

**Language and Causal Understanding:**

**There's Something About Mary**

**Asifa Majid**

Department of Psychology  
University of Glasgow

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

March 2001

© Asifa Majid

ProQuest Number: 13818983

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 13818983

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



12275-Copy 1

## Acknowledgements

Thanks to Ash Asudeh and Line Mikkelsen for the two years we spent in Edinburgh, drinking tea and talking about language – and everything else.

I am indebted to Chuck Clifton, Keith Rayner and Lyn Frazier for all the difficult questions they asked me in Amherst. Thanks to Kyle Johnson, Julie Johnson and Jane Ashby for the easy questions they asked.

I would like to thank those friends in Glasgow, who kept me sane these last years. Thanks especially to Frederic Gosselin, Nick Zangwill and Linda Moxey; and also to Roger van Gompel and Andrew Stewart. Thank you Al, Emily, Neil, Rory and Simon. Thanks to Kate Cavanagh for keeping me going over the last five months. Lastly, thank you Sandie Cleland for everything.

Finally, I would like to express my gratitude to my supervisors Tony Sanford and Martin Pickering. Thank you for all of the support and encouragement you have provided over the years. Thanks especially to you Tony for inspiring my love of psychology and teaching me the importance of asking big questions.

This research has been funded by the University of Glasgow.



## **Declaration**

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

# Abstract

Making causal inferences is a ubiquitous property of the cognitive system. This dissertation examines how people make causal attributions in the social domain. Two cues to causality are examined. The first is the implicit causality in verbs and the second is covariation information.

When people are presented with minimal sentences such as *Mary fascinated Ted* and then are asked what the cause of that event is, then people say something about Mary, for example, *Mary was interesting*. On the other hand, if they are presented with sentences such as *Mary liked Ted*, and then are asked about the cause, they say something about Ted, for example, *Ted is nice*. In the first example, cause is attributed to Mary or the first noun phrase (NP1); whereas in the second example, it is attributed to Ted, or the second noun phrase (NP2). This is called the implicit causality verb bias.

The implicit causality bias is reviewed in some detail in the first two chapters. This is followed by a test of whether it is present even when no causal question is asked. The results from Chapter 3 suggest that the verb bias is present in such circumstances.

Chapter 4 examined the relation between implicit causal information, such as that provided by implicit causality verbs, and explicit causal information, such as covariation theory. According to covariation theory, cause is determined by establishing what covaries with what. Chapter 4 demonstrated that both implicit and explicit sources of causal information are used to make attributions in production. However, Chapter 5 showed that while implicit causal information is also used in comprehension, the effect of explicit covariation information is weak.

In order to ascertain exactly which cues people make to use attributions from covariation information Chapters 6 and 7 contrast a frequency signalling account of covariation theory with a focussing account. According to the frequency signalling account, cause is attributed to that individual which is in the smallest group; while according to the focussing account, cause is attributed to the individual who is in focal attention. It is found that both frequency and focussing influence attributions – and in very systematic ways.

