to arise (i.e. that the backward looking Centre should preferentially be realised as a pronoun) are not met and no penalty occurs.

The second explanation evokes Vonk's position regarding thematic shifts which we also refer to in an attempt to explain the lack of an implicit causality congruency effect associated with the Name conditions. We shall come to this explanation below following a summary of other findings we obtained from this experiment.

The behaviour of repeat names and pronominal anaphora differs where inter-sentential reference is required, as it is in this Experiment, compared to the case where intra-sentential reference is employed (Experiments 2-4). In the Pronoun conditions, the implicit causality congruency effect of the same type as has been observed in our previous studies was found. The interaction between verb bias and referent is driven by the point corresponding to reference to the second mentioned character in the context of an NP1 biasing verb. As reported with respect to Experiments 2-4 in the preceding chapter, reference to the first mentioned character is always easy, reference to the second is only easy when supported by verb bias.

This implicit causality effect was not found in the Name conditions however. We consistently found the effect in our previous studies and never found an interaction between anaphor type and the implicit causality congruency effect as we observed in this experiment. As the only difference between this current experiment and those reported in Chapter 4 is in terms of the manner in which the information is presented (two separate main clause sentences rather than as main-subordinate clause order sentence) we can look for an explanation as a result of this structural difference.

How might recourse to Vonk et al's position regarding thematic shifts within the context of the 2-sentence pairs we examined account for both the lack of a repeat name penalty and the implicit causality congruency effect? Recall Vonk et al propose that over-specification of reference is used to explicitly signal the presence of a thematic shift. In section 2.2.5 above in light of the findings reported by Vonk et al and by Cloitre and Bever (1988) we attempted to explain the mechanism by which over-specification facilitates thematic shift comprehension by suggesting that over-specification results in a shift within the reader's discourse model from the conceptual to the more superficial level. A shift in the level of grain within the reader's discourse model is necessary for integration of the information following the thematic shift. Therefor shifting as a result of encountering an overspecific anaphor immediately prior to such a shift being required facilitates integration.

No repeat name penalty arises as the reader interprets what follows the anaphor in the second sentence as a separate topic and so doesn't relate it to the information contained within the previous sentence. The shift can be interpreted as being from the act of 'John blaming Bill' to the actual underlying cause of the event. In the case of the experiments reported in Chapter 4, the information about the cause of the event is embedded within the topic of the sentence as a whole. In the current experiment, each sentence effectively possesses a separate topic.

We propose this explanation can account for the lack of finding of an implicit causality congruency effect in the Name conditions in the current experiment. The repeat name anaphor is interpreted as a thematic shift indicator and the information predicated by it is not integrated with the information contained in the preceding sentence. A marginal implicit causality congruency effect was found in the pattern of data corresponding to the question response times indicating that at some point after reading the second sentence, subjects did attempt to integrate the explicit cause with the information contained in the first sentence. In other words in the Name conditions, during reading of the second sentence, the information contained within it is not integrated with that has been read previously. Some partial integration may take place later as indicated by the hint of an implicit causality effect in the question response time data.

We have argued up to this point that implicit causality exerts a processing influence during integration. If integrative processing does not occur, and we argue it doesn't when the repeat name anaphor is interpreted as a thematic shift indicator, implicit causality simply does not get a chance to exert a processing influence.

Analysis of subjects' question response accuracy indicates better accuracy following a sentence containing a repeat name anaphor than following a sentence containing a pronominal anaphor. We also found an interaction between anaphor type and the implicit causality congruency effect with the effect present in the Pronoun conditions but absent in the Name conditions.

Another possible explanation for not finding a congruency effect in the Name conditions is that subjects can answer the question without necessarily integrating the information contained within the two sentences. In this experiment, for the experimental materials the question always focused on information contained in the second sentence. Therefore simply attending to the second sentence is sufficient for subjects to answer the question. This also may explain why we didn't find any evidence of an implicit causality congruency effect in the Name conditions. Readers weren't integrating the information contained within the two sentences because they knew they didn't have to. The question could be answered simply by reading the second sentence. In the case of the Pronoun conditions however integration of information contained in the two sentences must take place in order for the question to be answered.

With respect to the implicit causality congruency effect in the Pronoun conditions, the pattern of question response accuracy data again supports our general conclusion. Recall that in reading time the condition driving the verb bias x referent interaction corresponds to reference to the second mentioned character in the context of an NP1 biasing verb. All other points were equivalently fast. This is again reflected in the pattern of response accuracy. The point corresponding to reference to the second mentioned character in the context of an NP1 biasing verb was responded to with less accuracy than for any of the other conditions which were responded to with equivalent high accuracy. This further indicates that there is a genuine problem in integrating reference to the second mentioned character when other cues (e.g. verb bias) indicating which antecedent is appropriate are not available in the preceding text.

Regarding our initial question of how the manner of presentation can account for the effects we have so far located, we can draw two general conclusions. The first is that the manner of presentation has no observable effect as far as the implicit causality congruency effect is concerned when pronominal reference links the two sentences. The second is that the repeat name penalty is dependent on the manner in which the information is presented as is the implicit causality

congruency effect when repeat name anaphors form the referential link between the two sentences. We propose the repeat name penalty is influenced by low level structural factors rather than by any higher level semantic information such as implicit causality. The lack of congruency effect in the Name conditions **may** be due to readers interpreting the repeat name as a thematic shift signal, but it may equally be due to some degree of strategic processing.

Up to this point we have examined the influence of implicit causality information in the absence of any other cues that may be useful for pronoun resolution. How might it influence processing when other constraints are present ?

5.3 Experiment 6 --

5.3.1 Introduction and Rationale

So far we have only looked at the influence of one particular type of non-structural cue which can be employed by the reader to facilitate resolution of ambiguous pronominal reference. Other non-structural cues are potentially available however. In English, pronouns can carry information in terms of number or gender marking to constrain reference. In this experiment we examine the influence of a gender cue when pitted against implicit causality.

In Examples (6) and (7) below, the pronoun can be resolved simply through recourse to the gender information it carries.

(6) John blamed Mary because she broke the window.

(7) John blamed Mary because he was in a bad mood.

With respect to our two-stage model of anaphor resolution, at what point might gender information be employed ? So far we have argued that co-indexation of pronouns is delayed until disambiguating information becomes available. Are pronominal anaphors co-indexed as soon as they are encountered when a gender cue is present ? The presence of a gender contrast certainly allows for co-indexation to

occur as the pronoun is encountered and there will be no misassignment if the system is sensitive to and utilises gender.

As far as implicit causality information goes, all the evidence so far points to it exerting an influence at the integrative stage of processing. Recall we found no interaction between verb bias and referent in Experiments 2-4 where we found evidence of a repeat name penalty. The penalty occurred independently of verb bias. Certainly some form of processing influence is arising as a result of reading the implicit causality verbs at a relatively early point, as evidenced by our consistently found main effect that sentence fragments containing NP2 biasing verbs take less time to read than those fragments containing NP1 biasing verbs. However, there is nothing in our data to suggest that this information is having any influence when a pronoun is encountered.

Evidence that gender information is employed as soon as a pronoun is encountered to inform initial assignment will be in the form of an ambiguity x congruency interaction. In the presence of a gender cue, if gender information is used to co-index a pronoun as soon as it is read, it should be easier to recover from a continuation going against the bias of the verb. In the absence of a gender cue however it should be more difficult to recover from a continuation going against the bias of the verb as no gender information is available and misassignment (on the basis of implicit causality information) will occur. When a gender cue is present, if gender information is used immediately a pronoun is read, we expect to find a weaker implicit causality congruency effect compared to the conditions where it is absent.

What evidence is there in the literature regarding the manner of influence of gender information within the anaphor resolution system? As summarised in Chapter 2, Stevenson and Vitkovitch (1986) suggest that even when gender information is sufficient to identify a pronoun's referent, other factors influence processing. This is consistent with the position that gender information is not used immediately by the system but exerts its influence at the same point in time as other factors. Up to this point we have been proposing that

the initial co-indexation stage of anaphor resolution does not occur immediately when a pronoun is encountered. This is consistent with the data described by Stevenson and Vitkovitch (1986).

In this experiment we attempt to examine the role played by gender information. We again split the sentence presentation following the anaphor to try to separate any immediate (early) processing of the pronoun from other later effects. In Chapter 3 we described the position adopted by Garnham, Oakhill and Cruttenden (1992) proposing that the use of gender information is under the strategic control of the reader. In extreme cases, Greene, McKoon and Ratcliff (1992) suggested that readers do not always fully resolve pronouns. To be sure that the processes associated with anaphor resolution are being measured in this study, as in previous studies we added questions after both the experimental materials and filler items that could only be correctly answered if the anaphor had been successfully resolved.

Method

Subjects

Thirty two English speaking subjects participated.

Stimuli

We manipulated pronominal ambiguity, verb bias and referent. The pronoun could either be referentially ambiguous or marked for gender to unambiguously refer to one of the two participants. Verb bias could either be NP1 or NP2 biasing. Referent could either be the first or second character.

There were 48 sets of experimental materials (see appendix 3 for full set).

Unambiguous pronoun / NP1 verb / Referent character1

Barry fascinated Mary because he performed magic tricks.

Unambiguous pronoun / NP1 verb / Referent character2

Mary fascinated Derek because he was easily entertained.

Ambiguous Pronoun / NP1 verb / Referent character1

Barry fascinated Derek because he performed magic tricks.

Ambiguous Pronoun / NP1 verb / Referent character2

Barry fascinated Derek because he was easily entertained.

There were also 96 filler items (appendix 8). Each subject saw each verb twice. The experiment was divided into two halves with a break halfway through. The first occurrence of each verb was in the first half of the study and the second occurrence in the second half.

Procedure

The procedure was identical to Experiment 3.

Each subject participated in 10 practice trials similar in structure to the experimental items at the start of the experiment. The experiment lasted roughly 35 minutes. Before the experiment subjects were provided with both verbal and written instructions (see Chapter 4 for full instructions).

5.3.2 Results

We report analyses for the two fragments of the experimental sentences themselves and also for the question responses following those fragments.

Outlier Replacement

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For reading times to the first fragment we excluded response times that were below 300 msec or above 15 seconds. This accounted for 2.0% of the data. Times falling above or below two and a half standard deviations from the

mean for each subject were replaced by that point. 3.1% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.6% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the question response data we excluded response times that were below 300 msec or above 15 seconds. This accounted for 1.2% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.7% of the data was replaced in this way.

First Fragment Analysis _

We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a marginal main effect of ambiguity ($F(1,31)=3.55$, $p<0.07$; $F_2(1,46)=3.01$, $p<0.09$) corresponding to sentence fragments containing ambiguous pronouns being read more slowly than those containing unambiguous pronouns. We found a marginal main effect of verb bias that was not significant by ($F_1(1,31)=3.45$, $p<0.08$; $F_2<2$) suggesting sentence fragments containing NP2 biasing verbs being read more quickly than those containing NP1 biasing verbs. We also found a marginal main effect of referent significant by subjects only ($F_1(1,31)=4.98$, $p<0.05$; $F_2(1,46)=2.39$, $p<0.13$). As the Ambiguous Pronoun conditions are identical at this point we examined the Unambiguous Pronoun conditions separately revealing no main effect of referent (both $F_s<1$).

Table 4 - Reading times for first fragment, subject means for each condition with all times in msec.

Verb / Referent	Ambiguous Pronoun	Unambiguous Pronoun
NP1 / Char1	2379	2112
NP1 / Char2	2157	2176
NP2 / Char1	2235	2109
NP2 / Char2	2058	2069


Second Fragment Analysis

We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of ambiguity ($F(1,31)=12.50$, $p<0.005$; $F(1,46)=19.20$, $p<0.0001$) corresponding to fragments following unambiguous reference being read more quickly than fragments following ambiguous reference. We found a main effect of verb bias ($F(1,31)=9.53$, $p<0.005$; $F(1,46)=6.51$, $p<0.05$) with fragments following NP2 biasing verbs being read more quickly than fragments following NP1 biasing verbs. We also found a main effect of referent ($F(1,31)=5.27$, $p<0.05$; $F(1,46)=4.33$, $p<0.05$) corresponding to faster reading time following reference to the first mentioned character than following reference to the second mentioned character. One interaction was significant, that of verb bias x referent ($F(1,31)=24.36$, $p<0.0001$; $F(1,46)=10.23$, $p<0.005$) corresponding to the implicit causality congruency effect (see Figure 4 below). Ambiguity did not interact with any factor.

Table 5 - Reading times for second sentence, subject means for each condition with all times in msec.

Verb / Referent	Ambiguous Pronoun	Unambiguous Pronoun
NP1 / Char1	1983	1696
NP1 / Char2	2234	1980
NP2 / Char1	1893	1769
NP2 / Char2	1919	1641

Figure 4 - verb bias x referent



We performed 2-way ANOVAs for NP1 and NP2 biasing verbs separately. We also performed 2-way ANOVAs for Char1 and Char2 referents separately. For NP1 verbs we found that Char1 and Char2 referents differ ($F(1,31)=14.911$, $p<0.0005$; $F(1,46)=11.890$, $p<0.005$). For NP2 verbs we found that Char1 and Char2 referents are equivalent ($F_1<2$; $F_2<1$). For Char1 referents we found that NP1 and NP2 verbs are equivalent (both $F_s<1$). For Char2 referents we found that NP1 and NP2 verbs differ ($F(1,31)=36.12$, $p<0.0001$; $F(1,46)=18.65$, $p<0.0001$). This means that the condition corresponding to integration following reference to the second mentioned character is slow when not supported by verb bias. All other conditions are equally fast. Integration following reference to the first mentioned character regardless of verb bias and integration following reference to the second mentioned character in the context of an NP2 biasing verb is equivalently easy.

Question Response Time Analysis

For the question response time data we also performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and

items as random factors. We found a main effect of ambiguity ($F(1,31)=41.01$, $p<0.0001$; $F(1,31)=89.92$, $p<0.0001$) with faster response times following sentences containing unambiguous pronouns. We found a main effect of verb bias ($F(1,31)=13.12$, $p<0.001$; $F(1,46)=10.13$, $p<0.005$) with faster response times following sentences containing NP2 biasing verbs. We found a marginal main effect of referent ($F(1,31)=4.14$, $p<0.06$; $F(1,46)<2$) although this was not significant by items. One interaction was significant, that of verb bias x referent corresponding to the implicit causality congruency effect ($F(1,31)=21.10$, $p<0.0001$; $F(1,46)=23.32$, $p<0.0001$).

Table 6 - Question response times, subject means for each condition with all times in msec. Numbers in parentheses correspond to response accuracy expressed as percentage correct.

Verb / Referent	Ambiguous Pronoun	Unambiguous Pronoun
NP1 / Char1	2312 (87.1%)	1810 (96.3%)
NP1 / Char2	2813 (85.3%)	2156 (92.5%)
NP2 / Char1	2453 (90.1%)	1896 (93.8%)
NP2 / Char2	2091 (93.6%)	1757 (99.0%)

Question Response Accuracy Analysis

For the question response accuracy data we also performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. The response accuracy analysis reveals a similar pattern of results to that found for the response time analysis. We found a main effect of ambiguity ($F(1,31)=21.08$, $p<0.0001$; $F(1,46)=15.72$, $p<0.0005$) with greater response accuracy following a sentence containing unambiguous reference. We found a main effect of verb bias ($F(1,31)=6.83$, $p<0.05$; $F(1,46)=5.87$, $p<0.05$) with greater response accuracy following a sentence containing an NP2 biasing verb. We also found the implicit causality congruency effect, verb bias x referent interaction that was significant by subjects only ($F(1,31)=9.02$, $p<0.01$; $F(1,46)<2$).

5.3.3 Discussion

The pattern of data associated with reading time to the second fragment closely parallels that found previously (Experiments 3 and 4 and the Pronoun condition in Experiment 5). We found main effects of verb bias and referent of the same type as we have reported previously. In addition we found a main effect of ambiguity in reading time to fragment 2 with sentence fragments following ambiguous reference read more slowly than those following unambiguous reference. We also found a marginal main effect of ambiguity on fragment 1 reading times suggesting that to some degree readers are certainly sensitive to the ambiguity when it is first encountered. This result in fragment 1 may simply mean that it is more difficult to represent two characters of the same gender than it is to represent two characters of different gender rather than having something to do with difficulty in interpreting an ambiguous pronoun.

Pronominal ambiguity did not interact with any of the other factors. The fact that we did not find that the implicit causality congruency effect interacted with ambiguity indicates initial assignment was not occurring on the pronoun using gender information. If gender information was used immediately and the reader did attempt to identify the pronoun's antecedent, we would have expected a weaker implicit causality congruency effect when gender cue was present than when it was absent. Its presence should facilitate integration of the information contained within the second fragment as it could be used to overcome the congruency effect to some degree. The lack of an interaction between congruency and ambiguity suggests that either readers don't try and resolve a pronoun when it is encountered or they do try to resolve the pronoun but they don't use gender information to do so. The first explanation is consistent with the position we have adopted on the basis of the data reported so far.

We also found an implicit causality congruency effect of the same type as we have reported previously. Again our data reveal that integration following reference to the first mentioned character is always easy, integration following reference to the second is only easy when supported by verb bias.

With respect to our initial question of at what points within the anaphor resolution system do gender information and implicit causality information exert a processing influence, what can we conclude ? The lack of an interaction between verb bias and any other factor in our analysis of reading times to fragment 1 further supports the position that implicit causality information exerts the bulk of its influence at the integrative stage of processing. In order to further explore the time course of the effects we have reported, in our next experiment we adopt an eye-tracking methodology in an attempt to provide us with a more spatially and temporally fine grained measure of the processes associated with resolving pronouns in the context of both implicit causality and gender cues.

5.4 Experiment 8

5.4.1 Introduction and Rationale

Up to this point all our on-line examinations of processing associated with resolving anaphoric reference have employed a self-paced reading methodology. For this experiment we adopt an eye-tracking paradigm which allows us to obtain a fine grained measure of readers' eye-movements as they read the experimental materials. Two assumptions are made in order to interpret eye movement data.

During normal reading, readers can control the rate and manner in which they are presented with information through eye movements. The first assumption, the immediacy assumption, proposes that when a word is fixated by the reader, it is interpreted to the deepest level possible (Just and Carpenter, 1980) (see section 2.3.1 above). This is related to a general incrementality of processing viewpoint which construes the goal of the language system as one of achieving a message level interpretation of the input as soon as constraints allow (Marslen-Wilson and Tyler, 1981). Certainly however, on occasion full processing of a word which is ambiguous to some degree must be delayed until further disambiguating information becomes available.

The second assumption of Just and Carpenter, the eye-mind assumption, proposes that the eye remains fixated on a word until processing of that word terminates.

The degree of this eye-mind coupling was examined by Ehrlich and Rayner (1983). Rayner (1977, 1978) examined a position proposing that the relationship between eye-movements and cognitive processing should best be characterised as possessing a cognitive lag. Eye-movements effectively brought information into a buffer which was then operated on by the processing system. This was contrasted with the immediacy hypothesis which denies the existence of such a lag and instead claims that eye fixation durations are longer when fixating on words for which more processing was required. The immediacy hypothesis proposes that not just low level processing associated with lexical access is occurring while the eye is fixating a portion of text, but also higher level processing.

Ehrlich and Rayner (1983) chose to focus on the question of anaphoric pronoun resolution, a relatively high level of cognitive processing requiring the reader to relate the currently fixated pronoun with its antecedent in some earlier point in the text (see Chapter 2). Degree of difficulty was manipulated by varying the distance between pronoun and antecedent. In their texts, the pronoun was gender marked and so had only one antecedent. In other words, the pronoun's antecedent could potentially be identified when the pronoun was encountered. Contrary to the immediacy hypothesis, Ehrlich and Rayner (1983) found that although there was an increase in reading time for pronouns whose antecedent was distant, this was manifested as an increase for the region following the one in which the pronoun was contained. No evidence was found for an increase in time spent on the region containing the pronoun as a function of antecedent distance. In other words, although antecedent identification could be accomplished during fixation of the pronoun, this did not happen. The initiation of the processes associated with resolving reference may occur as the signal to do so is encountered, but in cases where this is not a trivial process and the antecedent must be recovered from some distant point, the completion of processing occurs while the eye is fixating some subsequent region. Reanalysis of the data reported in

Ehrlich (1983) by Ehrlich and Rayner (1983) supports this position. The increase in time to resolve the reference of pronouns with distant antecedents can be interpreted as a result of the antecedent falling out of the readers' discourse focus model. This was the argument taken by Sanford and Garrod (1981) who claimed that distant antecedents are less accessible as they no longer belong to the topic of the current segment of text. It is also consistent with our claim that under the circumstances we have so far examined, co-indexation of a pronoun does not occur as soon as it is encountered.

For our purposes we want to try to distinguish the different points in time during which implicit causality and gender information are used by the reader. One way in which this can be achieved through analysis of eye-movement data is by separately focusing on the amount of time the eye first spends in a defined region of text and the total amount of time the eye spends in that particular region.

This experiment closely parallels Experiment 6 with several notable exceptions. An extra region was added at the end of the experimental sentences as a buffer to accommodate end of sentence wrap up effects (Just and Carpenter, 1980).

(8) John blamed Bill because he broke the window, the day before yesterday.

The second modification was that questions followed a third of the experimental trials. An example question is 'Did John break a window?'. The questions could only be correctly answered if the pronoun has been successfully resolved.

Method

Subjects

Forty English speaking subjects participated.

Stimuli

We manipulated pronominal ambiguity, verb bias and ending. The pronoun could either be referentially ambiguous or marked for gender to unambiguously refer to one of the two participants. Verb bias could either be NP1 or NP2 biasing. Referent could either be the first or second character.

There were 24 sets of experimental materials (see appendix 4 for full set).

Unambiguous pronoun / NP1 verb / Referent character1

Barry fascinated Mary because he performed magic tricks, so I heard.

Unambiguous pronoun / NP1 verb / Referent character2

Mary fascinated Derek because he was easily entertained, so I heard.

Ambiguous Pronoun / NP1 verb / Referent character1

Barry fascinated Derek because he performed magic tricks, so I heard.

Ambiguous Pronoun / NP1 verb / Referent character2

Barry fascinated Derek because he was easily entertained, so I heard.

There were also 111 filler items. The 24 verbs used to construct the experimental materials were selected from the set of verbs examined in the off-line studies (see list above).

Apparatus

Eye movements were monitored by a Stanford Research Institute Dual Purkinje Generation 5.5 Eye-tracker made by Forward Technologies, California under license to the S.R.I. The eye-tracker has an angular resolution of 10 arc. Viewing was binocular with eye location recorded from the right eye. The passages were presented on a VDU interfaced with a PC which controlled the experiment. The VDU was located at a distance of 70 cms. The position of the subject's eye was

sampled every millisecond and analysed using software developed by Chuck Clifton at Umass which continuously monitors the output to establish the sequence of eye-fixations and their start and finish times to the nearest millisecond.

Procedure

Each subject was first seated comfortably at the eye-tracker with their head held still by a chin and forehead restraint. A bite bar further was also employed in order to reduce head movement. The system was locked onto the Purkinje images from the right eye and the subject was taken through a short calibration procedure. When calibration was completed satisfactorily, the experiment began.

Each subject participated in 2 practice trials similar in structure to the experimental items at the start of each experimental block. There were three such blocks. The experiment lasted between 45 minutes and 1 hour. Before the experiment subjects were provided with both verbal and written instructions.

Regions

We divided each sentence into 5 regions for purposes of analysis. Region 1 contained the first noun phrase, the verb and the second noun phrase. Region 2 contained the connective 'because'. Region 3 contained the subordinate clause. Region 4 contained the rest of the sentence up to a line break while Region 5 contained the final word(s) of the sentence on the next line.

Barry fascinated Mary/ because/ he performed magic tricks,/ so I/
heard.

Analyses

An error occurred with one of our materials. Reading times for that material were excluded from our analyses.

An automatic procedure pooled short contiguous fixations. The procedure incorporated fixations of less than 80 msec into larger fixations within one character and then deleted fixations of less than 40 msec that fell within three characters of any other region. Before analysing the data, trials where the subject failed to read the sentence or when tracker loss ensued were removed.

First-Pass Reading Time is the sum of the fixations occurring within a region before the first fixation outside the region. If the eye fixates on a point beyond the end of a region before landing in the region for the first time then the first-pass time for that region is zero. Total Reading Time is the sum of all fixations in a region.

The analyses reported below exclude 0-msec times that occurred when readers skipped a region.

5.4.2 Results _

First-Pass Reading Times

First-pass reading time corresponds to the amount of time the eye first spends in a designated region of text.

For Region 1, e.g. 'John blamed Bill', we performed 2 (same gender vs different gender) x 2 (NP1 verb vs NP2 verb) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias significant by subjects only ($F(1,39)=8.84$, $p<0.005$; $F(2,78)<2$) corresponding to faster reading times associated with reading NP2 biasing verbs.

Table 7 - First-pass reading times for region 1, e.g. 'John blamed Bill', subject means for each condition will all times in msec.

Verb bias	Characters of same gender	Characters of different gender
NP1	964	950
NP2	898	889

For Region 3, e.g. 'he broke the window', we performed 2 (same gender vs different gender) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found one interaction of verb bias x referent that was significant only by subjects ($F(1,39)=11.70$, $p<0.005$; $F(1,21)<1$) suggesting the implicit causality congruency effect.

Table 8 - First-pass reading times for region 3, e.g. 'he performed magic tricks', subject means for each condition with all times in msec.


Verb / Referent	Characters of same gender	Characters of different gender
NP1 / Char1	981	982
NP1 / Char2	1040	1004
NP2 / Char1	1030	1042
NP2 / Char2	932	1022

For Region 3, e.g. 'he broke the window', we examined the conditions corresponding to the characters being of the same or different gender separately.

Where characters are of the same gender

For Region 3, e.g. 'he broke the window', we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found one interaction of verb bias x referent that was significant only by subjects ($F(1,39)=8.69$, $p<0.01$; $F(1,21)=1.83$, $p<0.19$) (see Figure 5).

Figure 5 - verb bias x referent



Where characters are of different gender

For the region 'he broke the window', we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. No main effects or interactions approached significance.

Total Reading Time

Total reading time corresponds to the total amount of time the eye spends in a designated region of text.

For Region 1, e.g. 'John blamed Bill', we performed 2 (same gender vs different gender) x 2 (NP1 verb vs NP2 verb) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias significant by subjects only ($F(1,39)=11.71, p<0.005$; $F(1,21)=3.29, p<0.1$) corresponding to faster reading times associated with reading NP2 biasing verbs.

Table 9 - Total reading times for region 1, e.g. 'John blamed Bill', subject means for each condition will all times in msec.


Verb bias	Characters of same gender	Characters of different gender
NP1	1475	1413
NP2	1329	1295

For Region 3, e.g. 'he broke the window', we performed 2 (same gender vs different gender) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias significant by subjects only ($F(1,39)=4.89$, $p<0.05$; $F_2<1$) corresponding to subordinate clauses following NP2 biasing verbs being read more quickly than following NP1 biasing verbs. We also found a main effect of referent significant by subjects only ($F(1,39)=6.60$, $p<0.05$; $F_2<2$) corresponding to a subordinate clause containing reference to the first mentioned character being read more quickly than one containing reference to the second mentioned character. We found a verb bias x referent interaction significant by subjects only ($F(1,39)=13.22$, $p<0.001$; $F_2(1,21)=2.59$, $p<0.13$) (see Figure 6). The gender x referent interaction approached significance by subjects only ($F(1,39)=3.81$, $p<0.06$; $F_2(1,21)=2.30$, $p<0.15$).

Table 10 - Total Reading times for region 3 'he performed magic tricks', subject means for each condition with all times in msec.

Verb / Referent	Characters of same gender	Characters of different gender
NP1 / Char1	1207	1214
NP1 / Char2	1366	1464
NP2 / Char1	1312	1250
NP2 / char2	1201	1266

Figure 6 - verb bias x referent




Again we examined the conditions corresponding to the characters being of the same or different gender separately.

Where characters are of the same gender

For Region 3, e.g. 'he broke the window', we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a verb bias x referent interaction significant by subjects only ($F(1,39)=10.50, p<0.005$; $F(1,21)=2.38, p<0.14$) (see Figure 7).


Figure 7 - verb bias x referent



Where characters are of different gender

For Region 3, e.g. 'he broke the window', we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of referent significant by subjects but marginal by items ($F(1,39)=8.91$, $p<0.005$; $F(1,21)=3.15$, $p<0.1$) corresponding to a subordinate clause containing reference to the first mentioned character being read more quickly than one containing reference to the second mentioned character. We also found a verb bias x referent interaction significant by subjects only ($F(1,39)=5.58$, $p<0.05$; $F(1,21)=2.07$, $p<0.17$) (see Figure 8).

Figure 8 - verb bias x referent



5.4.3 Discussion

With respect to our Total Reading Time analyses, although we rarely found significant effects on an F2 analysis, the general pattern of significant effects found on an F1 analysis is consistent with what we have reported previously. In both our first-pass reading time analysis and our total time analysis for our first region we found evidence that NP2 biasing verbs are read more quickly than NP1 biasing verbs. This replicates what we found in our self-paced reading studies. For our third region, i.e. the subordinate clause, we found main effects of verb bias and ending. These correspond to subjects reading subordinate clauses more quickly following an NP2 biasing verb. They also read the clause more quickly if that clause contains reference to the first mentioned character in the main clause. We have already speculated as to why these effects might arise and we shall return to their interpretation in Chapter 8 which provides a more detailed interpretation of the set of experiments reported in this thesis when taken as a whole.

As our analyses treating items as random factors failed to produce significant results, interpretation of the First Pass Reading Time data

is equivocal. However, there does seem to be weak evidence to suggest that the presence or absence of a gender cue determines whether or not implicit causality information will be used by the system immediately during processing of the subordinate clause. If gender information is sufficient to identify the pronominal referent, the influence of implicit causality information is delayed slightly. If gender information isn't available to identify the pronominal referent, implicit causality information appears to exert an influence more quickly.

A caveat should be raised however as only rarely do our F2 analyses support our findings. There are two main reasons why we had difficulty finding effects significant by items. The first is simply that we examined half the number of experimental materials that we normally examine. This arose as our experiment was run alongside a separate study thus restricting the number of items we could present to subjects. One error occurred with one of our materials reducing the number of degrees of freedom associated with the F2 analysis from 46 in our self-paced reading studies to 21 in this current study. Another difference between the design of this eye-tracking study and our self-paced reading studies, although minor, may have contributed noise to our data. Recall we added a final region to the end of our materials so that any sentence wrap-up effects occurred after our final region of interest, see example (9). Although the intention was for these additional phrases to be relatively semantically empty this may not have always been the case and information that was contained within the phrase may have influenced how the reader ultimately interpreted the preceding main and subordinate clauses thus possibly clouding our total time data.

(9) John blamed Bill because he broke the window, the day before yesterday.

On occasion the final phrase may have affected the overall plausibility of the sentence. We did not control for this so it is not impossible that the plausibility of our materials was influenced in small but important ways thus adding noise to the eye-movement data we obtained.

Given these potential problems however, the pattern of eye-movement data we obtained is consistent with the pattern of reading time data we have previously obtained in our self-paced reading experiments. We find evidence of a processing influence of implicit causality and gender information during integration of the information contained within the subordinate clause. Although the pattern of reading time data is not clear, there does seem to be the suggestion that when a gender cue differentiating the two participants in the main clause is present, implicit causality information is not used immediately as the subordinate clause is read. Recall we found some evidence that implicit causality exerted a processing influence when we examined total reading time of the subordinate clause but no evidence that it exerted an influence when we examined first pass reading times. This was only true in the case where a gender cue was present. When this cue was absent, i.e. where the participants in the main clause were of the same gender, we did find some weak evidence of a processing influence of implicit causality information on first pass reading times as well as on total reading time of the subordinate clause.

5.5 General Discussion

The data gathered in Experiment 5, where subjects were presented with two separate sentences, suggests that when the manner of presentation of information is manipulated, the implicit causality congruency effect remains but only under conditions where the two separate utterances are linked by a pronominal anaphor. We claimed above that this was due to the reader not integrating the information contained within the two sentences when the referential link connecting them was a repeat name anaphor. We argued this on the basis of the position outlined by Vonk et al suggesting that over-specific anaphors were interpreted by the reader as signals of shifts in theme and therefore as an explicit cue **not** to integrate the information predicated by the repeat name anaphor with what had been read previously.

With respect to the time course of the processing influence of implicit causality information, the pattern of data reported in Experiments 6 and 7 in this chapter is consistent with our previously outlined position that implicit causality exerts an influence at the integrative stage of processing, i.e. the second stage within the model of anaphor resolution that we proposed earlier in this thesis.

The pattern of data we found in Experiment 6 suggests that gender information is not used as soon as a pronoun is encountered. The marginal main effect of ambiguity that we found on fragment 1 reading times suggests that the reader is sensitive to the ambiguity but as we proposed above this may be due to the ease in which two characters can be represented depending on whether they are of the same or different gender. We propose that it is the case that there is a delay between a pronoun being read and its antecedent being identified rather than the case that initial identification takes place when the pronoun is encountered but without reference to gender information. Sanford and Garrod (1989) summarised evidence also suggesting that pronouns aren't interpreted as soon as they are encountered. Our position is compatible with this.

We base our conclusion on the lack of evidence for an ambiguity x implicit causality congruency interaction. If gender cue is used by the system as soon as a pronoun is encountered, it should be easier to integrate information in the subordinate clause which goes against the bias of the verb when gender cue is present. We found no evidence of this. The gender cue should allow the system to identify a pronoun's antecedent as soon as the pronoun is encountered. The first stage of co-indexation would therefore be finished before the rest of the information in the subordinate clause is read. The only work remaining to be done during reading of this information is integration. In the case where gender cue is absent, both co-indexation and integration must be carried out while the information predicated by the pronoun in the subordinate clause is read and therefore there should be greater processing difficulty associated with reading a subordinate clause going against the bias of the verb. We found no interaction between ambiguity and the implicit causality

congruency effect suggesting that gender information is not used as soon as a pronoun is encountered to identify its antecedent.

Pronouns are not resolved immediately, even under circumstances where information sufficient to do so is present. This seems counter-intuitive but is compatible with a system that, for the majority of cases, requires resolution of a pronoun to be delayed until information predicated by that pronoun is read. If delay is beneficial to the system on the majority of occasions (and we propose that avoiding misassignment is beneficial) it would seem optimal for the system to always delay pronoun resolution. Only on a small number of occasions will this lead to sub-optimal behaviour but generally this behaviour will approach optimality.

The data from our eye-tracking study (Experiment 7) support this position. We do have some suggestion in the data however that when the subordinate clause is first read, the presence or absence of a gender cue determines the manner in which implicit causality information will be used by the system. If gender cue is absent, implicit causality information will be used immediately. If it is present, implicit causality information will not be used immediately. Recall though that our results in this study were rarely supported by significant F2 analyses. Although our Experiment 7 data allow us to do little more than speculate, it is possible that first pass reading times are giving us a window onto the level of co-indexation of pronominal anaphora, while the total time measure captures both this and integrative processing.

We shall return to interpretation of the studies reported in this chapter and all those in this thesis in Chapter 8.

CHAPTER 6

6.1 Introduction and Rationale

In Chapter 3 we discussed two sets of experiments reported by McDonald and MacWhinney (1995) and by Garnham, Traxler, Oakhill and Gernsbacher (1996). The data described by McDonald and MacWhinney is claimed to be supportive the position that implicit causality information exerts an early influence within the anaphor resolution system, specifically at the point at which a pronoun is encountered. This stands in contradiction to the data described by Garnham et al which suggests that implicit causality exerts a relatively late influence within the anaphor resolution system, the bulk of its influence arising during the process of integration. All the experiments reported in this thesis up to this point support the position outlined by Garnham et al. This chapter is an attempt at reconciling the positions of McDonald and MacWhinney and Garnham et al.

If it is correct, McDonald and MacWhinney's claim for non-structural factors influencing processing at an early temporal point is important. It suggests that non-structural information is available within the system to guide processing at an early stage. If this is at an initial stage of processing, it follows that the anaphor resolution system must be behaving in a nonrestricted manner. However, the McDonald and MacWhinney data stand alone within the literature on the claim that such higher level factors exert an influence early. The experiments reported in this thesis suggest that implicit causality information exerts an effect at a later, integrative stage of sentence processing. This is based on data gathered using a reading time measure. As McDonald and MacWhinney use a probe task rather than reading time as their measure of the influence of implicit causality, it may be the case that the probe task is simply a more sensitive measure for this type of effect. Because of this possible difference in sensitivity of the measure, it is not easy to directly compare the McDonald and MacWhinney data with data generated by a self-paced reading measure. However, as the probe task was also used by Garnham et al, we can use their data in a more direct comparison.

Although there is much in common between the experimental techniques employed by McDonald and MacWhinney and by Garnham et al, a number of differences remain.

The studies of McDonald and MacWhinney and Garnham et al differ in four important ways :

- 1. The nature of the probe task used.**
- 2. The thematic structures of the verbs used.**
- 3. Additional information about the causal relationship between the characters that is present in the McDonald and MacWhinney materials but absent in those examined by Garnham et al.**
- 4. The experimental conditions examined within each set of studies.**

1. The nature of the probe task used.

The particular nature of the probe task differs between the two sets of studies. In the case of McDonald and MacWhinney, subjects heard sentences and had to respond to a visually presented probe word. The probe was the repetition of the name of one of the participants mentioned in the sentence. Recall using a cross-modal naming task, Marslen-Wilson, Tyler and Koster (1993) found evidence for an early influence of high level pragmatic factors on pronoun resolution. In light of the Marslen-Wilson et al data, it is possible that cross modal tasks generally detect early influences of non-structural factors on pronoun resolution although this may arise as a consequence of the nature of the task itself and may reveal little about normal reading processes. Additionally there may simply be less processing load associated with hearing a sentence than there is with reading one.

The probe task employed by Garnham et al was not cross-modal. In the case of the set of studies reported by Garnham et al, subjects were presented with the sentences visually, word by word. The probes

were again presented visually and subjects had to make a recognition response. As the presentation rate associated with the individual words presented by Garnham et al was slower than that which would accompany normal reading, it would seem plausible to suggest that the extra time provided for the subjects to read each word would have given ample opportunity for implicit causality information to exert a processing influence. Garnham et al found no evidence for an early effect of this information despite these circumstances providing the most likely ones under which an early effect would be found. Artificially slow reading times would be expected to **increase** the possibility of implicit causality information exerting an influence at an early processing point. We reject point (1) as an explanation for the differences found in the time course of the effect.

2. The thematic structures of the verbs used.

McDonald and MacWhinney restricted their set of verbs exhibiting an implicit causality bias to those which can be classed as 'state' verbs. The materials of Garnham et al contained no such restriction and their set of verbs consisted of both state and action verbs. Recall, state verbs describe internal experiential states, for instance in the sentence 'John fascinated Bill because he...', 'fascinated' leads to the interpretation that some action or quality on the part of the character 'John' has lead to 'Bill' being in a state of fascination. 'John' occupies the thematic stimulus role and 'Bill' the thematic experiencer role. The verb 'fascinated' is an NP1 biasing verb. Similarly, state verbs such as 'admired' in a sentence like 'John admired Bill because he...' lead to an interpretation that 'John's' internal experience of admiration is caused by some quality of 'Bill'. 'John' occupies the thematic experiencer role and 'Bill' the thematic stimulus role. 'Admire' is classed as an NP2 biasing verb. Action verbs, such as 'kicked', describe actions occurring between people and possess associated agent and patient thematic roles. Although verbs of this sort were excluded from the set examined by McDonald and MacWhinney, they form a subset of those examined by Garnham et al.

As mentioned in section 2.7.1 above, Stevenson, Crawley and Kleinman (1994) tried to reduce the explanation needed to account

for the implicit causality bias to the level of thematic roles. In the case of state verbs, Stevenson et al claim that the locus of the ultimate cause for the described event lies with the character occupying the 'stimulus' role. Similarly for action verbs, Stevenson et al claim that the ultimate cause is interpreted as lying with the character occupying the patient role. Indeed, there is a large degree of correlation between verb bias and associated thematic structure but it appears to be no more than that. There are certainly some counter-examples, such as 'telephone' and 'punish'. Although their thematic structures are identical, the first is an NP1 biasing verb while the second NP2. In the first case, the character occupying the thematic agent role is interpreted as the causer, while in the second it is the character occupying the patient role. Also, the explanation proposed by Stevenson et al may claim to be able to account for the biases per se, but seems unable to explain the relative strengths of the biases associated with individual verbs. Clearly this is determined by some other factor and it is this which probably best accounts for the bias effect in general.

In light of this evidence that thematic structure does nothing more than correlate with verb bias, that the verbs examined by McDonald and MacWhinney and by Garnham et al differ at the level of thematic structure should not be taken as a serious candidate for explaining the contradiction in their findings.

3. Additional information about the causal relationship between the characters that is present in the McDonald and MacWhinney materials but absent in those examined by Garnham et al.

One factor common to both the experimental materials employed by McDonald and MacWhinney and by Garnham et al and not present in materials examined in previous research or in the experiments so far reported in this thesis is the presence of an additional phrase preceding the connective. In the case of those sentences examined by Garnham et al, this phrase is predominantly either a temporal or a locative prepositional phrase, (example (1) below).

(1) Sandra called Elaine before breakfast because she ...

The additional prepositional phrase provides no extra information relevant to what may have caused the described event. Conversely, the additional phrase present in the materials examined by McDonald and MacWhinney is, in the majority of cases, an adverbial phrase, see example (2) below.

(2) Beth disappointed Pam repeatedly because she ...

In contrast to the function of the additional phrase in the Garnham et al materials, we argue that this additional adverb **does** provide further information about the nature of the relationship between the two mentioned participants and can be used by the reader to form the basis of an inference which can facilitate comprehension of the underlying cause. Compare Example (2) with Example (3) below.

(3) Beth disappointed Pam because she ...

In the case of Example (2) we can ask ourselves whether the fact that the event of 'Beth' disappointing 'Pam' occurs repeatedly is likely to be due to some quality on 'Beth's' part or due to some quality on 'Pam's' part. The verb 'disappoint' is an NP1 biasing verb so without this additional phrase (Example (3)) the causal locus is attributed to 'Beth'. If the phrase places further emphasis on some quality of 'Beth's' as the reason behind the described event, this will result in an increase in the attribution to the character 'Beth' as causally responsible.

In conjunction with the implicit causality information, the extra adverbial information may strengthen the causal bias toward one or other of the participants. In other words, the adverb might also bias toward one of the two characters in addition to the verb bias. This may result in an increase in the absolute level of the bias as a product of the interaction between the verb and the adverb. Whether the result of this should lead to an increase in the overall activation differential between the two participants or whether it should lead to an influence of the bias at an earlier point in time is unclear. As McDonald and MacWhinney found an early influence of implicit

causality information in sentences of this type, it is possible that the large differential is actually leading to implicit causality influencing processing at an early point.

It may be the case that the additional adverbial also serves to strengthen the overall congruency effect. We examine both of these claims in Experiment 8 in this chapter.

It is possible that the additional information preceding the connective may provide extra time for the verb bias to exert a processing influence. If that was the case, however, Garnham et al should also have found evidence for an early influence of implicit causality on pronoun resolution.

4. The experimental conditions examined within each set of studies.

We feel that the most likely candidate for reconciling the inconsistent data presented by McDonald and MacWhinney and by Garnham et al lies at the level of an experimental confound. To recap, consider the standard procedure used to measure the implicit causality congruency effect: There will be a processing difficulty associated with reading a subordinate clause inconsistent with the direction of bias of the verb contained in the preceding main clause. In other words, an NP1 continuation will be read quickly following an NP1 biasing verb (Example (4)) but more slowly following an NP2 biasing verb (Example (5) below). Similarly, an NP2 continuation will be read quickly following an NP2 biasing verb (Example (6)) but more slowly following an NP1 biasing verb (Example (7) below).

(4) John fascinated Bill because he was interesting.

(5) John blamed Bill because he was in a bad mood.

(6) John blamed Bill because he broke the window.

(7) John fascinated Bill because he was easily entertained.

The subordinate clause in Examples (4) and (6) is consistent with the direction of verb bias. In Example (5) and (7) however the subordinate clause is inconsistent with the verb bias. Although the majority of studies examining the nature of the implicit causality effect (including the Garnham et al study) contain all 4 conditions (2 consistent with bias and 2 inconsistent with bias), McDonald and MacWhinney's studies contained only the consistent conditions. The lack of counter-balance leads to a possible confound. The effect of this is that subjects are potentially able to predict to which character the pronoun following the connective refers before that pronoun is encountered. Recall Au (1986) demonstrated that subjects display sensitivity to verb bias when asked to sort verbs into categories. Because subjects are aware of the verb bias it follows that upon encountering a verb, they should be inclined to expect one type of continuation over another. As subjects are never presented with counter-examples (i.e. inconsistent conditions), reinforcement of this predictive behaviour will occur during the experiment. Subjects will effectively 'know' that in the context of the experiment, an NP1 biasing verb will always be followed by a pronoun consistent with the bias and likewise for an NP2 biasing verb. McDonald and MacWhinney tried to counter this argument by demonstrating a lack of influence of verb bias in sentences of the sort :

(8) John amazed Bill time after time at the juggling competition.

They found no effect of anticipatory priming at any probe point (i.e. no implicit causality influence on antecedent activation levels) and concluded that it was the process of anaphor resolution itself which stimulated the effect. This fails to adequately challenge the priming argument however. Ehrlich (1980) demonstrated the necessity of the presence of the overt causal connective 'because' for the implicit causality information to exert a processing influence. As this necessary condition is lacking in McDonald and MacWhinney's second experiment, that priming due to an experimental confound gives rise to the early effect cannot be satisfactorily dismissed.

We believe that this experimental confound is responsible for the discrepancy between the McDonald and MacWhinney and Garnham et

al positions. It is very likely that subjects engaged in predictive processing as they read McDonald and MacWhinney's materials. Therefore their finding of an early effect of implicit causality information should be interpreted as evidence for nothing more than subjects' sensitivity to the verb bias and ability to predict the nature of subsequent information as a result of the experimental environment. The confound and lack of counter-examples in their materials leads to anticipatory priming.

6.2 Experiment 8

6.2.1 Introduction and Rationale

Experiment 8 following is an attempt to replicate the data reported by McDonald and MacWhinney but using a reading time measure rather than the probe task and with both consistent and inconsistent implicit causality conditions. We investigate our suggestion that the extra adverbial information strengthens the influence of implicit causality information and increases the activation differential between the two participants. On the basis of this proposal we either expect to find a stronger implicit causality effect than we have found previously, or the influence of the bias within the anaphor resolution system arising at an earlier point. Our study differs in a number of important ways from McDonald and MacWhinney's, not least in terms of the presence of the equivalent counter-examples corresponding to materials containing incongruent sentence continuations.

Method

Subjects

Thirty two English speaking subjects participated.

Stimuli

We manipulated anaphor, verb bias and ending. Anaphor could either be pronoun or repeat name. Verb bias could either be NP1 or NP2 biasing. Ending could either be to the first or second NP.

We examined the same experimental materials as were used in Experiment 3. The only difference was that in this study we added an adverb following the second Noun Phrase in the main clause. These adverbs were comparable to those present in the materials of McDonald and MacWhinney. There were 48 sets of experimental materials (see appendix 5 for full set).

Pronoun / NP2 verb / Referent character2

John scolded Bill severely because he / had damaged the mahogany table.

Pronoun / NP2 verb / Referent character1

John scolded Bill severely because he / was aware of the potential danger.

Name / NP2 verb / Referent character2

John scolded Bill severely because Bill / had damaged the mahogany table.

Name / NP2 verb / Referent character1

John scolded Bill severely because John / was aware of the potential danger.

As in Experiment 3, there were also 96 filler items. The experiment was divided into two halves. The first occurrence of the verb was in the first half of the study and the second in the second half.