# Table of Contents

**Acknowledgements.....2**

**Declaration.....3**

**Abstract.....4**

**Table of Contents .....5**

**List of Tables.....11**

**List of Figures .....12**

**INTRODUCTION .....14**

*0.1. Causal cognition.....14*

*0.2. Causality and the lexicon.....16*

**CHAPTER 1: REVIEW OF THE IMPLICIT CAUSALITY BIAS .....20**

*1.1. The breadth of the implicit causality bias.....20*

*1.2. Classifying implicit causality verbs.....22*

        1.2.1. Thematic Role Taxonomy .....23

        1.2.2. Linguistic Category Model.....27

        1.2.3. Discussion .....30

*1.3. Effects that moderate the implicit causality bias.....33*

        1.3.1. Manipulating presentation format.....33

            1.3.1.1. Active versus passive voice .....33

            1.3.1.2. Interrogative versus declarative form .....36

            1.3.1.3. Affirmative versus negative phrasing .....36

        1.3.2. Manipulating semantic properties of the verb .....37

            1.3.2.1. Semantic valence .....37

        1.3.3. Manipulating contextual factors .....41

            1.3.3.1. Gender of characters and participants and the implicit causality bias .....41

        1.3.4. Is the implicit causality bias “interpersonal”?.....43

        1.3.5. General context effects .....46

1.3.4. Summary .....	47
1.4. Overview .....	48
<b>CHAPTER 2: EXPLANATIONS OF IMPLICIT CAUSALITY.....</b>	<b>50</b>
2.1. Introduction .....	50
2.2. <i>Implicit causality is implicit covariation</i> .....	50
2.2.1. Implicit covariation theory .....	50
2.2.1.1. Causal attributions in social psychology.....	50
2.2.1.2. Covariation theory .....	51
2.2.1.3. Implicit covariation.....	53
2.2.1.4. Tests of the implicit covariation theory .....	64
2.3. <i>Implicit causality is implicit salience</i> .....	68
2.3.1. Implicit salience theory .....	68
2.3.2. Tests of the implicit salience theory .....	69
2.4 <i>Implicit causality and the priming hypothesis</i> .....	74
2.4.1. The priming hypothesis .....	74
2.4.2. Tests of the priming hypothesis.....	76
2.5. Conclusion .....	79
<b>CHAPTER 3: IS IMPLICIT CAUSALITY A FOCUSING PROPERTY OF VERBS? .....</b>	<b>83</b>
3.0 Introduction .....	83
3.1 <i>Experiment 1:</i> .....	89
3.1.1. Method.....	90
3.1.1.1 Participants .....	90
3.1.1.2 Materials .....	90
3.1.1.3 Procedure.....	91
3.1.2 Results .....	91
3.1.2.1 Coding .....	91
3.1.2.2 Data analyses .....	93
3.1.3 Discussion .....	99
3.2 <i>Experiment 2:</i> .....	101
3.2.1 Method.....	101
3.2.1.1 Participants .....	101
3.2.1.2 Materials and Procedure .....	102
3.2.2 Results .....	102
3.2.2.1 Coding .....	102
3.2.2.2 Data analyses .....	102
3.2.3 Discussion .....	107
3.3 <i>EXPERIMENT 3:</i> .....	109

3.3.1	Method.....	109
3.3.1.1	Participants .....	109
3.3.1.2	Materials and Procedure .....	109
3.3.2	Results .....	110
3.3.2.1	Coding criteria .....	110
3.3.2.2.	Data analyses .....	110
3.4	<i>General discussion</i> .....	112
<b>CHAPTER 4: COVARIATION INFORMATION IN DISCOURSE.....</b>		<b>115</b>
4.0.	<i>Introduction</i> .....	115
4.1.	<i>Experiment 4:</i> .....	118
4.1.1.	Method.....	119
4.1.1.1.	Participants .....	119
4.1.1.2.	Materials .....	119
4.1.1.3.	Procedure .....	120
4.1.2.	Results .....	121
4.2.1.	Covariation and the implicit causality bias .....	121
4.2.2.	Is covariation strong enough to over-ride the implicit causality bias?.....	122
4.1.3.	Summary .....	126
4.2.	<i>Experiment 5:</i> .....	127
4.2.1.	Method.....	128
4.2.1.1.	Participants .....	128
4.2.1.2.	Materials .....	128
4.2.1.3.	Procedure .....	129
4.2.2.	Results .....	129
4.3.1.	Covariation and the implicit causality bias .....	130
4.3.2.	Summary.....	131
4.3.3.	Does covariation over-ride the implicit causality bias? .....	132
4.3.4.	Summary.....	134
4.3.	<i>General Discussion</i> .....	134
<b>CHAPTER 5: COVARIATION THEORY IN COMPREHENSION .....</b>		<b>137</b>
5.0.	<i>Introduction</i> .....	137
5.1.	<i>Experiment 6:</i> .....	142
5.1.1.	Method.....	142
5.1.1.1.	Participants .....	142
5.1.1.2.	Materials .....	142
5.1.1.3.	Procedure .....	145
5.1.2.	Results .....	146
5.1.3.	Discussion .....	148

5.2. Experiment 7:	148
5.2.1. Method:	149
5.2.1.1. Participants	149
5.2.1.2. Materials and Procedure	149
5.2.2. Results	149
5.2.3. Discussion	150
5.3. Experiment 8:	150
5.3.1. Method:	151
5.3.1.1. Participants	151
5.3.1.2. Materials and Procedure	151
5.3.2. Results	152
Referent:	153
5.3.3. Discussion	153
5.4. Experiment 9:	154
5.4.1. Method:	154
5.4.1.1. Participants	154
5.4.1.2. Materials	154
5.4.1.3. Procedure	156
5.4.2. Results	156
5.4.2.1. Reading time for target sentence:	156
5.4.2.2. Question response accuracy	158
5.4.2.3. Question response time	159
5.4.3. Discussion	161
5.5. General discussion:	162
<b>CHAPTER 6: COVARIATION THEORY AND NATURAL LANGUAGE QUANTIFIERS</b>	<b>164</b>
6.0. Introduction	164
6.1. Experiment 10:	170
6.1.1. Method:	171
6.1.1.1. Participants	171
6.1.1.2. Materials	171
6.1.1.3. Procedure	172
6.1.2. Results	172
6.1.2.1. Baseline comparisons	173
6.1.2.2. Frequency signalling or focussing?	176
6.1.3. Discussion	178
6.2. Experiment 11:	179
6.2.1. Method:	181
6.2.1.1. Participants	181
6.2.1.2. Materials and Procedure	181

6.2.2. Results .....	182
6.2.2.1. Baseline comparisons .....	182
6.2.2.2. Is there a focus difference? .....	185
6.2.2.3. Congruency or Saliency? .....	188
6.2.3. Discussion .....	190
<b>CHAPTER 7: COVARIATION THEORY AND NUMERICAL AND PERCENTAGE</b>	
<b>INFORMATION .....</b>	<b>192</b>
 7.0. Introduction .....	 192
7.1. Experiment 12:.....	193
7.1.1. Method.....	193
7.1.1.1. Participants .....	193
7.1.1.2. Materials .....	193
7.1.1.3. Procedure .....	194
7.1.2. Results .....	194
7.1.3. Summary .....	195
7.2. Experiment 13:.....	195
7.2.1. Method.....	196
7.2.1.1. Participants .....	196
7.2.1.2. Materials .....	196
7.2.1.3. Procedure .....	197
7.2.2. Results .....	197
7.2.2.1. Is the focus difference due to a numerical difference? .....	198
7.2.2.2. Is there an overall effect of numerical information? .....	198
7.2.2.3. Baseline comparisons .....	199
7.2.3. Discussion .....	202
7.3. Experiment 14:.....	203
7.3.1. Method.....	203
7.3.1.1. Participants .....	203
7.3.1.2. Materials and Procedure .....	203
7.3.2. Results .....	203
7.3.3. Summary .....	204
7.4. Experiment 15:.....	204
7.4.1. Method.....	205
7.4.1.1. Participants .....	205
7.4.1.2. Materials .....	205
7.4.1.3. Procedure .....	206
7.4.2. Results .....	206
7.4.2.1. Is there an overall effect of proportional information? .....	206
7.4.2.2. Baseline comparisons .....	207

7.4.3. Discussion .....	210
7.5. <i>Comparing base-rate information</i> .....	211
7.5.1. Natural language quantifiers and numerical information .....	211
7.5.2. Natural language quantifiers and percentage information .....	212
7.5.3. Numerical information and percentage information .....	213
7.6. <i>General Discussion</i> .....	214
<b>General Discussion .....</b>	<b>216</b>
<b>References .....</b>	<b>220</b>
<b>Appendix 1 .....</b>	<b>230</b>
<b>Appendix 2 .....</b>	<b>234</b>
<b>Appendix 3 .....</b>	<b>236</b>
<b>Appendix 4 .....</b>	<b>237</b>
<b>Appendix 5 .....</b>	<b>238</b>
<b>Appendix 6 .....</b>	<b>240</b>