Procedure

Experiment 11 was run on an Apple Macintosh computer using the Psyscope experimental software. A button box was connected to the

computer which recorded subjects responses with millisecond accuracy.

We employed the basic experimental method reported for Experiment 3. The experiment lasted roughly 35 minutes.

6.2.2 Results

We report analyses for the two fragments of the experimental sentences themselves and also for the question responses following those fragments.

Outlier Replacement

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For reading times to the first fragment we excluded response times that were below 300 msec or above 16 seconds. This accounted for 1.2% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.1% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the data corresponding to reading times to the second fragment we excluded response times that were below 300 msec or above 11 seconds. This accounted for 1.4% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.8% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the question response data we excluded response times that were below 300 msec or above 13 seconds. This accounted for 0.5% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 3.3% of the data was replaced in this way.

First Fragment Analysis

We performed 2 (Pronoun anaphor vs Repeat Name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type ($F(1,31)=9.42$, $p<0.005$; $F(1,46)=4.74$, $p<0.05$) corresponding to fragments containing pronouns being read more quickly than those containing repeat names. We also found a main effect of referent significant by subjects only ($F(1,31)=6.39$, $p<0.05$; $F(1,46)=3.28$, $p<0.08$) suggesting faster reading time following reference to the second mentioned character. As fragment 1 is identical for the Pronoun conditions, we looked at the Name conditions separately. The main effect of referent was found corresponding to the repeat name penalty although it was marginal by items ($F(1,31)=8.82$, $p<0.01$; $F(1,46)=3.28$, $p<0.08$).

Table 1 - Reading times for first fragment, subject means for each condition with all times in msecs.

Verb / Referent	Name	Pronoun
NP1 / Char1	3711	3183
NP1 / Char2	3396	3354
NP2 / Char1	3550	3457
NP2 / Char2	3287	3201


Second Fragment Analysis

We performed 2 (Pronoun anaphor vs Repeat Name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type ($F(1,31)=19.99$, $p<0.0001$; $F(1,46)=25.76$, $p<0.0001$) corresponding to faster reading times following reference achieved through the use of a repeat name anaphor. We also found an interaction between verb bias and referent corresponding to the implicit causality congruency effect ($F(1,31)=7.09$, $p<0.05$; $F(1,46)=4.87$, $p<0.05$). (see Figure 1 below).

Table 2 - Reading times for second fragment, subject means for each condition with all times in msec.

Verb / Referent	Name	Pronoun
NP1 / Char1	1555	1825
NP1 / Char2	1633	1991
NP2 / Char1	1696	1931
NP2 / Char2	1620	1803

Figure 1 -verb bias x referent interaction




Although we did not find an interaction between the congruency effect and anaphor type, we nevertheless examined the Pronoun and Name conditions separately.

Pronoun Conditions

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found an interaction between verb bias and referent corresponding to the congruency effect ($F(1,31)=4.32, p<0.05$; $F(1,46)=4.46, p<0.05$). See Figure 2 below.

Figure 2 - verb bias x referent interaction (Pronoun conditions)



Name Conditions

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a marginal main effect of verb bias that was not significant by items ($F(1,31)=3.20$, $p<0.09$; $F_2<1$). We found a hint of an interaction between verb bias and referent that was marginal by subjects and not significant by items ($F(1,31)=3.13$, $p<0.09$; $F_2<2$).

Question Response Time Analysis

We performed 2 (Pronoun anaphor vs Repeat Name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type ($F(1,31)=50.17$, $p<0.001$; $F_2(1,46)=125.93$, $p<0.0001$) corresponding to faster response times following sentences containing repeat name anaphors. We also found a main effect of verb bias ($F(1,31)=12.66$, $p<0.005$; $F_2(1,46)=6.20$, $p<0.05$) corresponding to faster response times following sentences containing NP2 biasing verbs. We also found an interaction between verb bias and referent ($F(1,31)=8.45$, $p<0.01$; $F_2(1,46)=9.66$, $p<0.005$) corresponding to the implicit causality congruency effect. We found a 3-way interaction of

anaphor type x verb bias x referent ($F(1,31)=6.90$, $p<0.05$; $F(1,46)=4.24$, $p<0.05$) corresponding to the presence of the implicit causality congruency effect in the Pronoun conditions ($F(1,31)=10.28$, $p<0.005$; $F(1,46)=11.62$, $p<0.005$) but not in the Name conditions (both $F_s<1$). We also found a marginal interaction between anaphor type and verb bias ($F(1,31)=3.55$; $F(1,46)=2.87$, $p<0.1$).

Table 3 - Question response times, subject means for each condition with all times in msecs. Numbers in parentheses correspond to response accuracy expressed as percentage correct.

Verb / Referent	Name	Pronoun
NP1 / Char1	2026 (91.0%)	2666 (75.6%)
NP1 / Char2	2009 (92.1%)	3132 (51.1%)
NP2 / Char1	2011 (93.8%)	2740 (78.6%)
NP2 / Char2	1889 (95.8%)	2460 (89.6%)

Question Response Accuracy Analysis

We performed 2 (Pronoun anaphor vs Repeat name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type ($F(1,31)=96.86$, $p<0.0001$; $F(1,46)=80.24$, $p<0.0001$) corresponding to greater accuracy following sentences contain repeat name anaphors. We also found a main effect of verb bias ($F(1,31)=47.79$, $p<0.001$; $F(1,46)=38.95$, $p<0.0001$) corresponding to greater response accuracy following NP2 verbs. We found interactions between anaphor type and verb bias ($F(1,31)=35.61$, $p<0.0001$; $F(1,46)=15.53$, $p<0.0005$), between anaphor type and referent significant by subjects only ($F(1,31)=3.98$, $p<0.06$; $F(1,46)=2.12$, $p<0.16$), verb bias and referent ($F(1,31)=23.50$, $p<0.0001$; $F(1,46)=17.46$, $p<0.0001$) and a 3-way interaction between verb bias, referent and anaphor type ($F(1,31)=19.33$, $p<0.0001$; $F(1,46)=15.05$, $p<0.0005$) corresponding to the presence of the implicit causality congruency effect in the Pronoun conditions ($F(1,31)=37.34$, $p<0.0001$; $F(1,46)=20.14$, $p<0.0001$) compared to the Name conditions (both $F_s<1$)

6.2.3 Comparing Experiments 3 and 8

We also carried out a between experiment analysis to examine what effect the presence of the additional phrase in Experiment 8 might have had on the nature of the repeat name penalty and the implicit causality congruency effect. We compared fragment 1 and fragment 2 reading times for Experiment 8 with the corresponding reading times for Experiment 3.

For fragment 1, to examine how the repeat name penalty might vary between studies we performed 2 (Char1 vs Char2 referent) x 2 (Experiment 3 vs Experiment 8) ANOVAs for both subjects and items as random factors for the Name conditions. We did not find an interaction between referent and Experiment (both $F_s < 1$) indicating that the magnitude of the repeat name penalty did not vary between Experiments 3 and 8. Experiment did not interact with any factor (all $F_s < 1$).

For fragment 2, to examine how the congruency effect might vary between studies we performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) x 2 (Experiment 3 vs Experiment 8) ANOVAs for both subjects and items as random factors. We found an interaction between the congruency effect (verb bias x referent) and experiment ($F_1(1,62)=10.73$, $p < 0.005$; $F_2(1,92)=4.80$, $p < 0.05$). This corresponds to the implicit causality congruency being stronger in Experiment 3 than in Experiment 8. We also found an interaction of anaphor type x referent x experiment ($F_1(1,62)=6.84$, $p < 0.05$; $F_2(1,92)=5.04$, $p < 0.05$). Experiment did not interact with any factor (all $F_s < 1$).

6.2.4 Discussion

Experiment 8 failed to replicate McDonald and MacWhinney's findings. It supports the Garnham et al position that implicit causality information exerts its influence at the point of integration. On reading time to the first fragment we found only evidence of a repeat name penalty of the same type as previously reported, i.e. longer reading times associated with reading a repetition of the first mentioned

character's name. We found no hint of an interaction between verb bias and any other factor.

Evidence for the influence of verb bias was found on reading times to the second fragment. We found an implicit causality congruency effect. Although we did not find a 3-way interaction between the effect and anaphor type, the implicit causality congruency effect was present in the Pronoun conditions but there was only weak evidence for it in the Name conditions. Recall we found a very weak verb bias x referent interaction that was marginally significant by subjects but not by items following a repeat name anaphor.

One possible reason why there was only a weak hint of the congruency effect in the Name conditions is due to load restrictions. In the context of additional information which may be relevant for integration, i.e. information contained in the additional adverbial phrase, the influence of verb bias may be weakened somewhat. The stage of co-indexation proceeds as normal when the repeat name is encountered while the stage of semantic integration is now informed by both verb bias information and the information following the anaphor but also by the additional information contained within the adverbial phrase. Simply because a number of factors are now potentially relevant, load restrictions within the system may mean that each individual factor exerts a relatively weak influence on processing explaining why we only find a weak hint of an influence of implicit causality information. This is supported by our finding that overall the magnitude of the implicit causality congruency effect is weaker in this study than in Experiment 3. The fact that we did not find any between Experiment difference in the magnitude of the repeat name penalty further supports the position that this is informed by low level factors and so is uninfluenced by load restrictions which occur when more high level (semantic) information is present.

For repeat name anaphors we have argued that co-indexation occurs when the anaphor is encountered and without reference to verb bias information. For pronouns co-indexation is delayed. Verb bias information is important for pronominal co-indexation. In this

experiment, verb bias information itself still plays an important role during co-indexation following reading of the pronoun although its influence at the stage of integration may be weaker in the context of the additional information conveyed by the adverbial phrase.

There is certainly a weaker processing load associated with the Pronoun conditions. We argue that when a certain load threshold is reached, and under conditions where less attention can be paid to certain factors, the influence of implicit causality becomes reduced. This occurs in the Name conditions where antecedent identification can be achieved using low level information and where the questions following the sentences can be answered without integration of the main and subordinate clauses, but it does not occur in the Pronoun conditions where implicit causality is still important for antecedent identification. This helps explain why we found stronger evidence for an implicit causality congruency effect in the Pronoun conditions.

One effect that we have found previously (e.g. in Experiment 5) which we failed to find in this experiment is an interaction between anaphor type and referent. In Experiment 5 the interaction corresponded to integration being easy in the context of either a repeat name or pronominal anaphor referring to the first mentioned character. In the context of reference to the second mentioned character, integration was only easy when that reference was achieved through the use of a repeat name compared to a pronoun. In this experiment we found only a main effect of anaphor type suggesting that pronouns have somehow lost their privilege in facilitating integration following reference to the first mentioned character compared to similar behaviour following repeat names. This is probably due to an increase in processing load arising as a result of the extra information contained within the sentence. In section 2.5.1 we outlined a proposal that interpreted the first mention privilege as nothing more than a memory effect. The additional processing load arising from the phrase present in the experimental materials examined in this study may have interfered with this memory effect so leading to the disappearance of the first mention privilege. Alternatively this first mention privilege may simply decrease over time.

The aim of this experiment was to attempt to adjudicate between the conflicting positions proposed by McDonald and MacWhinney and by Garnham et al. In light of the evidence in our pattern of data associated with reading the second fragment on the experimental sentences, the position of Garnham et al, that implicit causality influences anaphor resolution at no point earlier than the point of integration, is supported. We suggest that the data reported by McDonald and MacWhinney have arisen as a result of an experimental confound.

CHAPTER 7

7.1 Introduction and Rationale

Up to this point all the experiments reported in this thesis have been concerned with examining the phenomenon of implicit causality and its influence within the anaphor resolution system. In this chapter we examine the influence of a factor we shall refer to as implicit consequentiality. In the same way that events have causes, they also have consequences. Causes of an event occurring at t arise at some point prior to t . Conversely, consequences arise at some point following t . In the same way that different types of events can have different types of causes, different types of events can also have different types of consequences. Recall our earlier example sentence 'John blamed Bill because he broke the window.' The act of NP1 blaming NP2 is usually preceded by the character occupying the NP2 position doing something to engender this outcome. This locus of cause gives rise to the implicit causality effect. Example (1) below describes an event and the consequences following. The act of NP1 punishing NP2 is usually followed by some consequence affecting the character occupying the NP2 position. At least we can propose this to be the case if we can demonstrate implicit consequentiality biases of a similar sort to implicit causality biases.

In example (1), what follows from John punishing Bill is Bill starting to cry.

(1) Because John punished Bill, Bill started to cry.

There is no restriction as to which character may experience the consequences of the event. In example (2), the same subordinate clause is followed by a consequence regarding the first mentioned character, 'John'.

(2) Because John punished Bill, John was accused of being too strict.

The basis for our examination of the implicit causality effect was that certain verbs contain causal biases imputing cause to one of the two

mentioned characters and that these verbs exert a processing influence on reading. We can now ask ourselves whether there are equivalent biases which, in sentences such as (1) and (2) above, impute a consequence to one of the two mentioned characters. Finding evidence for such a bias is interesting for a number of reasons. Recall that Garvey and Caramazza (1974) proposed that implicit causality biases were properties encoded at the verb root. If we can show that there are biases above and beyond implicit causality we can look at what basic properties are common to the biases and we can begin to speculate over what sort of mechanism may be giving rise to the effects. Biases may be defined generally as some property such that when that verb exhibiting the bias appears in a particular context it influences, in language production, some form of lexical selection or, in the case of language comprehension, interpretation.

Certainly, there may be some statistical preference giving rise to a bias toward interpreting a pronoun as coreferential with one character rather than another in the sorts of constructions we have looked at so far, but we need to ask ourselves what factors might have given rise to those statistical preferences. In other words, prior experience with a particular verb or with the types of events typically described by such verbs may lead to a predisposition to interpret the causal structure of events described by the verb in the same sorts of ways that have been done previously. In the sorts of sentences we examine in this thesis, this disposition may be reflected as a general preference to interpret pronouns as coreferential with one character rather than another. Although this can be encoded as some sort of bias, what needs addressing is the sort of mechanism that gave rise to such a bias. For instance, has it come about simply through experience of the particular verb or is it at a more abstract level of representation, partly determined by experience of the events typically described by such verbs?

Any explanation of the effect per se should go beyond merely a description of the conditions which are necessary for the effect to be observed. For the purposes of this thesis we are not so much concerned with the mechanism giving rise to the bias, but rather with

the way in which the bias exerts a processing influence within the language system.

In the same way that we performed an off-line plausibility study on our set of implicit causality materials, we performed an equivalent study on our set of verbs for which we were interested in examining consequential biases.

7.2 Experiment 9a

The goal of our first off-line study was to select implicit consequentiality verbs possessing strong biases.

Method

Subjects

22 English speaking subjects participated.

Stimuli

Experimental booklets were constructed containing 110 sentence fragments. We examined the same 50 verbs we examined in Experiment 1a and embedded each verb in a sentence of the form :

Because Sue VERBED John, ...

Names were varied. Each booklet contained 60 filler materials of a similar structure to the above fragment. Each implicit causality verb appeared only once. The gender of the two characters in each sentence were different.

Procedure

One random order was used for half the booklets while the reverse of this was order used for the other half. Subjects were simply instructed to provide a sensible continuation to the sentence fragments.

7.2.1 Results

The endings for each experimental sentence were scored on the basis of whether the main clause initially contained an anaphor referring to the first or second noun phrase. Whether subjects used a pronoun or repeat name anaphor was also scored. Sentences where the main clause contained an initial plural reference were scored as 'They' and sentences containing continuations which didn't fall into any of the above categories, such as 'Because Mary loved John, her parents were angry.' were scored as 'other'. For each verb, the total number of continuations which fell into each category of continuation are reported below :

Table 1 - Experiment 9a continuation results

	NP1 Pronoun	NP2 Pronoun	NP1 Name	NP2 Name	'They'	Other
consoled	0	18	0	1	1	2
concerned	1	19	0	1	0	1
amazed	2	18	0	1	0	1
noticed	9	8	0	0	0	5
called	2	10	0	1	0	9
aggravated	0	19	0	1	2	0
annoyed	6	13	0	2	1	0
scolded	0	19	0	2	1	0
approached	3	14	0	1	0	4
confused	0	20	0	0	1	1
amused	0	18	0	1	2	1
praised	0	19	0	1	0	2
loathed	19	2	0	0	0	1
punished	1	20	0	1	0	0
telephoned	3	13	0	1	2	3
bored	6	15	0	1	0	0
dumbfounded	3	18	0	0	0	1
dreaded	17	1	0	0	0	4
detested	20	1	0	0	1	0
troubled	4	17	0	1	0	0
confided in	4	12	0	0	3	3
flattered	0	19	0	0	0	3
inspired	3	15	0	1	0	3
thanked	0	22	0	0	0	0
deplored	15	5	0	0	0	2
congratulated	2	19	0	1	0	0
resented	18	2	0	0	1	1
appreciated	17	3	0	0	2	0
ridiculed	3	14	0	2	0	3
angered	1	18	0	2	0	1

apologised to	1	16	0	1	0	4
confessed to	4	14	0	1	1	2
believed	14	5	0	1	0	2
valued	18	1	0	0	1	2
despised	18	2	0	0	1	1
derided	1	16	0	0	0	5
complemented	1	19	0	0	0	2
disappointed	2	18	0	0	0	2
honoured	10	8	0	0	1	3
agitated	1	18	0	1	0	2
accused	1	10	0	1	1	9
questioned	2	16	0	0	0	4
infuriated	4	16	0	0	0	2
fascinated	2	20	0	0	0	0
admired	22	0	0	0	0	0
reassured	0	21	0	0	0	1
criticised	0	21	0	0	0	1
liked	19	1	0	0	0	2
adored	21	1	0	0	0	0
lied to	5	16	0	0	1	0

The data reported in the above table provides us with an indication of the relative strengths of the implicit consequentiality bias associated with each verb. The maximum number of continuations of each category for each verb is 22. The bias strength for the verbs we examined is reported below along with an analysis of the plausibility rating.

7.3 Experiment 9b

The goal of this off-line experiment was to ensure that the materials in each of the experimental conditions we intended to examine in our on-line tasks were of equivalent plausibility. As for our implicit causality materials, we wanted to be sure that the version containing the ending consistent with the verb bias and the version containing the ending inconsistent with the bias matched on a plausibility rating. We did this to exclude an explanation for the difference in reading time between the consistent and inconsistent conditions in terms of a simple difference in their levels of plausibility.

Using the verbs examined in Experiment 9a we constructed materials similar to examples (3) and (4) below. In other words, for each verb we constructed two versions, one containing reference to the first mentioned character, the other containing reference to the second.

(3) Because Harold dreaded Justin, Harold steadfastly refused to go back to school.

(4) Because Harold dreaded Justin, Justin was told to try acting less aggressively.

The verb 'dreaded' is NP1 biasing. Example (3) is consistent with the bias as the reference is to the first Noun Phrase. Example (4) contains a continuation inconsistent with the bias as reference is to the second Noun Phrase. We want the plausibility of each of these sentences to be equivalent. When the repeat name anaphor is replaced by a pronoun, the information predicated by the pronoun in the main clause should provide unambiguous evidence as to which character is the appropriate pronominal referent. Examples (5) and (6) below should be implausible interpretations.

(5) Because Harold dreaded Justin, Justin steadfastly refused to go back to school.

(6) Because Harold dreaded Justin, Harold was told to try acting less aggressively.

To measure the plausibility of these four possible interpretations we constructed two endings which for each verb produced a set of 4 items. We equated the length for each of these 4 conditions for each material.

(7) Because Harold dreaded Justin, Harold steadfastly refused to go back to school.

(8) Because Harold dreaded Justin, Justin was told to try acting less aggressively.

(9) Because Harold dreaded Justin, Justin steadfastly refused to go back to school.

(10) Because Harold dreaded Justin, Harold was told to try acting less aggressively.

To recap, we want Examples (7) and (8) to be rated as highly plausible while we want Examples (9) and (10) to be rated as highly implausible.

Method

Subjects

24 English speaking subjects participated.

Stimuli

Experimental booklets were constructed containing 96 sentences. We selected 24 verbs from our list exhibiting strong implicit consequentiality biases. Each subject saw 4 instances of each verb, once in each of the 4 conditions. Randomisation of presentation order occurred in 4 blocks to prevent any sentence pair containing the same verb being in close proximity. Half the subjects saw the sentences in one random order, for the other half this order was reversed. Subjects were simply asked to rate the plausibility for each sentence. The instructions given at the start of the booklet were of the same type as those used in Experiment 1a.

7.3.1 Results

We calculated the average plausibility rating for the experimental materials. The plausibility ratings should be statistically equal for the versions corresponding to plausible continuations for the first and second mentioned character conditions and for the versions corresponding to implausible continuations for the first and second mentioned character conditions. The plausibility scores range from 0 to 7 with 0 standing for the most implausible and 7 for the most plausible. The scoring category Plaus1 corresponds to what we want to be a plausible interpretation of the version of the sentence containing reference to the first mentioned character. Category Implaus1 corresponds to what we want to be an implausible interpretation of the version containing reference to the first mentioned character. The scoring category Plaus2 corresponds to

what we want to be a plausible interpretation of the version of the sentence containing reference to the second mentioned character. Category Implaus2 corresponds to what we want to be an implausible interpretation of the version containing reference to the second mentioned character.

We performed a separate ANOVA on the Plausibility ratings for the NP1 and the NP2 set of verbs. For the NP1 verbs we found a main effect of plausibility ($F(1,7)=410$, $p<0.0001$). There was no interaction ($F<1$) between plausibility and referent. In other words we can treat the plausibility ratings for Plaus1 and Plaus2 as statistically equivalent and likewise for Implaus1 and Implaus2.

For the NP2 verbs we found a main effect of plausibility ($F(1,7)=496$, $p<0.0001$). We found an interaction ($F(1,7)=9.33$, $p<0.05$) between plausibility and referent. This corresponds to the difference between the Implausible and Plausible 2 versions of the materials being greater than the difference between the Implausible and Plausible 1 versions. Figure 1 depicts this interaction.

Those verbs are reported in the tables following :

Table 2 - Experiment 9b plausibility results combined with experiment 9a results


NP1 Verbs

	NP1 continuation	NP2 continuation	Plaus1	Implaus2	Plaus2	Implaus1
dreaded	17	1	6.59	1.06	5.59	2.06
appreciated	17	3	6.0	2.71	6.29	3.47
adored	21	1	6.59	2.89	5.12	2.82
admired	22	0	5.0	2.06	6.35	1.65
loathed	19	2	5.2	1.9	5.05	1.68
resented	18	2	5.1	1.45	6.17	1.88
despised	18	2	5.35	2.37	6.1	1.65
liked	19	1	5.0	1.75	5.82	2.0
Mean	18.875	1.5	5.6	2.0	5.8	2.2

NP2 Verbs

	NP1 continuation	NP2 continuation	Plaus1	Implaus2	Plaus2	Implaus1
aggravated	0	20	5.94	1.08	6.29	1.82
annoyed	6	15	5.35	1.59	6.53	0.71
scolded	0	21	6.47	0.94	6.47	1.59
confused	0	20	6.35	0.88	6.0	1.0
punished	1	21	6.65	0.65	6.18	1.72
flattered	0	19	5.35	1.72	6.53	2.18
thanked	0	22	5.76	1.53	6.71	0.71
congratulated	2	19	5.29	1.76	6.18	2.35
Mean	1.1	19.6	5.9	1.3	6.4	1.5

Figure 1 - plausibility x NP



Although this interaction is significant, the size of the plausibility difference between the two continuations is so small that we feel confident in assuming that it will not noticeably influence any pattern of reading time data we obtain. The absolute plausibility difference between conditions for NP2 biasing verbs is 4.4 for the NP1 versions and 5.1 for the NP2 versions.

7.4 Experiment 10

7.4.1 Introduction and Rationale

Given the indication from our off-line experiments that implicit consequentiality seems to be a property conceptually similar to implicit causality, we want to investigate whether the on-line nature of its processing influence is also similar to the on-line processing influence of implicit causality. On the basis of the data generated by the experiments reported in the previous chapters we proposed that implicit causality exerts the bulk of its influence at integration. When repeat name anaphors are encountered co-indexation occurs immediately informed solely by low level information. For

pronominal anaphors however, co-indexation is delayed until the information following the pronoun has been read.