# List of Tables

**Table 1.....53**

**Table 2.....60**

**Table 3.....63**

**Table 4.....72**

**Table 5.....73**

**Table 6.....91**

**Table 7.....94**

**Table 8.....103**

**Table 9.....129**

**Table 10.....141**

**Table 11.....145**

**Table 12.....146**

**Table 13.....150**

**Table 14.....153**

**Table 15.....159**

**Table 16.....181**

**Table 17.....187**

**Table 18.....195**

**Table 19.....197**

**Table 20.....204**

**Table 21.....205**

# List of Figures

Figure 1.....54

Figure 2.....55

Figure 3.....58

Figure 4.....59

Figure 5.....61

Figure 6.....61

Figure 7.....64

Figure 8.....69

Figure 9.....75

Figure 10.....86

Figure 11.....95

Figure 12.....96

Figure 13.....97

Figure 14.....98

Figure 15.....101

Figure 16.....103

Figure 17.....104

Figure 18.....105

Figure 19.....106

Figure 20.....108

Figure 21.....111

Figure 22.....122

Figure 23.....124

Figure 24.....125

Figure 25.....126

Figure 26.....130

Figure 27.....131

Figure 28.....133

Figure 29.....134

Figure 30.....152

Figure 31.....157

<b>Figure 32.....</b>	<b>158</b>
<b>Figure 33.....</b>	<b>160</b>
<b>Figure 34.....</b>	<b>161</b>
<b>Figure 35.....</b>	<b>174</b>
<b>Figure 36.....</b>	<b>175</b>
<b>Figure 37.....</b>	<b>176</b>
<b>Figure 38.....</b>	<b>183</b>
<b>Figure 39.....</b>	<b>184</b>
<b>Figure 40.....</b>	<b>185</b>
<b>Figure 41.....</b>	<b>189</b>
<b>Figure 42.....</b>	<b>190</b>
<b>Figure 43.....</b>	<b>199</b>
<b>Figure 44.....</b>	<b>200</b>
<b>Figure 45.....</b>	<b>200</b>
<b>Figure 46.....</b>	<b>201</b>
<b>Figure 47.....</b>	<b>207</b>
<b>Figure 48.....</b>	<b>208</b>
<b>Figure 49.....</b>	<b>209</b>
<b>Figure 50.....</b>	<b>210</b>
<b>Figure 51.....</b>	<b>212</b>
<b>Figure 52.....</b>	<b>213</b>
<b>Figure 53.....</b>	<b>214</b>

# INTRODUCTION

## 0.1. Causal cognition

People have to be able to make judgements about cause in order to be effective in the world. As humans, we have to be able to predict the environment so that we can decide on the appropriate action. From an evolutionary perspective, we have to be able to avoid dangers, gather food, and find a receptive mate. In order to be able to do any of this we have to be able to represent causal regularities: that the big animal with the spotted fur caused the gash in my leg; that the bush on the other side of the forest provides abundant berries; that the female without a current partner is more likely to be susceptible to my advances.

In modern day psychology, many questions have been asked about the causal relations we make, and the cues that form the basis of our causal judgements. Just as Gestalt psychologists attempted to identify the spatial relations that made an object cohere, researchers today have attempted to identify the causal relations that make events cohere. Generally, psychologists have asked whether knowledge of causality is innate or learned, and whether it is domain-general or domain specific (see Corrigan & Denton, 1996; Sperber, Premack & Premack, 1995 for a review).

Given the above evolutionary reasons for the need to be sensitive to the environment, one question that has been asked is whether the causal processes underlying animal and human actions are the same (*e.g.* Dickinson & Shanks, 1995). From this research it has been established that at least one type of causal knowledge, associative learning based on contiguous events, as in classical conditioning (Pavlov, 1927) and instrumental conditioning (Skinner, 1953) does involve the same process all the way down the phylogenetic tree. Here the organism learns that the CS causes the UCS (or a mental representation of the UCS); or that an action causes a particular outcome.

For human beings, causal cognition is implicated in many different types of behaviour – not just whether two events are contiguous. Within the field of perception, researchers have attempted to establish whether causal information is used in order to recover the physical structure of the world. For example, in developmental psychology

there is evidence that infants as young as 6 months know that when an object is released in mid-air it should drop to the ground and not float in mid-air (Spelke, Jacobson, Keller & Sebba, submitted, reported in Baillargeon, Kotovsky & Needham, 1995). This demonstrates that causal knowledge is being used by them in order to make sense of the physical world.