The predictions we make for the on-line influence of implicit consequentiality information are motivated by our implicit causality findings. We expect the second fragment of example (11) below to be read more quickly than example (12) as it contains a continuation consistent with the implicit consequentiality bias. The verb 'dreaded' possesses an NP1 implicit consequentiality bias.

(11) Because Harold dreaded Justin, he / steadfastly refused to go back to school.

(12) Because Harold dreaded Justin, he / was told to try acting less aggressively.

In the conditions where the pronominal anaphor is replaced by a repeat name anaphor (example 13 and 14 below), on reading time to the first fragment we expect to find a repeat name penalty of the sort we have consistently found up to this point, i.e. a penalty associated with repetition of the first mentioned character's name independent of any other factor.

(13) Because Harold dreaded Justin, Harold / steadfastly refused to go back to school.

(14) Because Harold dreaded Justin, Justin / was told to try acting less aggressively.

Again on the basis of the results generated by our examination of the processing influence of implicit causality for pronominal anaphora we do not expect any difference in integration following a pronoun referring to the first mentioned character as a function of whether there is additional verb bias support. Recall, in our examinations of implicit causality we found that integration following pronominal reference to the first mentioned character is easy regardless of whether it is supported by verb bias. Integration following pronominal reference to the second mentioned character was easy only when supported by verb bias information. We do not expect the processing behaviour of implicit consequentiality to be any different.

It should be pointed out however that the difference between the following experiment and its implicit causality counterpart (Experiment 3) is not simply that the current one examines implicit consequentiality and Experiment 3 examined implicit causality, the order of the clauses presented to the reader is also different. While our previous single sentence examinations of the influence of verb bias were restricted to main-subordinate clause order sentences our examination of the implicit consequentiality verb bias requires sentences of a subordinate-main clause order type. In subordinate-main clause order sentences, the anaphor appears as the subject of the main clause, For main-subordinate clause order sentences however, the anaphor appears as the subject of the subordinate clause. There may be processing differences between these two types of constructions above and beyond differences arising as result of properties associated with the different types of verb biases themselves. Anaphors may be treating differently as a function of whether or not they appear in a main or subordinate clause. We shall return to these possible processing differences when we come to interpret the data.

Method

Subjects

Sixty four English speaking subjects participated.

Stimuli

We manipulated type of anaphor, verb bias and referent. The anaphor could either be a repeat name or ambiguous pronoun. Verb bias could either be NP1 or NP2 biasing. Referent could either be first or second mentioned character.

There were 32 sets of experimental materials (see appendix 6 for full set).

Name / NP1 verb / Referent character1

Because Harold dreaded Justin, Harold steadfastly refused to go back to school.

Name / NP1 verb / Referent character2

Because Harold dreaded Justin, Justin was told to try acting less aggressively.

Pronoun / NP1 verb / Referent character1

Because Harold dreaded Justin, he steadfastly refused to go back to school.

Pronoun / NP1 verb / Referent character2

Because Harold dreaded Justin, he was told to try acting less aggressively.

There were also 96 filler items (appendix 8). The 16 verbs used to construct the experimental materials were selected from the set of verbs examined in the off-line studies (see list above). Each subject saw each verb twice. The experiment was divided into two halves with a break halfway through. The first occurrence of each verb was in the first half of the study and the second occurrence in the second half.

Procedure

The procedure was identical to Experiment 3.

Each subject participated in 10 practice trials similar in structure to the experimental items at the start of the experiment. The experiment lasted roughly 35 minutes. Before the experiment subjects were provided with both verbal and written instructions (see Chapter 3 above for full instructions).

7.4.2 Results

We report analyses for the two fragments of the experimental sentences themselves and also for the question responses following those fragments.

Outlier Replacement

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For reading times to the first fragment we excluded response times that were below 300 msec or above 15 seconds. This accounted for 0.9% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.7% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the data corresponding to reading times to the second fragment we excluded response times that were below 300 msec or above 15 seconds. This accounted for 0.5% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.6% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the question response data we excluded response times that were below 300 msec or above 20 seconds. This accounted for 0.5% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 3.2% of the data was replaced in this way.

First Fragment Analysis _


We performed 2 (Pronoun anaphor vs Repeat name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type significant by subjects only ($F(1,63)=4.23$, $p<0.05$; $F(1,30)=3.77$, $p<0.1$) indicating faster reading time associated with reading fragments containing pronouns. We also found a main effect of verb bias significant by subjects only ($F(1,63)=6.07$, $p<0.05$; $F(2,2)$) corresponding to fragments containing NP1 biasing verbs being read more quickly than those containing NP2 biasing verbs. The anaphor type x referent interaction was significant by both subjects and items

($F(1,63)=9.18$, $p<0.005$; $F(1,30)=5.49$, $p<0.05$) corresponding to the repeat name penalty for repeat name anaphors. No other main effects or interactions approached significance. As the pronouns are identical for each condition in fragment 1 we also looked at the Name conditions separately. 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors revealed a main effect of referent ($F(1,63)=10.12$, $p<0.005$; $F(1,30)=5.27$, $p<0.05$) corresponding to the repeat name penalty (see Figure 2).

Table 3 - Reading times for first fragment, subject means for each condition with all times in msecs.

Verb / Referent	Name	Pronoun
NP1 / Char1	2701	2509
NP1 / Char2	2519	2508
NP2 / Char1	2881	2543
NP2 / Char2	2600	2662

Figure 2 - referent main effect (repeat name anaphor)



Second Fragment Analysis


We performed 2 (Pronoun anaphor vs Repeat name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both

subjects and items as random factors. We found a main effect of anaphor type ($F(1,63)=45.89, p<0.0001$; $F(1,30)=42.61, p<0.0001$) indicating faster reading time following reference through the use of a proper name. We found a main effect of verb bias ($F(1,63)=25.65, p<0.0001$; $F(1,30)=5.73, p<0.05$) indicating faster reading time following an NP2 biasing verb. We also found a marginal main effect of referent that was only significant by subjects ($F(1,63)=4.59, p<0.05$; $F(1,30)<2$) suggesting faster reading time following reference to the second mentioned character. We found an anaphor type x referent interaction ($F(1,63)=7.45, p<0.01$; $F(1,30)=4.57, p<0.05$) (see Figure 3 below) and a verb bias x referent interaction corresponding to an implicit consequentiality congruency effect that was only significant by subjects ($F(1,63)=5.50, p<0.05$; $F(1,30)<2$).

Table 4 - Reading times for second fragment, subject means for each condition with all times in msecs.

Verb / Referent	Name	Pronoun
NP1 / Char1	2328	2560
NP1 / Char2	2149	2766
NP2 / Char1	2191	2482
NP2 / Char2	1886	2348

Figure 3 - anaphor type x referent




Means comparisons reveal that the points corresponding to integration following pronominal reference to either the first or second mentioned character are equivalent (both $F_s < 1$). However, the points corresponding to integration following reference to either the first or second mentioned character using a repeat name anaphor are different ($F_1: p < 0.005$; $F_2: p < 0.01$). This is a persistence of our repeat name penalty corresponding to longer reading times following a repeat name referring to the first mentioned character compared to a repeat name referring to the second mentioned character.

We examined the Pronoun and Name conditions separately.

Pronoun Conditions Analysis

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias ($F_1(1,63)=15.15, p < 0.0005$; $F_2(1,30)=4.42, p < 0.05$) reflecting faster reading time following an NP2 biasing verb. We also found a verb bias x referent interaction corresponding to an implicit consequentiality congruency effect that was significant by subjects only ($F_1(1,63)=5.77, p < 0.05$; $F_2(1,30)=2.43, p < 0.13$) (see Figure 4 below).

Figure 4 - verb bias x referent (Pronoun conditions)



We performed 2-way ANOVAs for NP1 and NP2 biasing verbs separately. For NP1 verbs we found that Char1 and Char2 referents marginally differ by subjects but not by items ($F(1,63)=3.30$, $p<0.08$; $F_2<2$). For NP2 verbs we found that Char1 and Char2 referents are equivalent (both $F_s<2$). We also performed 2-way ANOVAs for Char1 and Char2 referents separately. For Char1 referents we found that NP1 and NP2 verbs are equivalent (both $F_s<1$). For Char2 referents we found that NP1 and NP2 verbs differ ($F(1,63)=18.82$, $p<0.0001$; $F_2(1,30)=6.73$, $p<0.05$). This means that for pronouns, reading times following reference to the first mentioned character are equivalent regardless of whether that reference is supported by verb bias. Following reference to the second mentioned character, reading is fast when supported by verb bias compared to when there is no verb bias support.

Name Conditions Analysis

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias that was marginal by items ($F(1,63)=14.52$, $p<0.0005$; $F_2(1,30)=3.76$, $p<0.07$) reflecting faster reading time following an NP2 biasing verb. We also found a main

effect of referent ($F(1,63)=17.34$, $p<0.001$; $F(1,30)=5.05$, $p<0.05$) reflecting faster reading time following reference to the second mentioned character. The verb bias x referent interaction did not approach significance (both $F_s<1$). In other words, we did not find any evidence of an implicit consequentiality congruency effect.

Question Response Time Analysis

We performed 2 (Pronoun anaphor vs Repeat name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type ($F(1,63)=73.99$, $p<0.0001$; $F(1,30)=118.22$, $p<0.0001$) corresponding to faster question response times following sentences containing repeat name anaphors. We also found a main effect of verb bias ($F(1,63)=13.15$, $p<0.001$; $F(1,30)=4.47$, $p<0.05$) and of referent ($F(1,63)=7.34$, $p<0.01$; $F(1,30)=4.42$, $p<0.05$). These correspond to faster question response times following sentences containing NP2 biasing verbs and reference to the first mentioned character respectively. One interaction was significant, that of anaphor type x referent ($F(1,63)=22.02$, $p<0.0001$; $F(1,30)=16.84$, $p<0.0005$).

Table 5 - Question response times, subject means for each condition with all times in msecs. Numbers in parentheses correspond to response accuracy expressed as percentage correct.

Verb / Referent	Pronoun	Name
NP1 / Char1	2821 (87.1%)	2344 (91.0%)
NP1 / Char2	3590 (57.4%)	2128 (94.1%)
NP2 / Char1	2676 (89.0%)	2033 (95.3%)
NP2 / Char2	3174 (82.7%)	1912 (95.3%)

Question Response Accuracy Analysis

We performed 2 (Pronoun anaphor vs Repeat name anaphor) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of anaphor type ($F(1,63)=101.40$, $p<0.0001$; $F(1,30)=84.64$, $p<0.0001$) corresponding to better question response accuracy following repeat

name anaphors. We found a main effect of verb bias ($F(1,63)=23.53$, $p<0.0001$; $F(1,30)=15.76$, $p<0.0005$) corresponding to better question response accuracy following sentences containing NP2 biasing verbs. We found a main effect of referent ($F(1,63)=27.99$, $p<0.0001$; $F(1,30)=18.15$, $p<0.0005$) corresponding to better question response accuracy following reference to the first mentioned character. We found an interaction of anaphor type x verb bias ($F(1,63)=10.44$, $p<0.005$; $F(1,30)=11.80$, $p<0.005$). We found an interaction of anaphor type x referent ($F(1,63)=32.86$, $p<0.0001$; $F(1,30)=28.03$, $p<0.0001$). We found an interaction of verb bias x referent ($F(1,63)=16.21$, $p<0.0005$; $F(1,30)=6.63$, $p<0.05$) corresponding to an implicit consequentiality congruency effect. We also found a 3-way interaction of anaphor type x verb bias x referent ($F(1,63)=17.36$, $p<0.0001$; $F(1,30)=12.45$, $p<0.005$). This interaction corresponds to the presence of a congruency effect in the Pronoun conditions ($F(1,63)=24.66$, $p<0.0001$; $F(1,30)=11.27$, $p<0.005$) but not in the Name conditions (both $F_s<1$).

7.4.3 Discussion

Our finding of a main effect of referent in the Name conditions in fragment 1 reading times appears to correspond to the repeat name penalty. There is a penalty associated with reference to the first mentioned character if that reference is achieved through the use of the repetition of that character's name. This penalty persisted in the pattern of data associated with reading times to the second fragment. In our examination of implicit causality we raised the possibility that this penalty was actually arising as a result of unambiguous reference. In the case of our examinations of implicit causality, we countered this by examining the behaviour of an unambiguous pronoun, where ambiguity was determined by the presence or absence of a gender cue. For our examination of implicit consequentiality we cannot yet exclude the possibility that the repeat name penalty is perhaps coming about as a result of the presence of an unambiguous anaphor. The next experiment sets out to examine this by looking at the effect of gender cue.

In line with our previous findings, we found no interaction between what for the moment we assume to be the repeat name penalty and verb bias.

In the introductory section to this experiment we raised the possibility of processing differences associated with subordinate-main clause order sentences compared to main-subordinate clause order sentences. We now turn to what might constitute these processing differences.

It has been suggested (Cooreman and Sanford, 1996) that a delay in processing of the subordinate clause occurs when that clause occurs in first position. This may be the case but processing at a level sufficient to give rise to what appears to be a repeat name penalty is certainly occurring. In main-subordinate clause order sentences the sentence topic is contained within the main clause. Main clauses advance the narrative of a discourse while subordinate clauses do not. When subordinate-main clause order sentences are read the topic can only be fully understood following reading of the main clause. The subordinate clause acts almost like a comment on that topic. There is some evidence to suggest that subordinate clauses are less deeply processed than main clauses (Bever and Townsend, 1979).

As the relevance of the information contained within the subordinate clause cannot be fully apprehended until the information in the main clause has been extracted, it seems reasonable to suggest that some processing of the subordinate clause is effectively delayed until information necessary for interpretation becomes available from the main clause (cf. cataphoric anaphors in subordinate clauses whose referent follows the anaphor). If processing of subordinate clauses does operate over a slightly larger time window than processing of comparable information contained within a main clause or if some degree of processing of the subordinate clause occurs in parallel with processing of the main clause then we would expect phenomena that result from processing of the subordinate clause to persist for a slightly longer period of time. If this holds it can account for the persistence of the repeat name penalty that we found in our analysis of the second fragment reading time data.

With respect to the implicit consequentiality congruency effect we found a marginal verb bias x referent interaction but only for the conditions containing pronouns. The pattern of data is similar to the pattern associated with the implicit causality congruency effect with reference to the first mentioned character being easy, reference to the second only easy when supported by verb bias.

Our implicit consequentiality congruency effect was found in the question response accuracy data but only for the Pronoun conditions as revealed by our 3-way interaction. It corresponds to equivalently high response accuracy following reference to the first mentioned character in either the context of an NP1 or NP2 biasing verb. Response accuracy was also high following reference to the second mentioned character but only when this reference was supported by verb bias.

We found evidence suggesting an implicit consequentiality congruency effect on both reading time to the second fragment and in question response accuracy. We only found this for the Pronoun conditions however. No analysis revealed the congruency effect arising for the Name conditions. The only effect found for the Name conditions was the repeat name penalty which occurred in both the reading time data for fragment one and fragment two. The most obvious possible reason for why we failed to find a congruency effect in the Name conditions is because subjects simply aren't fully integrating the information contained in the subordinate and main clauses. We shall propose a possible explanation for this pattern of data below in light of the data we report for Experiment 11.

We found the repeat name penalty associated with reference to the first mentioned character using a repeat name anaphor independent of verb bias. We found some evidence for an implicit consequentiality congruency effect in reading time to the second fragment but only in the Pronoun conditions. The congruency effect was stronger in the pattern of data associated with question response accuracy.

In light of the possibility that what appears to be the repeat name penalty may be resulting from the presence of an unambiguous referring device we shall delay relating the data generated by the above experiment to our two-stage model of anaphor resolution until this possible explanation is addressed.

7.5 Experiment 11

7.5.1 Introduction and Rationale

As we mentioned in the discussion to the previous experiment, Experiment 11 examines the behaviour of implicit consequentiality information in the presence and absence of a gender cue. When a gender cue is present, this is sufficient to uniquely identify a pronominal antecedent (example (11)). When gender information is sufficient to identify a pronoun's antecedent, what role does implicit consequentiality information play ?

(11) Because John punished Sue, she started to cry.

In the case of implicit causality we found evidence that gender does not exert an influence when a pronoun is initially encountered. However, the difference between processing of a main-subordinate clause order sentence where the pronoun occurs in the subordinate clause may differ in some important way from processing of a subordinate-main clause order sentence where the pronoun occurs in the main clause.

We suggested in our discussion for the previous experiment that some processing of the subordinate clause may still be occurring as the following main clause is read. It would be reasonable to propose that if processing of the pronoun is delayed as a result of the subordinate-main clause order and this processing doesn't occur before verb bias information is available, the fact that it is disambiguated by gender may have little impact. If this is the case we expect to simply find an implicit consequentiality congruency effect on reading time to the second fragment and no interaction with ambiguity.

If immediate processing of the subordinate clause occurs that includes co-indexation of the pronoun, this may occur solely with reference to gender information. If this is the case we expect to find no interaction between ambiguity and verb bias on reading time to the first fragment (with the split occurring following the pronoun). However, it should be easier to integrate information inconsistent with the verb bias if a gender contrast is available. We used this argument in Experiment 6 and proposed that if gender information is used immediately we would have expected an interaction between ambiguity and the verb bias congruency effect. We did not find such an interaction in the case of implicit causality but it may be the case that the difference in processing associated with subordinate-main clause order sentences gives rise to such a result.

In the experiments reported in Chapters 4 and 5 we only examined the processing of pronouns when they occurred in a subordinate clause. If, as has been suggested, processing of subordinate clauses in general is somehow shallower than processing of information in main clauses, processing of the pronoun may simply be delayed under those circumstances. If pronouns when they occur in main clauses are processed, or co-indexed immediately, we expect to find evidence of some degree of co-indexation in this experiment.

If pronouns in main clauses are co-indexed immediately and gender information is used in this assignment, we expect to find an ambiguity x congruency interaction in this study.

If, as we argue is the case for implicit causality and main-subordinate clause order sentences, implicit consequentiality exerts the bulk of its influence at the integrative stage of processing and co-indexation of a pronoun is always delayed, we expect to find congruency effects on reading time to the second fragment and no an interaction with ambiguity on reading time to the first fragment.

Method

Subjects

Thirty two English speaking subjects participated.

Stimuli

We manipulated pronominal ambiguity, verb bias and referent. The pronoun could either be referentially ambiguous or marked for gender to unambiguously refer to one of the two participants. Verb bias could either be NP1 or NP2 biasing. Referent could either be first or second mentioned character.

There were 32 sets of experimental materials (see appendix 7 for full set).

Unambiguous pronoun / NP1 verb / Referent character1

Because Harold dreaded Joanne, he steadfastly refused to go back to school.

Unambiguous pronoun / NP1 verb / Referent character2

Because Joanne dreaded Justin, he was told to try acting less aggressively.

Ambiguous pronoun / NP1 verb / Referent character1

Because Harold dreaded Justin, he steadfastly refused to go back to school.

Ambiguous pronoun / NP1 verb / Referent character2

Because Harold dreaded Justin, he was told to try acting less aggressively.

There were also 96 filler items (appendix 8). The 16 verbs used to construct the experimental materials were identical to those examined in the previous study. Each subject saw each verb twice. The experiment was divided into two halves with a break halfway through. The first occurrence of each verb was in the first half of the study and the second occurrence in the second half.

Procedure

The procedure was identical to the previous experiment.

Each subject participated in 10 practice trials similar in structure to the experimental items at the start of the experiment. The experiment lasted roughly 35 minutes. Before the experiment subjects were provided with both verbal and written instructions (see Chapter 3 above for full instructions).

7.5.2 Results

We report analyses for the two fragments of the experimental sentences themselves and also for the question responses following those fragments.

Outlier Replacement

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For reading times to the first fragment we excluded response times that were below 400 msec or above 14 seconds. This accounted for 1.0% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.8% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the data corresponding to reading times to the second fragment we excluded response times that were below 400 msec or above 15 seconds. This accounted for 0.7% of the data. Times falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 2.9% of the data was replaced in this way.

We removed data from the two tails of the distribution where there was a clear discontinuity in response times. For the question response data we excluded response times that were below 500 msec or above 15 seconds. This accounted for 0.9% of the data. Times

falling above or below two and a half standard deviations from the mean for each subject were replaced by that point. 3.3% of the data was replaced in this way.

First Fragment Analysis

We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. No main effects or interactions approached significance. We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors for just the ambiguous pronouns. No main effects or interactions approached significance.

Table 6 - Reading times for first fragment, subject means for each condition with all times in msec.

Verb / Referent	Ambiguous Pronoun	Unambiguous Pronoun
NP1 / Char1	2332	2159
NP1 / Char2	2174	2201
NP2 / Char1	2391	2230
NP2 / Char2	2334	2355

Second Fragment Analysis


We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of ambiguity ($F(1,31)=24.64$, $p<0.0001$; $F(1,30)=29.37$, $p<0.0001$) reflecting faster reading times following unambiguous reference. We found a main effect of verb bias ($F(1,31)=23.09$, $p<0.0001$; $F(1,30)=3.94$, $p<0.06$) reflecting faster reading time following fragments containing NP2 biasing verbs. We found an interaction between ambiguity and verb bias ($F(1,31)=4.55$, $p<0.05$; $F(1,30)=2.96$, $p<0.1$) that was marginal by items. We found an ambiguity x referent interaction but was significant by subjects only ($F(1,31)=4.87$, $p<0.05$; $F(1,30)<2$). We also found a verb bias x referent interaction which again was significant by subjects only

($F(1,31)=5.39$, $p<0.05$; $F(1,30)<1$) (see Figure 5 below). The 3-way interaction of verb bias x referent x ambiguity was significant by both subjects and items ($F(1,31)=4.26$, $p<0.05$; $F(1,30)=4.50$, $p<0.05$). We examine this 3-way interaction in detail below.

Table 7 - Reading times for second fragment, subject means for each condition with all times in msecs.

Verb / Referent	Ambiguous Pronoun	Unambiguous Pronoun
NP1 / Char1	2210	2208
NP1 / Char2	2566	2055
NP2 / Char1	2355	1809
NP2 / Char2	2241	1789


Figure 5 - verb bias x referent



We explored the nature of the 3-way interaction by performing analyses on the Ambiguous and Unambiguous Pronoun conditions separately. We shall firstly focus on the Ambiguous Pronoun conditions.

Ambiguous Pronoun Conditions

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a significant verb bias x referent interaction that was marginal by items ($F(1,31)=7.56$, $p<0.01$; $F(1,30)=3.09$, $p<0.1$) (see Figure 6 below).



Unambiguous Pronoun Conditions

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of verb bias ($F(1,31)=23.74$, $p<0.0001$; $F(1,30)=6.09$, $p<0.05$) corresponding to faster reading time following an NP2 biasing verb.

We can interpret the 3-way interaction as reflecting a difference in the implicit consequentiality congruency effect between the Ambiguous and Unambiguous Pronoun conditions. We found a congruency effect in the Ambiguous Pronoun conditions but failed to find one in the Unambiguous Pronoun conditions.

Question Response Time Analysis

We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of ambiguity ($F(1,31)=55.83$, $p<0.0001$; $F(1,30)=49.48$, $p<0.0001$) indicating faster question response times following a sentence containing unambiguous reference. We found a main effect of verb bias ($F(1,31)=8.57$, $p<0.01$; $F(1,30)=6.29$, $p<0.05$) indicating faster response times following a sentence containing an NP2 biasing verb. We also found a main effect of referent that was significant by subjects only ($F(1,31)=5.98$, $p<0.05$; $F(1,30)<2$) reflecting faster response times following reference to the first mentioned character. We found an ambiguity x referent interaction ($F(1,31)=12.26$, $p<0.005$; $F(1,30)=9.39$; $p<0.005$) and an ambiguity x verb bias x referent interaction significant by subjects only ($F(1,31)=4.43$, $p<0.05$; $F(1,30)=2.33$, $p<0.14$). This corresponds to a marginal verb bias x referent interaction in the Unambiguous Pronoun conditions significant by subjects only ($F(1,31)=4.35$, $p<0.05$; $F(1,30)<1$) and no such effect in the Ambiguous Pronoun conditions (both $F_s<2$).

Table 8 - Question response times, subject means for each condition with all times in msec. Numbers in parentheses correspond to response accuracy expressed as percentage correct.

Verb / Referent	Ambiguous Pronoun	Unambiguous Pronoun
NP1 / Char1	2323 (80.5%)	2083 (91.3%)
NP1 / Char2	2910 (50.4%)	1824 (92.9%)
NP2 / Char1	2229 (90.5%)	1730 (96.0%)
NP2 / Char2	2508 (74.0%)	1727 (93.7%)

Question Response Accuracy Analysis

For the question response accuracy data we also performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs for both subjects and items as random factors. We found a main effect of ambiguity

($F(1,31)=103.51$, $p<0.0001$; $F(1,30)=51.79$, $p<0.0001$) reflecting greater response accuracy following a sentence containing unambiguous reference. We found a main effect of verb bias ($F(1,31)=26.78$, $p<0.0001$; $F(1,30)=8.41$, $p<0.01$) reflecting greater response accuracy following a sentence containing an NP2 biasing verb. We also found a main effect of referent ($F(1,31)=34.74$, $p<0.0001$; $F(1,30)=21.25$, $p<0.0001$) reflecting greater response accuracy following a sentence containing reference to the first mentioned character. The ambiguity x verb bias interaction was significant ($F(1,31)=14.72$, $p<0.001$; $F(1,30)=6.56$, $p<0.05$). The ambiguity x referent interaction was also significant ($F(1,31)=24.04$, $p<0.0001$; $F(1,30)=16.27$, $p<0.0005$). The 3-way interaction of ambiguity x verb bias x referent was also significant but only by subjects ($F(1,31)=4.90$, $p<0.05$; $F(1,30)=2.40$, $p<0.14$).

7.5.3 Discussion

Our significant 3-way interaction found in our fragment 2 reading time data indicates a difference in the nature of the implicit consequentiality congruency effect between the Unambiguous and Ambiguous Pronoun conditions. We found a congruency effect similar to that found in Experiment 10 but again only significant by subjects. This congruency effect was only found for the Ambiguous Pronoun conditions. The Ambiguous Pronoun conditions in this study correspond to the Pronoun conditions in the previous experiment. In the current experiment no hint of a congruency effect was found in the case of the Unambiguous Pronoun conditions. This parallels our finding in Experiment 10 of a lack of evidence for a congruency effect in the Name conditions. Both the Unambiguous Pronoun and Name conditions can be interpreted as corresponding to unambiguous reference. We shall return to this lack of effect in the General Discussion section following.

7.6 General Discussion

The data of Experiments 10 and 11 are consistent with our two-stage model of anaphor resolution which we have proposed on the basis of the data generated by our examination of implicit causality

information. Again we find evidence of a repeat name penalty arising on reading the first fragment containing reference to the first mentioned character where this is achieved through the repetition of that character's name. This does not interact with any other factor and can't be considered simply a result of using an unambiguous anaphor as we didn't find a comparable effect in the Unambiguous Pronoun conditions in Experiment 11.

The basic question we posed at the outset of this chapter was whether we could find verb bias effects comparable to implicit causality effects but corresponding to implied consequences rather than implied causes. We have been partially successful in answering this. We did find some evidence that such an effect was present although this was rarely supported by significant F2 analyses. When it was supported by F2 analyses, it was in the pattern of question response accuracy data.

One intriguing finding that arose in our examinations of the influence of implicit consequentiality was the difference in its behaviour under conditions corresponding to unambiguous reference compared to conditions corresponding to ambiguous reference. This contrasts with the behaviour of implicit causality under equivalent conditions. The implicit causality congruency effect arose in contexts of both ambiguous and unambiguous reference.

Recall that with respect to reading time to the second fragment, we only found evidence for the implicit consequentiality congruency effect in the Ambiguous Pronoun conditions. We outlined in our introduction that if we found an effect between ambiguity and the congruency effect, we could take it as evidence that gender information was being used by the system before verb bias exerts an influence. The most obvious explanation for the difference between the influence of gender information in our implicit consequentiality studies compared to our implicit causality studies is that pronouns are treated differently depending on whether or not they appear in a main or subordinate clause. When they appear in a main clause, as they do in Experiments 10 and 11, an attempt at co-indexation is made immediately. Low level factors including gender information

but not verb bias informs this stage. When pronouns occur in a subordinate clause, as in the experiments reported in Chapters 4 and 5, co-indexation is delayed. We shall return to this in Chapter 8 when we come to discuss the set of experiments reported in this thesis as a whole.

Alternatively it may simply be that compared to implicit causality, implicit consequentiality is a weaker constraint. When reference can be established using alternative cues (i.e. gender information or matching of repeat proper names), implicit consequentiality information is simply not required in order to facilitate anaphoric reference resolution. If it is a weaker type of verb bias then integrating information contained within subordinate and main clauses that is inconsistent with the bias won't generally be problematic as the verb bias is simply a weak constraint that does not need satisfying provided other criteria (i.e. anaphor / antecedent co-indexation) are met. Although the magnitudes of the implicit consequentiality biases per se were comparable to the implicit causality ones, the actual implicit consequentiality constraint itself may be weaker.

There are a number of possible reasons why our data in general were not as clear as they were for our examinations of the influence of implicit causality in our previous chapters. Firstly we simply had fewer items which possessed strong implicit consequentiality biases (32 as opposed to 48 implicit causality materials). The main reason why this weakness is present is simply because we selected our implicit consequentiality verbs from our initial set of 50 implicit causality verbs. While many of these verbs are certainly appropriate to appear in sentences of the type we examined and possessed bias strengths similar in magnitude to the implicit causality verbs, we were unable to find enough to have the same number of good instances of implicit consequentiality verbs as we had implicit causality verbs. We were also unwilling to repeat each verb more than once.

Because of the striking similarity between the pattern of data we report in this chapter and the effects we obtained in our implicit

causality examinations, despite the equivocal nature of our data and in light of the possible problems we had with our materials we feel confident in proposing that implicit consequentiality is as real an effect as implicit causality and that it influences processing in a similar manner.

CHAPTER 8

This chapter provides a general summary of all the experimental work reported in this thesis. It begins by focusing on what common effects have been found through the examination of both implicit causality and implicit consequentiality verb biases. It then focuses on each of these phenomena in turn, relating the findings reported in Chapters 4-7 to our proposed two-stage model of anaphor resolution. Recall, we proposed that the initial stage of resolving anaphors consists of co-indexation between anaphor and antecedent while the second stage consists of semantic integration of the information following the anaphor with what has been read previously in the text. This chapter, and thesis, concludes with an expansion and re-statement of this proposed model.

8.1 When does verb bias influence anaphoric processing ?

Recall that in Experiments 3,4,6,8,10 and 11 when we found the verb bias congruency effect it was always on reading time to the second fragment, i.e. **following** the fragment in which the anaphor was encountered. In none of our examinations did we ever find the congruency effect on reading time to the first fragment which contained the anaphor. This occurred even under conditions which we suggested were likely to maximise our finding of such an early effect if it was present (Experiment 8). If verb bias information was being utilised by the system as soon as an anaphor was encountered, we would have expected to find a congruency effect on reading time to the first fragment when that fragment contained a repeat name anaphor. In Example (1) below, the anaphor 'Bill' provides strong indication that the information in the subordinate clause is going to continue against the direction of bias of the verb.

(1) John fascinated Bill because Bill / was easily entertained.

If implicit causality information was used by the system early, the knowledge that the verb 'fascinate' is NP1 biasing should influence interpretation of the repeat name anaphor. Therefore, we would have expected to find evidence of the implicit causality congruency effect

on reading time to the fragment containing the repeat name. In other words, we would have expected reading time to fragment 1 in Example (1) to be longer than reading time to fragment 1 in Example (2).

(2) John fascinated Bill because John / was very interesting.

We found a repeat name penalty on fragment 1 reading times that was independent of other factors (including verb bias). The penalty corresponds simply to an increase in reading time following repetition of the first mentioned character's name.

The nature of the repeat name penalty and the lack of a verb bias congruency effect also holds for our examination of those verbs exhibiting implicit consequentiality biases. In Experiment 10 we found a repeat name penalty of the type reported above on reading time to the first fragment. We did not find any evidence of an implicit consequentiality congruency effect on reading time to this fragment. As for implicit causality, any evidence we did find for an effect of verb bias was on reading time to the second fragment and in the question data.

As mentioned in previous chapters, we propose that pronouns are treated differently depending on whether they occur in a subject position in a main or subordinate clause. In our examination of the behaviour of implicit causality, the pronoun always occurred as the grammatical subject of a subordinate clause. In the case of our examination of implicit consequentiality, the pronoun always occurred as the grammatical subject of a main clause. With respect to when gender information is used to co-index a pronoun, we found a difference that may be due to the type of clause in which the pronoun is the grammatical subject. When a gender cue was present in our examination of implicit causality (Experiment 6), we found no evidence that gender information was used as soon as a pronoun was encountered. However, when we looked at implicit consequentiality (Experiment 11), we did find evidence that gender information was used immediately to co-index the pronoun.

It has been suggested (Cooreman and Sanford, 1996) that there may be important differences in the way in which subordinate and main clauses are processed. Indeed, the way in which pronouns, and anaphors more generally, are processed may differ substantially depending on whether they appear in a main or subordinate clause.

We must be careful to distinguish effects we propose arise from the processing influence of verb bias information from those effects that may arise as a result of processing differences between main-subordinate and subordinate-main clause order sentences. We first turn to the evidence we have gathered for the behaviour of implicit causality information and, by extension, its influence on the interpretation of anaphors when they appear as grammatical subjects of subordinate clauses.

8.1.1 Implicit Causality

We found a main effect of verb bias in our fragment 1 reading time data in our examinations of the influence of implicit causality but failed to find one in comparable examinations of the influence of implicit consequentiality. In the case of implicit causality, fragments containing NP2 biasing verbs were read more quickly than those containing NP1 biasing verbs.

As mentioned above we also found evidence for a repeat name penalty on reading times to the first fragment. This did not interact with verb bias. We suggest that it arises at the point of co-indexation and is driven by non-semantic factors. It corresponds to a penalty associated with a repetition of the first mentioned character's name. It appears a relatively short lasting effect as we find no evidence of its persistence in the pattern of reading time data associated with reading the second fragment of the sentences we examined. On the basis of the nature of the repeat name penalty, we propose that repeat name anaphors are co-indexed with their appropriate antecedent as soon as they are read and that this stage of co-indexation is informed solely by low level, structural information. We interpret the repeat name phenomenon as an indication that co-indexation occurs. In our examination of implicit causality, we found

that pronouns aren't co-indexed immediately. If co-indexation of pronouns did occur when they were encountered, in cases where the pronoun identified its antecedent unambiguously we would have expected to find a weaker congruency effect in the fragment 2 reading time data. If co-indexation occurred when the pronoun was encountered then when gender information was present, this should facilitate reading an incongruous continuation. We found no evidence of this and take that as indicating that co-indexation of pronouns is delayed.

On reading time to the second fragment we consistently found the implicit causality congruency effect. In our single sentence experiments we found that integration following a pronoun referring to the first mentioned character was always easy but integration following a pronoun referring to the second mentioned character was only easy when supported by verb bias. In the Name conditions in Experiment 3 we found that a mismatch of implicit and explicit causes led to a processing difficulty (i.e. there was no first mention privilege on fragment 2 reading times).

In Experiment 5 we examined the nature of the influence of implicit causality information when the information that we had previously presented in main-subordinate clause order sentences was presented in two separate sentences; see Examples (3) and (4).

(3) John fascinated Bill. This was because Bill / he was easily entertained.

(4) John fascinated Bill. This was because John / he was very interesting.

We found an interaction between the implicit causality congruency effect and anaphor type in the reading time data corresponding to reading the second sentence. Stated simply, the congruency effect was present in the Pronoun conditions, but absent in the Name conditions.

We argued that the repeat name anaphors are interpreted as thematic shift signals. This means that the reader does not initially integrate the information contained in the second sentence with that contained in the first. We did not find an interaction between the implicit causality congruency effect and anaphor type on the question response time data, suggesting that integration may be occurring for the Name conditions when it becomes apparent that it is appropriate to do so.

When pronouns are encountered in a subordinate clause we argued that co-indexation is delayed. In Experiment 6, where we examined the effect of presence or absence of a gender cue, we found no evidence that initial processing of the pronoun was occurring when a gender cue was present, even though it could be used to unambiguously identify the pronoun's antecedent. For sentences such as Examples (5) and (6) below, the pronoun is coreferential with the character 'John'. If gender information is used as soon as the pronoun is encountered, it should be easier to recover from a subordinate clause continuing against the direction of the verb bias (Example (5)) when a gender contrast is present than when it is absent (Example (6)). If the stage of co-indexation has been informed by gender information, only the stage of integration encounters a processing problem as the information going against the bias of the verb is integrated with that contained in the preceding clause. There will not be misassignment but there will be a clash of causes. When a gender contrast is not present, a processing problem arises at both the stage of co-indexation and integration. We would therefore expect the congruency effect to be reduced in conditions where a gender cue was present if that gender information was used by the processor as soon as a pronoun was encountered. We found no evidence that the processor attempted to co-index pronouns as soon as they were encountered.

(5) John blamed Mary because he / was in a bad mood.

(6) John blamed Bill because he / was in a bad mood.

On reading time to the second fragment in Experiment 6 we did not find any evidence that the presence or absence of the gender contrast affected the implicit causality congruency effect. We interpret this as meaning that even when information sufficient to identify a pronoun's antecedent is present, co-indexation between pronoun and antecedent does not occur when that pronoun is encountered. Co-indexation and integration both occur at some point following the pronoun. In other words, our reading time data associated with reading the fragment following a pronoun is effectively an index of both co-indexation and integration.

To summarise, we propose that co-indexation between a repeat name anaphor and its antecedent occurs as soon as the repeat name is read. For pronouns occurring as grammatical subjects in subordinate clauses, co-indexation is delayed. For both pronominal and repeat name anaphora, implicit causality exerts a processing influence during subsequent integration.

8.1.2 Implicit Consequentiality

Recall our sentences containing implicit consequentiality verbs were of a subordinate-main clause order type; see Example (7) below.

(7) Because Harold dreaded Justin, Harold / steadfastly refused to go back to school.

(8) Because Harold dreaded Justin, Justin / was told to try acting less aggressively.

The anaphor is the grammatical subject of the main clause. In the case of repeat name anaphors, again we found evidence for a repeat name penalty on reading times to the first fragment. This did not interact with verb bias. We suggest that the penalty arises at the point of co-indexation and is driven by low level (non-semantic) factors. It corresponds to a penalty associated with a repetition of the first mentioned character's name (Example (7) above). As we argue on the basis of the data gathered from examining the influence of implicit causality, we propose that the stage of co-indexation is

informed solely by low level information. In contrast however to our experiments on implicit causality, we find some evidence for the persistence of the repeat name penalty in reading time to the second fragment for sentences containing implicit consequentiality verbs. We propose the effect persists as a degree of processing of the subordinate clause is still occurring while the main clause is being read. We argued in Chapter 7 that a delay of some form is necessary and beneficial from a processing perspective. The information in the subordinate clause, acting as it does as a comment on the topic in the main clause, cannot be fully apprehended until the main clause has been read.

Also in contrast to our experiments on implicit causality, we failed to find any evidence for a congruency effect in reading time to fragment 2 in the Name conditions (Examples (7) and (8) above). This may have arisen as a result of the nature of the experimental task itself. We mentioned in Chapter 4 that, for the Name conditions in our examination of implicit causality, it was possible for subjects to answer the following question without having necessarily integrated the information in the conjoined clauses. In other words, subjects simply had to focus on the information contained within the subordinate clause in the Name conditions. Given a sentence such as (9) below, it is possible to answer the question 'Who was in a bad mood ? John or Bill ?' by simply focusing on the subordinate clause 'John was in a bad mood.' No integration is strictly necessary for subjects to answer the questions following the experimental items.

(9) John blamed Bill because John was in a bad mood.

As we did find evidence of a congruency effect in the Name conditions in our experiments reported in Chapter 4, we did not perceive this potential task effect to be of importance for our examinations of implicit causality. The congruency effect would not have arisen if readers were not integrating the information in the two clauses. However, it may be more of a problem for subordinate-main clause order sentences, i.e. our implicit consequentiality studies. As mentioned previously, there seems to be evidence that subordinate clauses are not processed as fully as main clauses. This fact,

combined with an experimental task which places further emphasis on the information contained in the main clause, may result in readers effectively not interpreting the repeat name in the main clause as anaphoric. In Example (10) below, the reader may simply focus on the information contained in the main clause, aware of the fact that that is sufficient for answering the following question.

(10) Because Harold dreaded Justin, Harold steadfastly refused to go back to school.

So then, a general tendency not to process subordinate clauses fully combined with task demands may account for our failure to find evidence of an implicit consequentiality congruency effect in the Name conditions in Experiment 10.

When we did find evidence for a congruency effect in the Pronoun conditions in Experiment 10 and the Ambiguous Pronoun conditions in Experiment 11, the items analyses were either non-significant (Experiment 10) or marginal (Experiment 11). We suggested two possible reasons for this. The first is that implicit consequentiality itself may simply be a weaker constraint than implicit causality. Violating the bias may not lead to great processing difficulty. Secondly, and perhaps more importantly, we simply examined fewer verbs with implicit consequentiality biases than verbs with implicit causality biases. The pattern of data we did find was similar to that found in our examination of implicit causality. For ambiguous pronouns, integration following reference which ultimately transpired to be to the first mentioned character was always easy while reference to the second was only easy when supported by verb bias. We have consistently found this in our examinations of the influence of implicit causality information (with the exception of Experiment 8).

8.1.3 Between Experiment Comparisons

Rather than simply focusing on the time course of the influence of low and high level factors in order to try to separate them, it is also possible to examine how they are affected by additional experimental manipulations. The argument for a dissociation between the way in

which low and high level factors influence anaphor resolution is supported by several findings reported in this thesis which have arisen through the direct manipulation of contextual factors. Our depth of processing manipulation in Experiments 3 and 4 influenced the strength of the implicit causality congruency effect but had no effect on the strength of the repeat name penalty. With shallower processing, the congruency effect was weaker. Similarly adding extra information in a clause preceding the connective (Experiment 8 compared with Experiment 3) resulted in a reduction in the implicit causality congruency effect but had no effect on the strength of the repeat name penalty. These two between experiment comparisons further support the position that the repeat name penalty arises because of low level (structural) factors. The magnitude of the penalty was not influenced by depth of processing or by the amount of load placed on the system. The manipulations should strongly influence higher level integrative processing. The implicit causality congruency effect was influenced by these factors indicating that it exerts its influence at a relatively high level of processing.

So then, we have been able to separate out the influence of low and high level factors using evidence both from our examination of the time course of such influences and from the above between experiment manipulations. Implicit causality and implicit consequentiality both appear to inform the anaphor resolution system during integrative processing.

8.2 Our Two-stage Account

We now return to our proposal of a two-stage model of anaphor processing outlined at the conclusion of Chapter 2. We proposed that the first stage involves co-indexation of anaphor with antecedent and is informed solely by low level factors. In other words, it behaves in a restricted manner. This stage is then followed by integrative processing linking together the information following the anaphor with what has been read previously. This stage of processing is informed by all potentially relevant factors and behaves in a non-restricted manner. In light of the data reported in this thesis, what support is there for this account?

Firstly we shall concentrate on repeat name anaphors. In our examinations of both implicit causality and implicit consequentiality we found evidence of a repeat name penalty associated with reading the sentence fragment containing that repeat name. It did not interact with verb bias. We interpret the penalty as an evidence for co-indexation. The penalty is encountered when a repeat name anaphor is used within a sentence to refer to the first mentioned character. We did not find the penalty in Experiment 5 where the information previously presented in a main-subordinate clause order sentence was presented in two separate sentences. It is possible that the penalty is restricted to within sentence reference; at least for the types of constructions we have examined. So, when a repeat name anaphor is encountered, co-indexation occurs immediately and is informed by low level information.

For pronouns appearing as grammatical subjects in subordinate clauses, we find no evidence that co-indexation occurs immediately, even under conditions where a gender contrast is present to unambiguously identify the pronominal antecedent; see Example (11) below.

(11) John fascinated Mary because she / was easily entertained.

We found no evidence on reading time to the second fragment that ambiguity interacted with the strength of the verb bias congruency effect. We did find a main effect of ambiguity on reading time to the first fragment, but this may have arisen simply as a result of it being easier to represent two characters differentiated by gender than it is to represent two characters which share this feature. So then, for pronouns appearing as grammatical subjects in subordinate clauses, it suggests co-indexation is delayed. In Chapter 3 when we outlined our 2-stage account we proposed that an overall delay in processing of a pronoun didn't violate the integrity of our model. We argued that even under conditions of delay, antecedent identification had to precede semantic integration. We reiterate this point and interpret the data associated with reading the fragment following the pronoun as measuring both stages of processing.

For pronouns appearing as grammatical subjects in main clauses, the picture may be somewhat different. We do find evidence that co-indexation is occurring as the pronoun is read, provided that information in the form of a gender contrast to unambiguously identify the pronominal referent is present. In Experiment 11 we found an interaction between ambiguity and the verb bias congruency effect on reading time to the second fragment in sentences such as (12) below. This suggests that gender information may have been used as soon as the pronoun was read to co-index the pronoun and its antecedent.

(12) Because Harold dreaded Joanna, he steadfastly refused to go back to school.

Regardless of the difference between pronouns as a function of the type of clause in which they occur, one thing does appear clear : the initial stage of co-indexation is informed by low level information (and not verb bias). At the end of Chapter 2 we proposed that structurally determined principles such as Parallel Function Strategy or Subject Assignment Strategy may inform this stage in a manner analogous to the way in which Late Closure and Minimal Attachment are considered to inform initial parsing decisions. There appears to be clear evidence for such a position in the pattern of data we find for both the Ambiguous and Unambiguous Pronoun conditions in the case of implicit causality and for the Unambiguous Pronoun conditions in the case of implicit consequentiality. We found that integration following reference to the first mentioned character in the preceding clause was always easy, regardless of verb bias. This is consistent with an initial commitment to interpret the pronoun as coreferential with the first mentioned character, as predicted by both PFS and SAS.

It is also possible however that the first mentioned character occupies some privileged position within the reader's discourse model, as Gernsbacher proposes, and it might simply be the case that it is easier to refer to that character rather than the second mentioned because of a greater saliency associated with the first mentioned participant. We did not attempt to separate a position focusing on the application of principles such as PFS and SAS from one focusing on the possible

privileged role occupied by first mentioned characters in a text which may subsequently influence processing associated with those characters. Separating these possibilities is worthy of further investigation.

So then, we propose that pronominal and repeat name anaphors are treated differently. We also propose that pronouns are treated differently depending on whether they appear as the grammatical subject in a main or subordinate clause. To summarise, repeat name anaphors are co-indexed as soon as they are encountered. Pronominal anaphors are only co-indexed when encountered if they appear in a main clause and if gender information sufficient to uniquely identify their antecedent is present. Where this gender cue is absent, co-indexation is delayed. When appearing in a subordinate clause co-indexation is always delayed.

Initial co-indexation behaves in a restricted manner informed solely by low level factors; possibly according to some principle such as Parallel Function Strategy or Subject Assignment Strategy. Even when processing may be delayed, the stage of co-indexation behaves in a restricted manner. The output from this stage is fed into the integrative stage of processing. This is informed by both verb bias information and by the semantic information following the anaphor. When these two sources of information conflict there may be a processing difficulty associated with integrating the semantic information. This is reduced however in cases where the appropriate antecedent is highly focused (e.g. first mentioned character). So then, integration following this stage of co-indexation behaves in a non-restricted manner and is informed by all information that is potentially relevant. The stage of integration is the earliest point at which both implicit causality and implicit consequentiality exert a processing influence.