Another strand of research in perception, which has received a lot of attention, is whether causality is perceived, in the same way that, for example, depth is perceived. That is, is causal perception fast, automatic, irresistible and stimulus driven (Scholl & Tremoulet, 2000). This line of research follows from the work by Michotte (1946/1963) and Heider & Simmel (1944). In a typical experiment, 2-D geometric shapes are shown moving on a screen. From these moving shapes people perceive causal relations. For example, two squares are shown on a screen. Box A moves in a straight line until it reaches Box B, at which point A stops moving while B moves along the same trajectory. People perceive this sequence of events as A causing B's motion. Further, Heider & Simmel found that similar stimuli also led to people attributing intentions to the geometric shapes, so for example, people would see B as *wanting to move away from A*, or *B escaping from A*.

This characteristic of making attributions about intentions, and what causes behaviour, is ubiquitous in human behaviour. Again, developmental psychology has shown that infants are sensitive to the difference between animate and inanimate things, and appear to realise that animates have intentions. Further, infants of only 12 months are surprised when a person acts in a way that is inconsistent with the emotion that they are displaying (Spelke, Philips & Woodward, 1995).

In social cognition, researchers have attempted to identify the cues people use in order to make causal attributions, such as covariation information (Kelley, 1967), abnormal conditions (Hilton & Slugoski, 1986), causal mechanisms (Ahn, Kalish, Medin & Gelman, 1995), *etc.*. Or, they have attempted to find what people attribute behaviour to: is it attributed to something internal to the person; or something external in the environment (Heider, 1958). This attempt to locate the source of an effect has led to the discovery biases in attribution process, such as the fundamental attribution error. The fundamental attributions error is the tendency towards dispositional (internal) rather than situational (external) causal attributions (Heider, 1958; Ross, 1977).

Social psychologists have also shown that the type of causal explanation people use depends on the social category of the individual (*e.g.* Grier & McGill, 2000). At the same time that category membership affects causal explanation, researchers have shown

that category construction and membership itself relies on a causal understanding of the world. For example, Murphy & Medin (1985) argue that categorisation cannot be done by feature matching alone, but instead it requires domain-specific theories. These domain-specific theories can only be acquired with knowledge of causation.

So, it can be seen that causal cognition has been a widely researched topic in many areas, such as learning, perception, development, social cognition, categorisation, to name but a few. People search for causal relations in order to make sense of the world. When events in a narrative have strong causal dependencies with one another, this eases comprehension (Trabasso & Sperry, 1985; Trabasso & van den Broek, 1985), and improves subsequent recall (Myers, Shinjo & Duffy, 1987). When causal information is provided with problems then this leads to enhanced problem solving and decision making (Tversky & Kahneman, 1986). In fact, causal information is fundamental to all aspects of cognition.

## **0.2. Causality and the lexicon**

Because people are so sensitive to causal information, it has been argued that cause is a conceptual primitive – it is used in order to understand concepts that are more complex (Jackendoff, 1983; 1990; Wierzbecka, 1992; 1996). A traditional way of representing the meaning of words or concepts is to decompose them into smaller units. So, for example, the meaning of the *bachelor* can be expressed as a concatenation of smaller conceptual units such as ADULT, MALE, UNMARRIED (Katz & Fodor, 1963; Miller & Johnson-Laird, 1976). However, there are accepted problems in expressing meaning in this manner (*e.g.* Murphy & Medin, 1985). Despite this, there is a dominant perspective in meaning today, which holds that there are some basic primitive conceptual units, which are used as the building blocks for more complex concepts.

One example of complex concepts being analysed into smaller conceptual units is in the domain of emotions (*e.g.* Johnson-Laird & Oatley, 1989; Wierzbecka, 1992). In order to understand a complex emotion such as regret or embarrassment, one can express that concept in terms of basic emotions. There are possibly five (although the actual number is in dispute ranging anywhere between 2 and 11, Power & Dalgleish, 1997) basic emotions. These can be denoted by the words happiness, sadness, fear, anger, disgust, (and possibly surprise). For example, feeling regret means something like: “feeling sad as a result of evaluating a past action as harmful or wrong in relation

to one's current standards" (Johnson-Laird & Oatley, 1989, 103). Regret entails the basic emotion of sadness.

It is assumed that basic emotions have a biological significance. The perception of these emotions in others helps us to be able to interact with the world in a more efficient way. Famously, Ekman (1973) showed that people associated a limited number of emotions to different facial expressions, and that the emotion terms associated with the faces are the same regardless of the cultural background of the person. Ekman argued that identification of these emotions was important because it helped modification of behaviour in an appropriate way. So, for example, if someone appears angry then the appropriate response would be to flee from that person.

Just as there have been attempts to build up complex emotions from basic ones, there have been numerous attempts to build up complex events from basic ones. Jackendoff (1983; 1990; Norman, Rumelhart, & L. N. R. Research Group, 1975) argued that complex events can be built from basic concepts such as GO, STAY, CAUSE, *etc.* Once again, these conceptual primitives are ones that we need in order to be able to negotiate the world. We have to be able to understand CAUSE, otherwise we will be eaten by that spotted furry animal, will not go to the same location to collect food, and will not realise that our advances will be more successful in certain situations.

Conceptual primitives are those that have biological significance. So, apart from basic concepts such as CAUSE and GO, there are others such as DIE (Goddard & Wierzbicka, 1998, cited in The Annual Report of the Max Planck Institute for Psycholinguistics, 1999). DIE is a universal semantic primitive because it is something that all human beings experience, and so it has significance. By using these conceptual primitives, events that are more complex can be expressed. For example, we can have events such as kill, which entail the basic events CAUSE and DIE.

Hence, it has been shown that causality plays a large part in human cognition, so much so that it has become lexicalised. More recently, researchers have found that people make causal attributions for verbs that do not entail causality.

There is a set of verbs that denote an interaction between people. When people are presented with a minimal sentence, which includes one of these verbs, *e.g. Mary fascinated Ted*, then they make consistent causal attributions to one of the characters in the sentence. In the case of *Mary fascinated Ted*, people think that there's something about Mary, which caused the event. So, when asked why this is the case, people will typically say something like: *Mary was interesting, Mary can play golf, etc.* On the other hand, another set of verbs exists where a consistent causal attribution is made to

Ted instead of Mary. For example, when people are asked to explain why *Mary liked Ted*, people will say something about Ted, *e.g. Ted was nice, Ted gave Mary's brother a present*.

The interesting thing about these verbs is that they do not *entail* a causal relation, yet people make consistent and reliable attributions to one or other character. They think that there is something about Mary, or Ted, which caused the event. This has been dubbed the implicit causality verb bias. This thesis will examine the causal attributions people make, and investigate the cues people use to make such judgements.

In the first chapter, an overview of the experimental literature on the implicit causality bias will be provided. The chapter will begin by examining two different classification schemes that have been proposed in order to distinguish verbs by the bias that they exhibit. This will be followed by a review of factors, other than the verb, that have been shown to influence the implicit causality bias.

Chapter 2 reviews three explanations of the implicit causality bias, and the problems associated with each of them. The first theory is implicit covariation theory. According to the implicit covariation theory general cues, which social psychologists have shown to be used in making attributions, have become lexicalised. The second account is implicit salience theory, which is also based on findings from social cognition. According to this theory, most causal weight is given to the most salient argument. And finally the priming hypothesis claims that the implicit causality bias is the result of spreading activation from the verb to other related words in an associative memory.

This leads into the first experimental chapter where the boundaries of the implicit causality bias are examined. There have been numerous claims, since the discovery of the implicit causality verb bias is not solely a property of the verb, but that it is only apparent when the verb is presented with a causal question. Chapter 3 examines this claim. Participants are presented with minimal sentences, and are not asked any type of causal question. Even under these circumstances, we find that there is evidence of an implicit causality bias.

The remainder of the thesis seeks to examine the interaction between implicit causal information provided by verbs and explicit causal cues provided through context. Chapters 4 and 5 present covariation information, which has been shown to influence causal attributions, as well as implicit causality verbs. In production (Chapter 4), both sources of information independently influence attributions, with a very intriguing effect of context on the implicit causality bias. However, in comprehension (Chapter 5)



there is a different pattern of results. Here the explicit covariation information is a much weaker cue.