References

Aaronson, D. and Ferrer, S. (1983) Lexical categories and reading tasks. *Journal of Experimental Psychology: Human perception and performance*, 9, 675-699.

Aaronson, D. and Scarborough, H.S. (1976) Performance theories for sentence coding: Some qualitative evidence. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 56-70.

Abelson, R.P. and Kanouse, D.E. (1966) Subjective acceptance of verbal generalisations. In S. Feldman (ed.) *Cognitive consistency (171-197)*. New York: Academic Press.

Anderson, A.A, Garrod, S.C. and Sanford, A.J. (1983) The accessibility of pronominal antecedents as a function of episode shifts in narrative text. *Quarterly Journal of Experimental Psychology*, 35, 427-440.

Ariel, M. (1990) *Accessing Noun-Phrase Antecedents*. Routledge.

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.

Balota, D. (1997) *Processing Pathways*. To appear in Garrod, S.C and Pickering, M.J. (eds.) *Language Processing*. UCL and MIT Press.

Bever, T.G. and Townsend, D.J. (1979) Perceptual mechanisms and formal properties of main and subordinate clauses. In W. Cooper and E.C.T. Wallace (eds.) *Sentence processing in psycholinguistic studies: Presented to Merrill Garrett*, 159-226. Hillsdale, N.J. : Erlbaum.

Bock, J.K. (1986) Syntactic persistence in language production. *Cognitive Psychology*, 18, 355-387.

Branigan, H.P., Pickering, M.J., and Stewart, A.J. (submitted) Syntactic priming of ambiguous sentences in language comprehension. Accepted subject to revision, *Language and Cognitive Processes*.

Broadbent, D.E. (1973) In defence of empirical psychology. London: Methuen.

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

Carlson, G.N. and Tanenhaus, M.K. (1988) Thematic roles and language comprehension. In W.Wilkins (ed.) *Thematic relations*, 263-300. *Syntax and Semantics Vol.21*. New York: Academic Press.

Chafe, W.L. (1972) Discourse structure and human knowledge in J.B. Carroll and R.O. Freedle (eds.) *Language comprehension and the acquisition of knowledge*. Winston and Sons, Washington D.C.

Chafe, W.L. (1974) Language and consciousness. *Language*, 50, 111-133.

Chang, F. (1980) Active memory processes in visual sentence comprehension: Clause effects and pronominal reference. *Memory and Cognition*, 8, 58-64.

Clark, H.H., and Sengul, C.J. (1979) In search of referents for nouns and pronouns. *Memory and Cognition*, 7, 35-41.

Cloitre, M. and Bever, T.G. (1988) Linguistic anaphors, levels of representation and discourse. *Language and Cognitive Processes*, 3, 293-322.

Cohen, J.D., MacWhinney, B., Flatt, M. and Provost, J. (1993) *PsyScope: A new graphic interactive environment for designing psychology experiments*. *Behavioural Research Methods, Instruments and Computers*, 25, 257-271.

Cooreman, A and Sanford, A.J. (1996) Processing subordinate and main clauses. Published in the *Proceedings of the Lancaster Conference on Anaphora*.

Corbett, A.T. and Chang, F.R. (1983) Pronoun disambiguation: accessing potential antecedents. *Memory and Cognition*, 11, 282-294.

Crawley, R.A., Stevenson, R.J. and Kleinman, D. (1990) The use of heuristic strategies in the interpretation of pronouns. *Journal of Psycholinguistic Research*, 4, 245-264.

Cutler, A. and Foss, D.J. (1977) On the role of sentence stress in sentence processing. *Language and Speech*, 21, 1-10.

Dell, G.S., McKoon, G. and Ratcliff, R. (1983) The activation of antecedent information during the processing of anaphoric reference in reading. *Journal of Verbal Learning and Verbal Behaviour*, 22, 121-132.

Ehrlich, K., and Rayner, K. (1983) Pronoun assignment and semantic integration during reading : Eye movements and immediacy of processing. *Journal of Verbal Learning and Verbal Behaviour*, 22, 75-87.

Ehrlich, K. (1980) Comprehension of pronouns. *Quarterly Journal of Experimental Psychology*, 32, 247-255.

Ferreira, F. and Clifton, C. (1986) The independence of syntactic processing. *Journal of Memory and Language*, 25, 348-368.

Fodor, J.A. (1983) *Modularity of Mind*. Cambridge, MA: Bradford.

Forster, K.I. (1979) Levels of processing and the structure of the language processor. In W. Cooper and E.C.T. Wallace (eds.) *Sentence processing in psycholinguistic studies: Presented to Merrill Garrett*, 159-226. Hillsdale, N.J. : Erlbaum.

Frazier, L. (1979) *On comprehending sentences: Syntactic parsing strategies*. Bloomington, IN: Indiana University Linguistics Club.

Frazier, L. (1987) Sentence processing: a tutorial review. In M. Coltheart (ed.) Attention and Performance XII: The psychology of reading. Hillsdale, N.J.: Erlbaum.

Frederikson, J.R. (1981) Understanding anaphora: Rules used by readers in assigning pronominal referents. *Discourse processes*, 4, 323-347.

Garnham, A. (1992) Minimalism versus constructionism : a false dichotomy in theories of inference during reading. *Psychology* 92.3.63.

Garnham, A. and Oakhill, J. (1985) On-line resolution of anaphoric pronouns: Effects of inference making and verb semantics. *British Journal of Psychology*, 76, 385-393.

Garnham, A., Oakhill, J. and Cruttenden, H. (1992) The role of implicit causality and gender cue in the interpretation of pronouns. *Language and Cognitive Processes*, 7, 231-255.

Garnham, A., Traxler, M.J., Oakhill, J. and Gernsbacher, M.A. (1996) The locus of implicit causality effects in comprehension. *Journal of Memory and Language*, 4, 517-543.

Garrod, S.C., Freudenthal, D. and Boyle, E (1994) The role of different types of anaphor in the online resolution of sentences in a discourse. *Journal of Memory and Language*, 33, 39-68

Garrod, S.C. and Sanford, A.J. (1985) On the real-time character of interpretation during reading. *Language and Cognitive Processes*, 1, 43-59.

Garvey, C. and Caramazza, A. (1974) Implicit causality in verbs. *Linguistic Inquiry*, 5, 459-464.

Garvey, C., Caramazza, A. and Yates, J. (1976) Factors underlying assignment of pronoun antecedents. *Cognition*, 3, 227-243.

Gernsbacher, M.A. (1989) Mechanism that improve referential access. *Cognition*, 32, 99-156.

Gernsbacher, M.A. and Hargreaves, D.J. (1988) Accessing Sentence Participants: The Advantage of First Mention. *Journal of Memory and Language*, 27, 699-717.

Gernsbacher, M.A., Hargreaves, D.J. and Beeman, M. (1989) Building and accessing clausal representations: The advantage of first mention versus the advantage of clause recency. *Journal of Memory and Language*, 28, 735-755.

Gibson, E. (1991) A computational theory of linguistic processing : memory limitations and processing breakdown. PhD thesis, Carnegie Melon University, Pittsburgh, PA.

Givon, T. (1983) Topic continuity in discourse: A quantitative cross-language study (Typological studies in language, Vol. 3). Amsterdam: J. Benjamins.

Glenberg, A. (1993) Comprehension while missing the point: more on minimalism and models. *Psychology*, 93.4.31.

Gordon, P.C., Grosz, B.J. and Gilliom, L.A. (1993) Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, 17, 311-347.

Gordon, P.C. and Searce, K.A. (1995) Pronominalisation and discourse coherence, discourse structure and pronoun interpretation. *Memory and Cognition*, 23, 313-323.

Greene, S.B. and McKoon, G. (1995) Telling something we can't know: Experimental approaches to verbs exhibiting implicit causality. *Psychological Science*, 5, 262-270.

Greene, S.B., McKoon, G. and Ratcliff, R. (1992) Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18, 266-283.

Grober, E.H., Beardsley, W. and Carmazza, A. (1978) Parallel function in pronoun assignment. *Cognition*, 6, 227-243.

Hankamer, J. and Sag, J.A. (1976) Deep and surface anaphora. *Linguistic Inquiry*, 7, 391-428.

Heider, F. (1958) *The psychology of interpersonal relations*. New York: Wiley.

Hirst, W. and Brill, G.A. (1980) Contextual aspects of pronoun assignment. *Journal of Verbal Learning and Verbal Behaviour*, 19, 168-175.

Hobbs, J.R. (1976) A computational approach to discourse analysis. Research report 76-2, Dept of Computer Sciences, City College, City University of New York, December 1976.

Hoffman, C. and Tchir, M.A. (1990) Interpersonal verbs and dispositional adjectives: The psychology of causality embodied in language. *Journal of Personality and Social Psychology*, 55, 530-778.

Just, M. A., and Carpenter, P.A. (1980) A theory of reading : From eye fixations to comprehension. *Psychological Review*, 87, 329-354.

Kanouse, D.E. (1972) Language, labelling and attribution. In E.E. Jones, D.E. Kanouse, H.H. Kelley, R.E. Nisbett, S. Valins and B. Weiner (eds.). *Attribution : perceiving the causes of behaviour* (121-125) Morristown, N.J. : General Learning Press.

Kelley, H.H. (1967) Attribution theory in social psychology. In D. Levine (ed.) *Nebraska Symposium on Motivation*. Lincoln: University of Nebraska Press.

Kutas, M., Van Petten, C. and Besson, M. (1988) Event-related potential asymmetries during the reading of sentences. *Electroencephalography and Clinical Neurophysiology*, 69, 218-233

Lascarides, A., Asher, N. and Oberlander, J. (1992) Inferring discourse relation in context. In Proceedings of the 30th Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, N.J., 1-8.

Leiman, J.M. (1982) A chronometric analysis of referent assignment to pronouns. Unpublished doctoral dissertation, Wayne State University, Detroit, MI.

MacDonald, M.C. and MacWhinney, B. (1990) Measuring inhibition and facilitation from pronouns. *Journal of Memory and Language*, 29, 469-492.

McDonald, J.L. and MacWhinney, B. (1995) The time course of pronoun resolution: Effects of implicit verb causality and gender. *Journal of Memory and Language*, 34, 543-566.

McKoon, G. and Ratcliff, R. (1992) Inference during reading. *Psychological Review*, 99, 440-466.

Marslen-Wilson, W.D., Levy, E. and Tyler, L.K. (1982) Producing interpretable discourse: the establishment and maintenance of reference. In R.J. Jarvella and W.Klein (eds.) *Speech, Place and Action*. Chichester: Wiley.

Marslen-Wilson, W.D. and Tyler, L.K. (1980) Towards a psychological basis for a theory of anaphora. In J. Kreiman and A.E. Ojeda (eds.) *Papers from the parasession on pronouns and anaphor* (258-286). Chicago: Chicago Linguistics Society.

Marslen-Wilson, W. Tyler, L.K., and Koster, C. (1993) Integretative processes in utterance resolution. *Journal of Memory and Language*, 32, 647-666.

Marslen-Wilson, W.D., Tyler, L.K. and Seidenberg, M. (1978) Sentence processing and the clause boundary. In W.J.M. Levelt and G.B. Flores d'Arcais (eds.) *Studies in the perception of language* (219-246). London: Wiley.

Millis, K.K. and Just, M.A. (1994) The influence of connectives on sentence comprehension. *Journal of Memory and Language*, 33, 128-147.

Minsky, M. (1975) A framework for representing knowledge. In P.H. Winston (ed.) *The Psychology of Computer Vision*. New York: McGraw-Hill.

Murdock, B.B (1962) The serial position effect in free recall. *Journal of Experimental Psychology*, 62, 482-488.

Neath, I. (1993) Distinctiveness and serial position effects in recognition. *Memory and Cognition*, 21, 689-698.

Neath, I. and Knoedler, A.J. (1994) Distinctiveness and serial position effects in recognition and sentence processing. *Journal of Memory and Language*, 33, 776-795.

Posner, M.I. and Snyder, C.R. (1975) Attention and cognitive control. In R.L. Solso (ed.) *Information processing and cognition*. Hillsdale, N.J.: Erlbaum.

Rayner, K. (1977) Visual attention in reading : Eye movements reflect cognitive processes. *Memory and Cognition*, 4, 443-448.

Rayner, K. (1978) Eye movements in reading and information processing. *Psychological Bulletin*, 85, 618-660.

Rayner, K., Carlson, M. and Frazier, L. (1983) The interaction of syntax and semantics during sentence processing. *Journal of Verbal Learning and Verbal Behaviours*, 22, 358-374.

Rudolph, U. and Forsterling, F. (1997) The psychological causality implicit in verbs: A review. *Psychological Bulletin*, 121, 192-218.

Sanford, A.J., and Garrod, S.C. (1981) Understanding written language: Explorations in comprehension beyond the sentence. Chichester, England: Wiley & Sons.

Sanford, A.J. and Garrod, S.C. (1989) What, when and how ? : Questions of immediacy in anaphoric reference resolution. *Language and Cognitive Processes*, 4, 235-262.

Sanford, A.J., Moar, K. and Garrod, S.C. (1988) Proper names as controllers of discourse focus. *Language and Speech*, 31, 43-56.

Sheldon, A.L. (1974) The role of parallel function in the acquisition of relative clauses in English. *Journal of Verbal Learning and Verbal Behaviour*, 13, 272-281.

Shillcock, R.C. (1982) The on-line comprehension of pronominal anaphora. *Language and Speech*, 25, 385-401.

Smyth, R. (1994) Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research*, 18, 5-19.

Stevenson, R.J., Crawley, R.A. and Kleinman, D. (1994) Thematic roles, focus and the representation of events. *Language and Cognitive Processes*, 9, 519-548.

Stevenson, R.J., Nelson, A.W.R. and Stenning, K. (1995) The role of parallelism in strategies of pronoun comprehension. *Language and Speech*, 38, 393-418.

Stevenson, R.J. and Vitkovitch, M. (1986) The comprehension of anaphoric relations. *Language and Speech*, 29, 335-357.

Swinney, D.A. (1979) Lexical activation during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behaviour*, 18, 681-689.

Traxler, M.J., Bybee, M.D. and Pickering, M.J. (1997) Influence of connectives on language comprehension: Eye-tracking evidence for

incremental interpretation. *Quarterly Journal of Experimental Psychology Section A - Human Experimental Psychology*, 50, 481-497.

Traxler, M.J., Sanford, A.J., Aked, J.P. and Moxey, L.M. (1997) Processing causal and diagnostic statements in discourse. *Journal of Experimental Psychology: Learning, Memory and Cognition* 23, 88-101.

Trueswell, J.C., Tanenhaus, M.K. and Garnsey, S.M. (1994) Semantic influences on parsing: Use of thematic information in syntactic disambiguation. *Journal of Memory and Language*, 33, 285-318.

Tyler, L.K. and Marslen-Wilson, W.D. (1977) The on-line effects of semantic context on syntactic processing. *Journal of Verbal Learning and Verbal Behaviour*, 16, 683-692.

Vonk, W., Hustinx, L.G.M.M. and Simons, W.H.G. (1992) The use of referential expressions in structuring discourse. *Language and Cognitive Processes*, 7, 301-333.

Appendix 1

Materials used in Experiments 2-4

Materials 1-12 contain NP1 biasing verbs.

Materials 13-24 contain NP2 biasing verbs.

Following each passage the first two questions are the 'resolving' questions used in all Experiments apart from 3. The final question in each triplet is the 'non-resolving' question used in Experiment 3.

1. Barry fascinated Derek because Barry performed magic tricks.

Barry fascinated Derek because Derek was easily entertained.

Barry fascinated Derek because he performed magic tricks.

Barry fascinated Derek because he was easily entertained.

Who performed magic tricks ?

Who was easily entertained ?

Did Barry fascinate Derek ?

2. Philip confessed to Callum because Philip had stolen the money.

Philip confessed to Callum because Callum would not be judgmental.

Philip confessed to Callum because he had stolen the money.

Philip confessed to Callum because he would not be judgmental.

Who had stolen the money ?

Who would not be judgmental ?

Was it a cheque book that had been stolen ?

3. Daniel apologised to Arnold because Daniel had been behaving selfishly.

Daniel apologised to Arnold because Arnold didn't deserve the criticism.

Daniel apologised to Arnold because he had been behaving selfishly.

Daniel apologised to Arnold because he didn't deserve the criticism.

Who had been behaving selfishly ?

Who didn't deserve the criticism ?

Did Daniel apologise to Arnold ?

4. Ann infuriated Liz because Ann had broken the promise.
Ann infuriated Liz because Liz hated being deceived.
Ann infuriated Liz because she had broken the promise.
Ann infuriated Liz because she hated being deceived.

Who had broken the promise ?
Who hated being deceived ?
Did Liz infuriate Ann ?

5. Rose disappointed Joan because Rose failed to appear.
Rose disappointed Joan because Joan had high standards.
Rose disappointed Joan because she failed to appear.
Rose disappointed Joan because she had high standards.

Who had failed to appear ?
Who had high standards ?
Did Rose disappoint Joan ?

6. Amy troubled Sue because Amy was starting to behave rather strangely.
Amy troubled Sue because Sue hated seeing others feeling very sad.
Amy troubled Sue because she was starting to behave rather strangely.
Amy troubled Sue because she hated seeing others feeling very sad.

Who was starting to behave rather strangely ?
Who hated seeing others feeling very sad ?
Did Sue trouble Amy ?

7. Henry inspired Terry because Henry had managed to beat the odds.
Henry inspired Terry because Terry needed someone to look up to.
Henry inspired Terry because he had managed to beat the odds.
Henry inspired Terry because he needed someone to look up to.

Who had managed to beat the odds ?
Who needed someone to look up to ?
Did Henry inspire Terry ?

8. Jake telephoned Luke because Jake wanted to ask a favour.
Jake telephoned Luke because Luke wouldn't remember to call.
Jake telephoned Luke because he wanted to ask a favour.
Jake telephoned Luke because he wouldn't remember to call.

Who wanted to ask a favour ?
Who wouldn't remember to call ?
Did Luke telephone Jake ?

9. James amused Craig because James performed hilarious impressions.
James amused Craig because Craig was very easily entertained.
James amused Craig because he performed hilarious impressions.
James amused Craig because he was very easily entertained.

Who performed hilarious impressions ?
Who was very easily entertained ?
Were the impressions hilarious ?

10. Susan concerned Carol because Susan was starting to behave erratically.
Susan concerned Carol because Carol hated seeing friends in trouble.
Susan concerned Carol because she was starting to behave erratically.
Susan concerned Carol because she hated seeing friends in trouble.

Who was starting to behave erratically ?
Who hated seeing friends in trouble ?
Was the behaviour predictable ?

11. Caroline amazed Florence because Caroline passed the exam.
Caroline amazed Florence because Florence was easily impressed.
Caroline amazed Florence because she passed the exam.
Caroline amazed Florence because she was easily impressed.

Who passed the exam ?
Who was easily impressed ?
Did Caroline amaze Florence ?

12. Diana called Nancy because Diana had found the telephone number.

Diana called Nancy because Nancy couldn't make outgoing calls.

Diana called Nancy because she had found the telephone number.

Diana called Nancy because she couldn't make outgoing calls.

Who had found the telephone number ?

Who couldn't make outgoing calls ?

Did Nancy call Diana ?

13. Jack liked Tony because Jack was made to feel quite at home.

Jack liked Tony because Tony was full of incredibly helpful advice.

Jack liked Tony because he was made to feel quite at home.

Jack liked Tony because he was full of incredibly helpful advice.

Who was made to feel quite at home ?

Who was full of incredibly helpful advice ?

Was the advice incredibly helpful ?

14. Michael appreciated Richard because Michael needed the extra help.

Michael appreciated Richard because Richard had offered to help.

Michael appreciated Richard because he needed the extra help.

Michael appreciated Richard because he had offered to help.

Who needed the extra help ?

Who had offered to help ?

Did Richard appreciate Michael ?

15. Trevor detested Gordon because Trevor hated being taken advantage of.

Trevor detested Gordon because Gordon was completely unreliable.

Trevor detested Gordon because he hated being taken advantage of.

Trevor detested Gordon because he was completely unreliable.

Who hated being taken advantage of ?

Who was completely unreliable ?

Did Trevor detest Gordon ?

16. Jean congratulated Rita because Jean was very impressed.
Jean congratulated Rita because Rita had won the championship.
Jean congratulated Rita because she was very impressed.
Jean congratulated Rita because she had won the championship.

Was was very impressed ?
Who had won the championship ?
Did Rita congratulate Jean ?

17. Catherine despised Elizabeth because Catherine had felt very let down.
Catherine despised Elizabeth because Elizabeth seemed to lie constantly.
Catherine despised Elizabeth because she had felt very let down.
Catherine despised Elizabeth because she seemed to lie constantly.

Who had felt very let down ?
Who seemed to lie constantly ?
Did Catherine despise Elizabeth ?

18. Cathy thanked Sally because Cathy had appreciated the present.
Cathy thanked Sally because Sally had brought the present.
Cathy thanked Sally because she had appreciated the present.
Cathy thanked Sally because she had brought the present.

Who had appreciated the present ?
Who had brought the present ?
Did Sally thank Cathy ?

19. Mick loathed Paul because Mick was starting to feel upstaged.
Mick loathed Paul because Paul had very little integrity.
Mick loathed Paul because he was starting to feel upstaged.
Mick loathed Paul because he had very little integrity.

Who was starting to feel upstaged ?
Who had very little integrity ?
Did Mick loathe Paul ?

20. Ted praised Bob because Ted was impressed by the project.
Ted praised Bob because Bob behaved very courageously.
Ted praised Bob because he was impressed by the project.
Ted praised Bob because he behaved very courageously.

Who was impressed by the project ?
Who behaved very courageously ?
Was the project unimpressive ?

21. Ray scolded Rob because Ray was aware of the potential danger.
Ray scolded Rob because Rob had damaged the mahogany table.
Ray scolded Rob because he was aware of the potential danger.
Ray scolded Rob because he had damaged the mahogany table.

Who was aware of the potential danger ?
Who had damaged the mahogany table ?
Did Ray scold Rob ?

22. Anna noticed Emma because Anna was always exceedingly observant.
Anna noticed Emma because Emma wore a remarkably colourful dress.
Anna noticed Emma because she was always exceedingly observant.
Anna noticed Emma because she wore a remarkably colourful dress.

Who was always exceedingly observant ?
Who wore a remarkably colourful dress ?
Did Emma notice Anna ?

23. Gemma punished Ellen because Gemma had had enough.
Gemma punished Ellen because Ellen had been very trying.
Gemma punished Ellen because she had had enough.
Gemma punished Ellen because she had been very trying.

Who had had enough ?
Who had been very trying ?
Did Gemma punish Ellen ?

24. Emily admired Tracy because Emily needed a role model.
Emily admired Tracy because Tracy was very motivated.
Emily admired Tracy because she needed a role model.
Emily admired Tracy because she was very motivated.

Who needed a role model ?
Who was very motivated ?
Did Tracy admire Emily ?

Appendix 2

Materials used in Experiment 5

Materials 1-12 contain NP1 biasing verbs.

Materials 13-24 contain NP2 biasing verbs.

1. Barry fascinated Derek. This was because Barry performed magic tricks.

Barry fascinated Derek. This was because Derek was easily entertained.

Barry fascinated Derek. This was because he performed magic tricks.

Barry fascinated Derek. This was because he was easily entertained.

Who performed magic tricks ?

Who was easily entertained ?

2. Philip confessed to Callum. This was because Philip had stolen the money.

Philip confessed to Callum. This was because Callum would not be judgmental.

Philip confessed to Callum. This was because he had stolen the money.

Philip confessed to Callum. This was because he would not be judgmental.

Who had stolen the money ?

Who would not be judgmental ?

3. Daniel apologised to Arnold. This was because Daniel had been behaving selfishly.

Daniel apologised to Arnold. This was because Arnold didn't deserve the criticism.

Daniel apologised to Arnold. This was because he had been behaving selfishly.

Daniel apologised to Arnold. This was because he didn't deserve the criticism.

Who had been behaving selfishly ?

Who didn't deserve the criticism ?

4. Ann infuriated Liz. This was because Ann had broken the promise.

Ann infuriated Liz. This was because Liz hated being deceived.

Ann infuriated Liz. This was because she had broken the promise.

Ann infuriated Liz. This was because she hated being deceived.

Who had broken the promise ?

Who hated being deceived ?

5. Rose disappointed Joan. This was because Rose failed to appear.

Rose disappointed Joan. This was because Joan had high standards.

Rose disappointed Joan. This was because she failed to appear.

Rose disappointed Joan. This was because she had high standards.

Who had failed to appear ?

Who had high standards ?

6. Amy troubled Sue. This was because Amy was starting to behave rather strangely.

Amy troubled Sue. This was because Sue hated seeing others feeling very sad.

Amy troubled Sue. This was because she was starting to behave rather strangely.

Amy troubled Sue. This was because she hated seeing others feeling very sad.

Who was starting to behave rather strangely ?

Who hated seeing others feeling very sad ?

7. Henry inspired Terry. This was because Henry had managed to beat the odds.

Henry inspired Terry. This was because Terry needed someone to look up to.

Henry inspired Terry. This was because he had managed to beat the odds.

Henry inspired Terry. This was because he needed someone to look up to.

Who had managed to beat the odds ?

Who needed someone to look up to ?

8. Jake telephoned Luke. This was because Jake wanted to ask a favour.

Jake telephoned Luke. This was because Luke wouldn't remember to call.

Jake telephoned Luke. This was because he wanted to ask a favour.

Jake telephoned Luke. This was because he wouldn't remember to call.

Who wanted to ask a favour ?

Who wouldn't remember to call ?

9. James amused Craig. This was because James performed hilarious impressions.

James amused Craig. This was because Craig was very easily entertained.

James amused Craig. This was because he performed hilarious impressions.

James amused Craig. This was because he was very easily entertained.

Who performed hilarious impressions ?

Who was very easily entertained ?

10. Susan concerned Carol. This was because Susan was starting to behave erratically.

Susan concerned Carol. This was because Carol hated seeing friends in trouble.

Susan concerned Carol. This was because she was starting to behave erratically.

Susan concerned Carol. This was because she hated seeing friends in trouble.

Who was starting to behave erratically ?

Who hated seeing friends in trouble ?

11. Caroline amazed Florence. This was because Caroline passed the exam.

Caroline amazed Florence. This was because Florence was easily impressed.

Caroline amazed Florence. This was because she passed the exam.

Caroline amazed Florence. This was because she was easily impressed.

Who passed the exam ?

Who was easily impressed ?

12. Diana called Nancy. This was because Diana had found the telephone number.

Diana called Nancy. This was because Nancy couldn't make outgoing calls.

Diana called Nancy. This was because she had found the telephone number.

Diana called Nancy. This was because she couldn't make outgoing calls.

Who had found the telephone number ?

Who couldn't make outgoing calls ?

13. Jack liked Tony. This was because Jack was made to feel quite at home.

Jack liked Tony. This was because Tony was full of incredibly helpful advice.

Jack liked Tony. This was because he was made to feel quite at home.

Jack liked Tony. This was because he was full of incredibly helpful advice.

Who was made to feel quite at home ?

Who was full of incredibly helpful advice ?

14. Michael appreciated Richard. This was because Michael needed the extra help.

Michael appreciated Richard. This was because Richard had offered to help.

Michael appreciated Richard. This was because he needed the extra help.

Michael appreciated Richard. This was because he had offered to help.

Who needed the extra help ?

Who had offered to help ?

15. Trevor detested Gordon. This was because Trevor hated being taken advantage of.

Trevor detested Gordon. This was because Gordon was completely unreliable.

Trevor detested Gordon. This was because he hated being taken advantage of.

Trevor detested Gordon. This was because he was completely unreliable.

Who hated being taken advantage of ?

Who was completely unreliable ?

16. Jean congratulated Rita. This was because Jean was very impressed.

Jean congratulated Rita. This was because Rita had won the championship.

Jean congratulated Rita. This was because she was very impressed.

Jean congratulated Rita. This was because she had won the championship.

Was was very impressed ?

Who had won the championship ?

17. Catherine despised Elizabeth. This was because Catherine had felt very let down.

Catherine despised Elizabeth. This was because Elizabeth seemed to lie constantly.

Catherine despised Elizabeth. This was because she had felt very let down.

Catherine despised Elizabeth. This was because she seemed to lie constantly.

Who had felt very let down ?

Who seemed to lie constantly ?

18. Cathy thanked Sally. This was because Cathy had appreciated the present.

Cathy thanked Sally. This was because Sally had brought the present.

Cathy thanked Sally. This was because she had appreciated the present.

Cathy thanked Sally. This was because she had brought the present.

Who had appreciated the present ?

Who had brought the present ?

19. Mick loathed Paul. This was because Mick was starting to feel upstaged.

Mick loathed Paul. This was because Paul had very little integrity.

Mick loathed Paul. This was because he was starting to feel upstaged.

Mick loathed Paul. This was because he had very little integrity.

Who was starting to feel upstaged ?

Who had very little integrity ?

20. Ted praised Bob. This was because Ted was impressed by the project.

Ted praised Bob. This was because Bob behaved very courageously.

Ted praised Bob. This was because he was impressed by the project.

Ted praised Bob. This was because he behaved very courageously.

Who was impressed by the project ?

Who behaved very courageously ?

21. Ray scolded Rob. This was because Ray was aware of the potential danger.

Ray scolded Rob. This was because Rob had damaged the mahogany table.

Ray scolded Rob. This was because he was aware of the potential danger.

Ray scolded Rob. This was because he had damaged the mahogany table.

Who was aware of the potential danger ?

Who had damaged the mahogany table ?

22. Anna noticed Emma. This was because Anna was always exceedingly observant.

Anna noticed Emma. This was because Emma wore a remarkably colourful dress.

Anna noticed Emma. This was because she was always exceedingly observant.

Anna noticed Emma. This was because she wore a remarkably colourful dress.

Who was always exceedingly observant ?

Who wore a remarkably colourful dress ?

23. Gemma punished Ellen. This was because Gemma had had enough.

Gemma punished Ellen. This was because Ellen had been very trying.

Gemma punished Ellen. This was because she had had enough.

Gemma punished Ellen. This was because she had been very trying.

Who had had enough ?

Who had been very trying ?

24. Emily admired Tracy. This was because Emily needed a role model.

Emily admired Tracy. This was because Tracy was very motivated.

Emily admired Tracy. This was because she needed a role model.

Emily admired Tracy. This was because she was very motivated.

Who needed a role model ?

Who was very motivated ?

Appendix 3

Materials used in Experiment 6

Materials 1-12 contain NP1 biasing verbs.

Materials 13-24 contain NP2 biasing verbs.

1. Barry fascinated Derek because he performed magic tricks.
Barry fascinated Lydia because he performed magic tricks.
Lydia fascinated Derek because he was easily entertained.
Barry fascinated Derek because he was easily entertained.

Who performed magic tricks ?

Who was easily entertained ?

2. Philip confessed to Callum because he had stolen the money.
Philip confessed to Tracey because he had stolen the money.
Tracey confessed to Callum because he would not be judgmental.
Philip confessed to Callum because he would not be judgmental.

Who had stolen the money ?

Who would not be judgmental ?

3. Daniel apologised to Arnold because he had been behaving selfishly.
Daniel apologised to Joanne because he had been behaving selfishly.
Joanne apologised to Arnold because he didn't deserve the criticism.
Daniel apologised to Arnold because he didn't deserve the criticism.

Who had been behaving selfishly ?

Who didn't deserve the criticism ?

4. Ann infuriated Liz because she had broken the promise.
Ann infuriated Ben because she had broken the promise.
Ben infuriated Liz because she hated being deceived.
Ann infuriated Liz because she hated being deceived.

Who had broken the promise ?

Who hated being deceived ?

5. Rose disappointed Joan because she failed to appear.
Rose disappointed John because she failed to appear.
John disappointed Joan because she had high standards.
Rose disappointed Joan because she had high standards.

Who had failed to appear ?
Who had high standards ?

6. Amy troubled Sue because she was starting to behave rather strangely.
Amy troubled Tim because she was starting to behave rather strangely.
Tim troubled Sue because she hated seeing others feeling very sad.
Amy troubled Sue because she hated seeing others feeling very sad.

Who was starting to behave rather strangely ?
Who hated seeing others feeling very sad ?

7. Henry inspired Terry because he had managed to beat the odds.
Henry inspired Carol because he had managed to beat the odds.
Carol inspired Terry because he needed someone to look up to.
Henry inspired Terry because he needed someone to look up to.

Who had managed to beat the odds ?
Who needed someone to look up to ?

8. Jake telephoned Luke because he wanted to ask a favour.
Jake telephoned Nancy because he wanted to ask a favour.
Nancy telephoned Luke because he wouldn't remember to call.
Jake telephoned Luke because he wouldn't remember to call.

Who wanted to ask a favour ?
Who wouldn't remember to call ?

9. James amused Craig because he performed hilarious impressions.
James amused Donna because he performed hilarious impressions.
Donna amused Craig because he was very easily entertained.

James amused Craig because he was very easily entertained.

Who performed hilarious impressions ?

Who was very easily entertained ?

10. Susan concerned Carol because she was starting to behave erratically.

Terry concerned Carol because she was starting to behave erratically.

Susan concerned Terry because she hated seeing friends in trouble.

Susan concerned Carol because she hated seeing friends in trouble.

Who was starting to behave erratically ?

Who hated seeing friends in trouble ?

11. Caroline amazed Florence because she passed the exam.

Jonathon amazed Florence because she passed the exam.

Caroline amazed Jonathon because she was easily impressed.

Caroline amazed Florence because she was easily impressed.

Who passed the exam ?

Who was easily impressed ?

12. Diana called Nancy because she had found the telephone number.

Craig called Nancy because she had found the telephone number.

Diana called Craig because she couldn't make outgoing calls.

Diana called Nancy because she couldn't make outgoing calls.

Who had found the telephone number ?

Who couldn't make outgoing calls ?

13. Jack liked Tony because he was made to feel quite at home.

Jack liked Mary because he was made to feel quite at home.

Mary liked Tony because he was full of incredibly helpful advice.

Jack liked Tony because he was full of incredibly helpful advice.

Who was made to feel quite at home ?

Who was full of incredibly helpful advice ?

14. Michael appreciated Richard because he needed the extra help.
Michael appreciated Kathryn because he needed the extra help.
Kathryn appreciated Richard because he had offered to help.
Michael appreciated Richard because he had offered to help.

Who needed the extra help ?

Who had offered to help ?

15. Trevor detested Gordon because he hated being taken advantage of.

Trevor detested Phoebe because he hated being taken advantage of.

Phoebe detested Gordon because he was completely unreliable.

Trevor detested Gordon because he was completely unreliable.

Who hated being taken advantage of ?

Who was completely unreliable ?

16. Jean congratulated Rita because she was very impressed.

Jean congratulated John because she was very impressed.

John congratulated Rita because she had won the championship.

Jean congratulated Rita because she had won the championship.

Was was very impressed ?

Who had won the championship ?

17. Catherine despised Elizabeth because she had felt very let down.

Catherine despised Jonathon because she had felt very let down.

Jonathon despised Elizabeth because she seemed to lie constantly.

Catherine despised Elizabeth because she seemed to lie constantly.

Who had felt very let down ?

Who seemed to lie constantly ?

18. Cathy thanked Sally because she had appreciated the present.

Cathy thanked Harry because she had appreciated the present.

Harry thanked Sally because she had brought the present.

Cathy thanked Sally because she had brought the present.

Who had appreciated the present ?

Who had brought the present ?

19. Mick loathed Paul because he was starting to feel upstaged.

Mick loathed Lisa because he was starting to feel upstaged.

Lisa loathed Paul because he had very little integrity.

Mick loathed Paul because he had very little integrity.

Who was starting to feel upstaged ?

Who had very little integrity ?

20. Ted praised Bob because he was impressed by the project.

Ted praised Sue because he was impressed by the project.

Sue praised Bob because he behaved very courageously.

Ted praised Bob because he behaved very courageously.

Who was impressed by the project ?

Who behaved very courageously ?

21. Ray scolded Rob because he was aware of the potential danger.

Ray scolded Liz because he was aware of the potential danger.

Liz scolded Rob because he had damaged the mahogany table.

Ray scolded Rob because he had damaged the mahogany table.

Who was aware of the potential danger ?

Who had damaged the mahogany table ?

22. Anna noticed Emma because she was always exceedingly observant.

Anna noticed Bill because she was always exceedingly observant.

Bill noticed Emma because she wore a remarkably colourful dress.

Anna noticed Emma because she wore a remarkably colourful dress.

Who was always exceedingly observant ?

Who wore a remarkably colourful dress ?

23. Gemma punished Ellen because she had had enough.

Gemma punished Paul because she had had enough.

Paul punished Ellen because she had been very trying.
Gemma punished Ellen because she had been very trying.

Who had had enough ?
Who had been very trying ?

24. Emily admired Tracy because she needed a role model.
Emily admired Grant because she needed a role model.
Grant admired Tracy because she was very motivated.
Emily admired Tracy because she was very motivated.

Who needed a role model ?
Who was very motivated ?

Appendix 4

Materials used in Experiment 7

Materials 1-12 contain NP1 biasing verbs.

Materials 13-24 contain NP2 biasing verbs.

1. Barry fascinated Derek because he performed magic tricks, according to his father.

Barry fascinated Derek because he was easily entertained, according to his father.

Barry fascinated Sarah because he performed magic tricks, according to his father.

Sarah fascinated Derek because he was easily entertained, according to his father.

2. Roy infuriated Ian because he had broken the promise, according to our mutual friend.

Roy infuriated Ian because he hated being deceived, according to our mutual friend.

Roy infuriated Liz because he had broken the promise, according to our mutual friend.

Ann infuriated Ian because he hated being deceived, according to our mutual friend.

3. Rose disappointed Joan because she failed to appear, or at least so I heard several days later.

Rose disappointed Joan because she had high standards, or at least so I heard several days later.

Rose disappointed Doug because she failed to appear, or at least so I heard several days later.

Doug disappointed Joan because she had high standards, or at least so I heard several days later.

4. Philip confessed to Callum because he had stolen the money, or so everyone else had believed.

Philip confessed to Callum because he would not be judgmental, or so everyone else had believed.

Philip confessed to Joanne because he had stolen the money, or so everyone else had believed.

Joanne confessed to Callum because he would not be judgmental, or so everyone else had believed.

5. Daniel apologised to Arnold because he had been behaving selfishly, as we had all thought for some time.

Daniel apologised to Arnold because he didn't deserve the criticism, as we had all thought for some time.

Daniel apologised to Sylvia because he had been behaving selfishly, as we had all thought for some time.

Sylvia apologised to Arnold because he didn't deserve the criticism, as we had all thought for some time.

6. Henry inspired Terry because he had managed to beat the odds, much to our general surprise.

Henry inspired Terry because he needed someone to look up to, much to our general surprise.

Henry inspired Emily because he had managed to beat the odds, much to our general surprise.

Emily inspired Terry because he needed someone to look up to, much to our general surprise.

7. Amy troubled Sue because she was starting to behave rather strangely, or so I thought at the time.

Amy troubled Sue because she hated seeing others feeling very sad, or so I thought at the time.

Amy troubled Roy because she was starting to behave rather strangely, or so I thought at the time.

Roy troubled Sue because she hated seeing others feeling very sad, or so I thought at the time.

8. Jean telephoned Rita because she wanted to ask a favour, for one reason or another.

Jean telephoned Rita because she wouldn't remember to call, for one reason or another.

Jean telephoned Rick because she wanted to ask a favour, for one reason or another.

Rick telephoned Rita because she wouldn't remember to call, for one reason or another.

9. James amused Craig because he performed hilarious impressions, according to my favourite cousin.

James amused Craig because he was very easily entertained, according to my favourite cousin.

James amused Carol because he performed hilarious impressions, according to my favourite cousin.

Carol amused Craig because he was very easily entertained, according to my favourite cousin.

10. Susan concerned Carol because she was starting to behave erratically, as I had always thought.

Susan concerned Carol because she hated seeing friends in trouble, as I had always thought.

Susan concerned Keith because she was starting to behave erratically, as I had always thought.

Grant concerned Carol because she hated seeing friends in trouble, as I had always thought.

11. Caroline amazed Florence because she passed the exam, much to everybody's surprise.

Caroline amazed Florence because she was easily impressed, much to everybody's surprise.

Caroline amazed Alasdair because she passed the exam, much to everybody's surprise.

Geoffrey amazed Florence because she was easily impressed, much to everybody's surprise.

12. Diana called Nancy because she had found the telephone number, or so I had guessed.

Diana called Nancy because she couldn't make outgoing calls, or so I had guessed.

Diana called Steve because she had found the telephone number, or so I had guessed.

Lloyd called Nancy because she couldn't make outgoing calls, or so I had guessed.

13. Catherine despised Elizabeth because she had felt very let down, for one reason or another.

Catherine despised Elizabeth because she seemed to lie constantly, for one reason or another.

Catherine despised Sebastian because she had felt very let down, for one reason or another.

Alexander despised Elizabeth because she seemed to lie constantly, for one reason or another.

14. Jack liked Tony because he was made to feel quite at home, according to my mother.

Jack liked Tony because he was full of incredibly helpful advice, according to my mother.

Jack liked Nell because he was made to feel quite at home, according to my mother.

Jane liked Tony because he was full of incredibly helpful advice, according to my mother.

15. Michael appreciated Richard because he needed the extra help, according to my friend.

Michael appreciated Richard because he had offered to help, according to my friend.

Michael appreciated Kathryn because he needed the extra help, according to my friend.

Suzanne appreciated Richard because he had offered to help, according to my friend.

16. Jake congratulated Luke because he was very impressed, contrary to all our expectations.

Jake congratulated Luke because he had won the championship, contrary to all our expectations.

Jake congratulated Dora because he was very impressed, contrary to all our expectations.

Jill congratulated Luke because he had won the championship, contrary to all our expectations.

17. Cathy thanked Sally because she had appreciated the present, or so my brother told me.

Cathy thanked Sally because she had brought the present, or so my brother told me.

Cathy thanked Simon because she had appreciated the present, or so my brother told me.

Gerry thanked Sally because she had brought the present, or so my brother told me.

18. Trevor detested Gordon because he hated being taken advantage of, which came as no surprise to me.

Trevor detested Gordon because he was completely unreliable, which came as no surprise to me.

Trevor detested Carole because he hated being taken advantage of, which came as no surprise to me.

Melony detested Gordon because he was completely unreliable, which came as no surprise to me.

19. Mick loathed Paul because he was starting to feel upstaged, according to my neighbour.

Mick loathed Paul because he had very little integrity., according to my neighbour.

Mick loathed Lisa because he was starting to feel upstaged, according to my neighbour.

Anne loathed Paul because he had very little integrity, according to my neighbour.

20. Ted praised Bob because he was impressed by the project, according to my informant.

Ted praised Bob because he behaved very courageously, according to my informant.

Ted praised Eve because he was impressed by the project, according to my informant.

Sue praised Bob because he behaved very courageously, according to my informant.

21. Anna noticed Emma because she was always exceedingly observant, or so I had been told.

Anna noticed Emma because she wore a remarkably colourful dress, or so I had been told.

Anna noticed Noel because she was always exceedingly observant, or so I had been told.

John noticed Emma because she wore a remarkably colourful dress, or so I had been told.

22. Ray scolded Rob because he was aware of the potential danger, or so I heard the next day.

Ray scolded Rob because he had damaged the mahogany table, or so I heard the next day.

Ray scolded Meg because he was aware of the potential danger, or so I heard the next day.

Jan scolded Rob because he had damaged the mahogany table, or so I heard the next day.

23. Terry punished Ellen because she had had enough, which I could easily imagine.

Terry punished Ellen because she had been very trying, which I could easily imagine.

Terry punished Kevin because she had had enough, which I could easily imagine.

Donna punished Ellen because she had been very trying, which I could easily imagine.

24. Gemma admired Tracy because she needed a role model, or so I had been led to believe.

Gemma admired Tracy because she was very motivated, or so I had been led to believe.

Gemma admired Peter because she needed a role model, or so I had been led to believe.

David admired Tracy because she was very motivated, or so I had been led to believe.

Appendix 5

Materials used in Experiment 8

Materials 1-12 contain NP1 biasing verbs.

Materials 13-24 contain NP2 biasing verbs.

1. Barry fascinated Derek lots and lots because he performed magic tricks.

Barry fascinated Derek lots and lots because he was easily entertained.

Barry fascinated Derek lots and lots because Barry performed magic tricks.

Barry fascinated Derek lots and lots because Derek was easily entertained.

Who performed magic tricks ?

Who was easily entertained ?

2. Ann infuriated Liz greatly because she had broken the promise.

Ann infuriated Liz greatly because she hated being deceived.

Ann infuriated Liz greatly because Ann had broken the promise.

Ann infuriated Liz greatly because Liz hated being deceived.

Who had broken the promise ?

Who hated being deceived ?

3. Rose disappointed Joan deeply because she failed to appear.

Rose disappointed Joan deeply because she had high standards.

Rose disappointed Joan deeply because Rose failed to appear.

Rose disappointed Joan deeply because Joan had high standards.

Who had failed to appear ?

Who had high standards ?

4. Philip confessed to Callum right away because he had stolen the money.

Philip confessed to Callum right away because he would not be judgmental.

Philip confessed to Callum right away because Philip had stolen the money.

Philip confessed to Callum right away because Callum would not be judgmental.

Who had stolen the money ?

Who would not be judgmental ?

5. Daniel apologised to Arnold profusely because he had been behaving selfishly.

Daniel apologised to Arnold profusely because he didn't deserve the criticism.

Daniel apologised to Arnold profusely because Daniel had been behaving selfishly.

Daniel apologised to Arnold profusely because Arnold didn't deserve the criticism.

Who had been behaving selfishly ?

Who didn't deserve the criticism ?

6. Henry inspired Terry intensely because he had managed to beat the odds.

Henry inspired Terry intensely because he needed someone to look up to.

Henry inspired Terry intensely because Henry had managed to beat the odds.

Henry inspired Terry intensely because Terry needed someone to look up to.

Who had managed to beat the odds ?

Who needed someone to look up to ?

7. Amy troubled Sue enormously because she was starting to behave rather strangely.

Amy troubled Sue enormously because she hated seeing others feeling very sad.

Amy troubled Sue enormously because Amy was starting to behave rather strangely.

Amy troubled Sue enormously because Sue hated seeing others feeling very sad.

Who was starting to behave rather strangely ?

Who hated seeing others feeling very sad ?

8. Jean telephoned Rita straight away because she wanted to ask a favour.

Jean telephoned Rita straight away because she wouldn't remember to call.

Jean telephoned Rita straight away because Jean wanted to ask a favour.

Jean telephoned Rita straight away because Rita wouldn't remember to call.

Who wanted to ask a favour ?

Who wouldn't remember to call ?

9. James amused Craig a lot because he performed hilarious impressions.

James amused Craig a lot because he was very easily entertained.

James amused Craig a lot because James performed hilarious impressions.

James amused Craig a lot because Craig was very easily entertained.

Who performed hilarious impressions ?

Who was very easily entertained ?

10. Susan concerned Carol a great deal because she was starting to behave erratically.

Susan concerned Carol a great deal because she hated seeing friends in trouble.

Susan concerned Carol a great deal because Susan was starting to behave erratically.

Susan concerned Carol a great deal because Carol hated seeing friends in trouble.

Who was starting to behave erratically ?

Who hated seeing friends in trouble ?

11. Caroline amazed Florence completely because she passed the exam.

Caroline amazed Florence completely because she was easily impressed.

Caroline amazed Florence completely because Caroline passed the exam.

Caroline amazed Florence completely because Florence was easily impressed.

Who passed the exam ?

Who was easily impressed ?

12. Diana called Nancy quickly because she had found the telephone number.

Diana called Nancy quickly because she couldn't make outgoing calls.