From these three chapters, it appears that the implicit causality bias is a very strong cue when making causal attributions: it is apparent when no causal question is asked, and even weighed more strongly when other cues to causality are provided. However, it is not clear what information participants use from explicit covariation information in order to make causal attributions. In order to answer this question Chapters 6 and 7 describe experiments which manipulate the way in which covariation information is presented. Chapter 6 presents different types of natural language quantifiers, while Chapter 7 presents numerical or proportional information. From these studies, it is concluded that the rhetorical value of the context statement is more important than the frequency information that it denotes.

First, let us examine the implicit causality bias in a bit more detail.

# CHAPTER 1:

## REVIEW OF THE IMPLICIT CAUSALITY BIAS

### 1.1. The breadth of the implicit causality bias

It has been over 25 years since the implicit causality verb bias was first reported by Garvey & Caramazza (1974). Since then there has been a constant interest in the implicit causality bias. When participants were asked to explain why an interaction between two people occurred they displayed a consistent and reliable “bias” to refer to one individual. Further, the preferred referent changed when the verb was changed. For example, when asked why *Mary fascinates Ted* participants attributed the cause of fascinating to Mary. Verbs, such as *fascinate*, that bias towards the first mentioned individual are called NP1 biasing verbs. However, when asked why *Mary liked Ted* participants attributed the cause to Ted. Verbs, such as *like*, that bias towards the second mentioned individual, are called NP2 biasing verbs.

The implicit causality verb bias is evident across a wide range of tasks. It is observed using forced choice procedures, such as rating scales, where the participant has to rate each character for causality (Brown & Fish, 1983a; Corrigan, 1988; 1992; Hoffman & Tchir, 1990; Kasof & Lee, 1993; Lafrance, Brownell & Hahn, 1997; McArthur, 1972; Semin & Marsman, 1994). Another forced choice procedure is to use a disambiguation task, where the participant is given a sentence with an ambiguous pronoun, and then has to choose the referent of the pronoun. For example, the participant is presented with *Ted telephoned Dominic because he wanted sympathy*, and then has to indicate whether *he* refers to *Ted* or *Dominic* (Caramazza, Grober, Garvey & Yates, 1977; Garvey, Caramazza & Yates, 1976; Mannetti & DeGrada, 1991).

An implicit causality bias is also manifest in less constrained situations, such as completion tasks. Typically, the participant is given a sentence fragment with a causal connective and the participant has to provide a sensible ending to that fragment. For example, the fragment *Ted telephoned Mary because...* would be provided and the participant would be free to complete the sentence with whatever they wanted. Even in this relatively open-ended task, participants display a bias towards to either NP1 or NP2

(Au, 1986; Brown & Fish, 1983a; 1983b; Brown & VanKleeck, 1989; Franco & Arcuri, 1990; Garvey, Caramazza & Yates, 1976).

Further, the implicit causality bias has been shown to influence ease of comprehension. Sentences with an explicit cause that is congruent with the verb's implicit causality bias are read faster than sentences with a cause incongruent with the bias. For example, *apologised to* is a NP1 biasing verb. When *apologised to* is presented in a sentence with an explicit cause referring to NP1, e.g. *Ted apologised to Mary because he had been behaving selfishly*, then it is read faster than when the same verb is presented with an ending referring to NP2, e.g. *Ted apologised to Mary because she didn't deserve the criticism* (Garnham, Oakhill & Cruttenden, 1992; Pynte, Kennedy, Murray & Courrieu, 1998; Stewart, Pickering & Sanford, 2000;).

Finally, when natural text is examined the implicit causality verb bias is evident there too. Long & De Ley (2000; also see Burgess, Livesay & Lund, 1998) conducted a corpus study, where they found that NP1 biasing verbs which appeared with a causal connective were predominantly followed by a reference to NP1; while NP2 biasing verbs were followed by NP2 referents. Therefore, across a number of different tasks, and even naturalistic observation, verbs display a consistent and reliable bias.

Not only is the implicit causality bias evident across a range of tasks, it is also evident across a variety of languages. Apart from English, the implicit causality bias has been reported in other Indo-European languages such as Afrikaans (Vorster, 1985), Dutch (Semin & Marsman, 1994), French (Pynte, Kennedy, Murray & Courrieu, 1988), German (Fiedler & Semin, 1988; Rudolph, 1997), Greek (Natsopoulos, Grouios, Bostantzopoulou, Mentenopoulos, Katsarou & Logothetis, 1993) and Italian (Franco & Arcuri, 1990; Mannetti & De Grada, 1991). Further, even in non-Indo-European languages such as Chinese (Brown & Fish, 1983b) and Japanese (Brown, 1986) people show a strong preference for attributing the cause of an interpersonal verb to a particular character.