Diana called Nancy quickly because Diana had found the telephone number.

Diana called Nancy quickly because Nancy couldn't make outgoing calls.

Who had found the telephone number ?

Who couldn't make outgoing calls ?

13. Catherine despised Elizabeth passionately because she had felt very let down.

Catherine despised Elizabeth passionately because she seemed to lie constantly.

Catherine despised Elizabeth passionately because Catherine had felt very let down.

Catherine despised Elizabeth passionately because Elizabeth seemed to lie constantly.

Who had felt very let down ?

Who seemed to lie constantly ?

14. Jack liked Tony considerably because he was made to feel quite at home.

Jack liked Tony considerably because he was full of incredibly helpful advice.

Jack liked Tony considerably because Jack was made to feel quite at home.

Jack liked Tony considerably because Tony was full of incredibly helpful advice.

Who was made to feel quite at home ?

Who was full of incredibly helpful advice ?

15. Michael appreciated Richard very much because he needed the extra help.

Michael appreciated Richard very much because he had offered to help.

Michael appreciated Richard very much because Michael needed the extra help.

Michael appreciated Richard very much because Richard had offered to help.

Who needed the extra help ?

Who had offered to help ?

16. Jake congratulated Luke vigorously because he was very impressed.

Jake congratulated Luke vigorously because he had won the championship.

Jake congratulated Luke vigorously because Jake was very impressed.

Jake congratulated Luke vigorously because Luke had won the championship.

Was was very impressed ?

Who had won the championship ?

17. Cathy thanked Sally wholeheartedly because she had appreciated the present.

Cathy thanked Sally wholeheartedly because she had brought the present.

Cathy thanked Sally wholeheartedly because Cathy had appreciated the present.

Cathy thanked Sally wholeheartedly because Sally had brought the present.

Who had appreciated the present ?

Who had brought the present ?

18. Trevor detested Gordon utterly because he hated being taken advantage of.

Trevor detested Gordon utterly because he was completely unreliable.

Trevor detested Gordon utterly because Trevor hated being taken advantage of.

Trevor detested Gordon utterly because Gordon was completely unreliable.

Who hated being taken advantage of ?

Who was completely unreliable ?

19. Mick loathed Paul thoroughly because he was starting to feel upstaged.

Mick loathed Paul thoroughly because he had very little integrity.

Mick loathed Paul thoroughly because Mick was starting to feel upstaged.

Mick loathed Paul thoroughly because Paul had very little integrity.

Who was starting to feel upstaged ?

Who had very little integrity ?

20. Ted praised Bob enthusiastically because he was impressed by the project.

Ted praised Bob enthusiastically because he behaved very courageously.

Ted praised Bob enthusiastically because Ted was impressed by the project.

Ted praised Bob enthusiastically because Bob behaved very courageously.

Who was impressed by the project ?

Who behaved very courageously ?

21. Anna noticed Emma at once because she was always exceedingly observant.

Anna noticed Emma at once because she wore a remarkably colourful dress.

Anna noticed Emma at once because Anna was always exceedingly observant.

Anna noticed Emma at once because Emma wore a remarkably colourful dress.

Who was always exceedingly observant ?

Who wore a remarkably colourful dress ?

22. Ray scolded Rob severely because he was aware of the potential danger.

Ray scolded Rob severely because he had damaged the mahogany table.

Ray scolded Rob severely because Ray was aware of the potential danger.

Ray scolded Rob severely because Rob had damaged the mahogany table.

Who was aware of the potential danger ?

Who had damaged the mahogany table ?

23. Gemma punished Ellen harshly because she had had enough.

Gemma punished Ellen harshly because she had been very trying.

Gemma punished Ellen harshly because Gemma had had enough.

Gemma punished Ellen harshly because Ellen had been very trying.

Who had had enough ?

Who had been very trying ?

24. Emily admired Tracy unreservedly because she needed a role model.

Emily admired Tracy unreservedly because she was very motivated.

Emily admired Tracy unreservedly because Emily needed a role model.

Emily admired Tracy unreservedly because Tracy was very motivated.

Who needed a role model ?

Who was very motivated ?

Appendix 6

Materials used in Experiment 10

Materials 1-8 contain NP1 biasing verbs.

Materials 9-16 contain NP2 biasing verbs.

1. Because Harold dreaded Justin, Harold steadfastly refused to go back to school.

Because Harold dreaded Justin, Justin was told to try acting less aggressively.

Because Harold dreaded Justin, he steadfastly refused to go back to school.

Because Harold dreaded Justin, he was told to try acting less aggressively.

Who refused to go back to school ?

Who was told to act less aggressively ?

2. Because Geoff appreciated Simon, Geoff thanked for all the hard effort.

Because Geoff appreciated Simon, Simon felt the work had been worth it.

Because Geoff appreciated Simon, he thanked for all the hard effort.

Because Geoff appreciated Simon, he felt the work had been worth it.

Who thanked for the effort ?

Who felt the work had been worth it ?

3. Because Eleanor adored Valerie, Eleanor baked a special birthday cake.

Because Eleanor adored Valerie, Valerie became very popular at school.

Because Eleanor adored Valerie, she baked a special birthday cake.

Because Eleanor adored Valerie, she became very popular at school.

Who became popular at school ?

Who baked a special birthday cake ?

4. Because Anne admired Joan, Anne was impressed by what had been said.

Because Anne admired Joan, Joan finally started to feel appreciated.

Because Anne admired Joan, she was impressed by what had been said.

Because Anne admired Joan, she finally started to feel appreciated.

Who started to feel appreciated ?

Who was impressed by what had been said ?

5. Because Deborah loathed Suzanne, Deborah was punished for being spiteful.

Because Deborah loathed Suzanne, Suzanne felt rather abandoned at school.

Because Deborah loathed Suzanne, she was punished for being spiteful.

Because Deborah loathed Suzanne, she felt rather abandoned at school.

Who was punished ?

Who felt abandoned at school ?

6. Because Mick resented Ross, Mick was thought to be of a jealous nature.

Because Mick resented Ross, Ross was starting to feel quite persecuted.

Because Mick resented Ross, he was thought to be of a jealous nature.

Because Mick resented Ross, he was starting to feel quite persecuted.

Who was thought to be of a jealous nature ?

Who was starting to feel persecuted ?

7. Because Roger despised Terry, Roger was considered to be unfriendly.

Because Roger despised Terry, Terry began to feel utterly miserable.

Because Roger despised Terry, he was considered to be unfriendly.

Because Roger despised Terry, he began to feel utterly miserable.

Who began to feel utterly miserable ?

Who was considered to be unfriendly ?

8. Because Sarah liked Nikki, Sarah obviously did not hate everybody.

Because Sarah liked Nikki, Nikki felt popular at the dinner party.

Because Sarah liked Nikki, she obviously did not hate everybody.

Because Sarah liked Nikki, she felt popular at the dinner party.

Who felt popular at the dinner party ?

Who did not hate everybody ?

9. Because Arthur aggravated Trevor, Arthur was banned from the sports club.

Because Arthur aggravated Trevor, Trevor refused to go to the youth club.

Because Arthur aggravated Trevor, he was banned from the sports club.

Because Arthur aggravated Trevor, he refused to go to the youth club.

Who was banned from the sports club ?

Who refused to go to the youth club ?

10. Because Linda annoyed Polly, Linda was punished by the art teacher.

Because Linda annoyed Polly, Polly complained to the maths teacher.

Because Linda annoyed Polly, she was punished by the art teacher.

Because Linda annoyed Polly, she complained to the maths teacher.

Who was punished by the art teacher ?

Who complained to the maths teacher ?

11. Because Jimmy scolded Brian, Jimmy was accused of being too severe.

Because Jimmy scolded Brian, Brian went to sulk in the living room.

Because Jimmy scolded Brian, he was accused of being too severe.

Because Jimmy scolded Brian, he went to sulk in the living room.

Who went to sulk in the living room ?

Who was accused of being too severe ?

12. Because Diana confused Nancy, Diana tried to explain the solution more clearly.

Because Diana confused Nancy, Nancy had to ask the chemistry lecturer for help.

Because Diana confused Nancy, she tried to explain the solution more clearly.

Because Diana confused Nancy, she had to ask the chemistry lecturer for help.

Who had to ask the lecturer for help ?

Who tried to explain the solution more clearly ?

13. Because Tom punished Bob, Tom received criticism for being too strict.

Because Tom punished Bob, Bob began to feel angry about the situation.

Because Tom punished Bob, he received criticism for being too strict.

Because Tom punished Bob, he began to feel angry about the situation.

Who received criticism ?

Who began to feel angry about the situation ?

14. Because Liz flattered Sue, Liz earned a supportive reputation.
Because Liz flattered Sue, Sue felt more confident about life.
Because Liz flattered Sue, she earned a supportive reputation.
Because Liz flattered Sue, she felt more confident about life.

Who earned a supportive reputation ?

Who felt more confident about life ?

15. Because Rob thanked Jim, Rob was praised for being very polite.
Because Rob thanked Jim, Jim finally began to feel appreciated.
Because Rob thanked Jim, he was praised for being very polite.
Because Rob thanked Jim, he finally began to feel appreciated.

Who began to feel appreciated ?

Who was praised for being very polite ?

16. Because Beth congratulated Lisa, Beth was considered easily impressed.
Because Beth congratulated Lisa, Lisa started to feel motivated again.
Because Beth congratulated Lisa, she was considered easily impressed.
Because Beth congratulated Lisa, she started to feel motivated again.

Who started to feel motivated again ?

Who was considered easily impressed ?

Appendix 7

Materials used in Experiment 11

Materials 1-8 contain NP1 biasing verbs.

Materials 9-16 contain NP2 biasing verbs.

1. Because Barry dreaded Derek, he steadfastly refused to go back to school.

Because Barry dreaded Derek, he was told to try acting less aggressively.

Because Barry dreaded Lydia, he steadfastly refused to go back to school.

Because Lydia dreaded Barry, he was told to try acting less aggressively.

Who refused to go back to school ?

Who was told to act less aggressively ?

2. Because Philip appreciated Callum, he thanked for all the hard effort.

Because Philip appreciated Callum, he felt the work had been worth it.

Because Philip appreciated Tracey, he thanked for all the hard effort.

Because Tracey appreciated Philip, he felt the work had been worth it.

Who thanked for the effort ?

Who felt the work had been worth it ?

3. Because Ann adored Liz, she baked a special birthday cake.

Because Ann adored Liz, she became very popular at school.

Because Ann adored Ben, she baked a special birthday cake.

Because Ben adored Ann, she became very popular at school.

Who became popular at school ?

Who baked a special birthday cake ?

4. Because Rose admired Joan, she was impressed by what had been said.

Because Rose admired Joan, she finally started to feel appreciated.

Because Rose admired John, she was impressed by what had been said.

Because John admired Rose, she finally started to feel appreciated.

Who started to feel appreciated ?

Who was impressed by what had been said ?

5. Because Henry loathed Terry, he was punished for being spiteful.
Because Henry loathed Terry, he felt rather abandoned at school.
Because Henry loathed Norma, he was punished for being spiteful.
Because Norma loathed Henry, he felt rather abandoned at school.

Who was punished ?

Who felt abandoned at school ?

6. Because Jake resented Luke, he was thought to be of a jealous nature.

Because Jake resented Luke, he was starting to feel quite persecuted.

Because Jake resented Lisa, he was thought to be of a jealous nature.

Because Lisa resented Jake, he was starting to feel quite persecuted.

Who was thought to be of a jealous nature ?

Who was starting to feel persecuted ?

7. Because Carol despised Linda, she was considered to be unfriendly.

Because Carol despised Linda, she began to feel utterly miserable.

Because Linda despised James, she was considered to be unfriendly.

Because James despised Linda, she began to feel utterly miserable.

Who began to feel utterly miserable ?

Who was considered to be unfriendly ?

8. Because Nikki liked Susan, she obviously did not hate everybody.

Because Nikki liked Susan, she felt popular at the dinner party.

Because Susan liked Terry, she obviously did not hate everybody.

Because Terry liked Susan, she felt popular at the dinner party.

Who felt popular at the dinner party ?

Who did not hate everybody ?

9. Because Jack aggravated Tony, he was banned from the sports club.

Because Jack aggravated Tony, he refused to go to the youth club.

Because Jack aggravated Rita, he was banned from the sports club.

Because Rita aggravated Jack, he refused to go to the youth club.

Who was banned from the sports club ?

Who refused to go to the youth club ?

10. Because Michael annoyed Richard, he was punished by the art teacher.

Because Michael annoyed Richard, he complained to the maths teacher.

Because Michael annoyed Kathryn, he was punished by the art teacher.

Because Kathryn annoyed Michael, he complained to the maths teacher.

Who was punished by the art teacher ?

Who complained to the maths teacher ?

11. Because Jean scolded Rita, she was accused of being too severe.

Because Jean scolded Rita, she went to sulk in the living room.

Because Jean scolded Brian, she was accused of being too severe.

Because Brian scolded Jean, she went to sulk in the living room.

Who went to sulk in the living room ?

Who was accused of being too severe ?

12. Because Irene confused Rosie, she tried to explain the solution more clearly.

Because Irene confused Rosie, she had to ask the chemistry lecturer for help.

Because Irene confused Kenny, she tried to explain the solution more clearly.

Because Kenny confused Irene, she had to ask the chemistry lecturer for help.

Who had to ask the lecturer for help ?

Who tried to explain the solution more clearly ?

13. Because Ted punished Bob, he received criticism for being too strict.

Because Bob punished Ted, he began to feel angry about the situation.

Because Bob punished Liz, he received criticism for being too strict.

Because Liz punished Bob, he began to feel angry about the situation.

Who received criticism ?

Who began to feel angry about the situation ?

14. Because Ray flattered Tim, he earned a supportive reputation.

Because Ray flattered Tim, he felt more confident about life.

Because Ray flattered Sue, he earned a supportive reputation.

Because Sue flattered Ray, he felt more confident about life.

Who earned a supportive reputation ?

Who felt more confident about life ?

15. Because Anna thanked Emma, she was praised for being very polite.

Because Anna thanked Emma, she finally began to feel appreciated.

Because Emma thanked Bill, she was praised for being very polite.

Because Bill thanked Emma, she finally began to feel appreciated.

Who began to feel appreciated ?

Who was praised for being very polite ?

16. Because Ellen congratulated Gemma, she was considered easily impressed.

Because Gemma congratulated Ellen, she started to feel motivated again.

Because Ellen congratulated Grant, she was considered easily impressed.

Because Grant congratulated Ellen, she started to feel motivated again.

Who started to feel motivated again ?

Who was considered easily impressed ?

Appendix 8

Filler items used in all spr studies

The following sentences were used as filler materials for all of the self-paced reading studies reported in this thesis. The second question following each material was used in Experiment 4 while the first was used in all other spr experiments.

Ben punched Douglas in the face as he tried to dodge the blow.
Who was punched ?
No question.

Jennifer bought Lydia the record but she already had it.
Who bought the record ?
No question.

Because of the dreadful weather Tom thought Raymond had become stranded.
Who had become stranded ?
Was the weather good ?

Because Jon was smiling, Susan knew he had passed his driving test.
Who was smiling ?
No question.

Because the concert had sold out, Ted told Doug he couldn't get in.
Who couldn't get in ?
Was the concert sold out ?

Mary distrusted Ellen since she told people about her family's secret.
Who was mistrusted ?
Was there a family secret ?

Walter gave Ronald a dirty look since he thought he had betrayed him.
Who was given a dirty look ?
No question.

Cindy threw the ball to Janet since she was in a better position to score.
Who threw the ball ?
No question.

Because Greg owed Neil so much money, he decided to take him to court.

Who owed money ?

No question.

Harold told Trevor to meet him at the new pub.

Who was told ?

Was it an old pub ?

Alice thought Jenny a fool as she had lent some money to the tramp.

Who was thought to be a fool ?

Was a tramp lent money ?

Patsy and Becky raced to the bottom of the hill and Becky twisted her ankle.

Who twisted her ankle ?

No question.

Joe beat Brendon at tennis and then they had a game of football.

Who was beaten ?

No question.

Elaine waved for Sue to stop but she never saw the signal.

Who waved ?

Was there a signal ?

Neil played the piano for Greg but kept forgetting the tune.

Who played the piano ?

No question.

Penny laughed at Beth's joke before telling one herself.

Who laughed at the joke ?

No question.

Lucy disagreed with Wendy but she couldn't explain why.

Who disagreed ?

No question.

Try as he might, Alex never managed to beat Steven at badminton.

Who couldn't be beaten ?

No question.

Iain worried about Dave since he was always blacking out.

Who was always blacking out ?

No question.

Despite the bad weather, Liz asked Deb to play croquet.

Who was asked ?

Was the weather bad ?

While he still had some money left, Larry repaid his debt to Bob.

Who repaid the debt ?

No question.

Julie always supported Sharon as she always did likewise.

Who was supported ?

No question.

George joked to Henry about the time he got stuck in the lift.

Who got stuck in a lift ?

Did someone get stuck in a lift ?

Although she was her best friend, Rita was often nasty about Debbie.

Who was nasty ?

No question.

Bobby kicked the ball to Lloyd as he was near the goal.

Who was near the goal ?

No question.

Judy liked to go to the theatre with Jill as she liked similar plays.

Who liked to go to the theatre ?

No question.

Mark discussed the book with Barry as he couldn't understand it.

Who couldn't understand it ?

No question.

Evette thanked Debbie although she didn't really mean it.

Who was thanked ?

No question.

Jerry tricked Keith as he enjoyed playing practical jokes.

Who was tricked ?

Did someone play a practical joke ?

Although the sun was setting, Penny took a photograph of Susannah.

Who took the photograph ?

Was the sun rising ?

Since she was a good cook, Cathy cooked dinner for Liz.

Who was a good cook ?
No question.

Simon bought Rick the book as a birthday present.
Who bought the book ?
Was the book a Christmas present ?

Kenny chased Benny around the playground but couldn't catch him.
Who was chased ?
No question.

Joanne did Pamela's homework as she was a genius at maths.
Who was a genius at maths ?
No question.

Lenny tried to convince Joe but he refused to listen to him.
Who refused to listen ?
No question.

Rachel beat Helen in the race but was disqualified.
Who was disqualified ?
No question.

Eddie signalled to Earl as the train arrived at the station.
Who was signaled to ?
Was it a bus that arrived at the station ?

Tina was wearing a large hat and was spotted by Yvonne.
Who was wearing a large hat ?
Was it a small hat ?

Jonathan corrected his mistake but Peter was still unimpressed.
Who was unimpressed ?
No question.

Laurence sometimes annoyed Douglas although it was unintentional.
Who was annoyed ?
No question.

As she always said the wrong thing, Dorothy apologised for Amanda.
Who apologised ?
No question.

Since he noticed it was sunny, Rob asked Tim if he wanted to play cricket.
Who was asked to play cricket ?

Was it sunny ?

Kate invited Sue to her birthday although she didn't like her.

Who was invited ?

Was it a Halloween party ?

Eric gave the painting to Tom although it was his favourite.

Who was given the painting ?

No question.

Nancy arranged the bouquet for Ursula since she was so artistic.

Who was artistic ?

No question.

Because of her tolerance, Raquel was able to live with Lisa.

Who was tolerant ?

No question.

Although he hated sport, Blair played squash with Alec.

Who hated sport ?

Did they play table tennis ?

Joanna drove to visit Kirsten as they were best friends.

Who drove ?

No question.

Kirstie wasn't able to talk to Patricia because of her swollen gums.

Who had swollen gums ?

No question.

Ellen made Mary laugh as she performed the humorous impression.

Who performed the humorous impression ?

Was she in a bad mood ?

Was the impression unamusing ?

Although Gavin didn't really like Bill, he admired his determination.

Who was determined ?

No question.

Noel ridiculed Henry although he was supposed to be his friend.

Who was ridiculed ?

No question.

Michelle tried to talk to Helena but she was in a bad mood.

Who was in a bad mood ?

Was she in a bad mood ?

Although Simon crashed his car, he lied to Phillip about the accident.

Who crashed the car ?

Was there a car crash ?

Dominic always beat Trevor at darts but put it down to luck.

Who always won ?

No question.

Tom tuned the piano for Richard but damaged it in the process.

Who damaged the piano ?

Was the piano damaged ?

Steven easily beat Ian in the race although he was quite unfit.

Who was unfit ?

No question.

Susan babysat for Samantha although she had other commitments.

Who had other commitments ?

No question.

Because of the power cut, Edna couldn't bake Gaynor the cake.

Who couldn't bake the cake ?

Was there a power cut ?

Disturbed by his snoring, Gavin threw a pillow at Paul.

Who threw the pillow ?

No question.

Awoken by her loud music, Louise shouted at Fiona.

Who was shouted at ?

Was the music quiet ?

Nora videoed Betty as it was her fortieth birthday.

Whose birthday was it ?

Was it a thirtieth birthday ?

David drank the bleach and William had to rush him to hospital.

Who was rushed to hospital ?

Was whisky drunk ?

Elizabeth bought Anna a present and then treated her to a meal.

Who was bought a present ?

No question.

Ron questioned Michael about his increasingly extravagant lifestyle.

Who was questioned ?
Was his lifestyle simple ?

Nicholas shoved Brian into the street as he was becoming irritating.
Who was irritating ?
No question.

Since Sophie was so generous, she bought Andrea an expensive gift.
Who was bought a gift ?
Was the gift cheap ?

Mike was in awe of Timothy although he didn't always believe him.
Who was in awe ?
No question.

Ross ran after Henry as he saw him drop his wallet.
Who dropped his wallet ?
Was a wallet dropped ?

Accusing Frank, James tried to convince his boss to sack him.
Who was accused ?
No question.

Amy bought Jennifer a drink as she still owed her some money.
Who owed money ?
No question.

Calling for help, John was punched by Dougal as he lay on the ground.
Who was punched ?
No question.

Joey wondered if Kirk would remember to leave him his keys.
Who had the keys ?
No question.

Patrick tripped over Paul since he was rather uncoordinated.
Who was uncoordinated ?
No question.

Natasha wondered whether Lydia would meet her as she was unreliable.
Who was unreliable ?
Was someone unreliable ?

Maggie saw the film with Elaine although she had read bad reviews.
Who had read bad reviews ?

Were the reviews good ?

Because of the drought, Clare told Pippa she would have to conserve water.

Who was told to conserve water ?

No question.

Margaret hinted to Claire that she should watch less television.

Who hinted ?

No question.

Debbie advised Justine that she should choose her friends with care.

Who was advised ?

No question.

Carole thanked Rosie for the advice before she left the office.

Who left the office ?

No question.

After Mark broke his arm, Jeremy drove him to the hospital.

Who drove to hospital ?

Was an ambulance called ?

Try as she might, Barbara couldn't persuade Judith to lend her any money.

Who tried to borrow money ?

No question.

Sarah frightened Phoebe as the house was very spooky.

Who was frightened ?

Was the house spooky ?

Kate sat beside Gemma and tried to copy her homework.

Who tried to copy ?

No question.

Although they used to be friends, Mary lost contact with Jean.

Who lost contact ?

No question.

Troubled as he was by the question, Brian asked Clive for advice.

Who was asked for advice ?

Was the question troubling ?

Alan bought the car from Sid but had nothing but trouble with it.

Who was the car bought from ?

No question.

Peter recited the speech to Hugh but he was underawed by it.

Who recited the speech ?

No question.

Stephen spoke as Jon tried to concentrate on listening to him.

Who spoke ?

No question.

Belinda practised the dance in front of Suzy who was full of advice.

Who was full of advice ?

No question.

Since it was nearly Christmas, Anna asked Joan to book a holiday for her.

Who was asked ?

No question.

Diving into the pool, Carl landed on Justin who hurt his arm.

Who hurt his arm ?

No question.

Kathleen pressured Eliza as she needed an answer.

Who needed an answer ?

No question.

Matt tricked Thomas as he had organised a suprise party for him.

Who organised a suprise party ?

No question.

Sheena wrote to Sally although she hadn't heard from her in years.

Who wrote ?

No question.

Lindsey was met by Marisa by the clocktower.

Who was met ?

No question.