The interesting thing about the implicit causality verb bias is that the verb itself does not entail causal information (although it may *imply* it). In the examples given so far: *fascinate*, *adore*, *telephone*, *apologise to*, there is nothing in the verb itself that entails a cause. Unlike *kill*, CAUSE is not a part of the semantics of the word, because the meaning to the word can be understood without ever knowing about the underlying causal relation. Despite this, people across different languages and different tasks show knowledge of exactly such relations. It is this knowledge that has continued to interest researchers for over 25 years. Why do participants attribute causality to Mary in the sentence *Ted loves Mary*? Is the preference for characters in a certain sentential position

(NP1 versus NP2), or is it something else? Is it only the verb that affects attributions? Which other factors influence the implicit causality bias? What is the implicit causality bias? How can we explain it? The remainder of this chapter, and the following chapter address these questions.

## 1.2. Classifying implicit causality verbs

Prior to Garvey & Caramazza's (1974) paper, researchers noticed that groups of verbs showed systematic differences. Abelson & Kanouse (1966; Kanouse, 1972) demonstrated a very intriguing finding that simple sentences of the form *subject verb object* provided implicit information about the proportion of objects which are "verbed". For example, Kanouse (1972) presented participants with sentences such as *Ted bends paper clips*, and then asked what the smallest proportion of paper clips being bent would be for the unquantified sentence to be true. Using this procedure, he found differences between verb groups. Subjective verbs, which "express sentiments or subjective orientations (*e.g. like, love, understand*)" produced larger group generalisations than manifest verbs, which "express objective actions or manifest relationships (*e.g. have, buy produce*)" (Kanouse, 1972, 141). By dividing verbs along this subjective-manifest dimension, certain properties of a particular verb could be predicted.

Other researchers used elaborated versions of this distinction in order to classify verbs. Based on established linguistic categories McArthur (1972) used four categories of verbs in a seminal paper on causal attributions. Verbs were either emotions (*e.g. afraid*), opinions (*e.g. think*), accomplishments (*e.g. translate*) or actions (*e.g. contribute*). Again, a distinction was made between "subjective" verbs, such as emotions and opinions and "manifest" verbs, such as accomplishments and actions. This division is still apparent today.

When classifying implicit causality verbs, a distinction is made between action (manifest) and state (subjective) verbs. Action verbs are those that refer to behavioural interactions typically involving voluntary muscles; whereas states refer to mental interactions that typically give rise to involuntary states (Rudolph & Forsterling, 1997).

Apart from this overt definition of actions and states, a variety of semantic and grammatical tests have been used in order to identify action and state verbs. For example, Brown & Fish (1983a; Semin & Fiedler, 1988, 1992; Semin & Marsman, 1994) used the following tests in order to distinguish action and state verbs. First, it is

claimed that action verbs can be used in the imperative form *e.g. Hit Mary!*, but state verbs cannot, as shown by the strangeness of *Adore Mary!*. Second, action verbs can be used in the progressive tense *e.g. Ted is hitting Mary*, but state verbs cannot. The sentence *Ted is adoring Mary* seems very strange. A third diagnostic is that states are substitutable in the frames “BEING X” and “FEELING X”. For example, *happy* is a state as both ‘*being happy*’ and ‘*feeling happy*’ are accepted as referring to emotions. However, *ignored* is not a state because although ‘*feeling ignored*’ is acceptable as an emotion, ‘*being ignored*’ is not (Johnson-Laird & Oatley, 1989).

Early studies found that action verbs elicited NP1 attributions, whereas state verbs elicited NP2 attributions (*e.g. McArthur, 1972*). But, later work demonstrated a more complex picture of attributions. Brown & Fish (1983a) found that some state verbs such as *fascinate* and *bore* were actually NP1 biasing, and Au (1986) showed that some action verbs, *e.g. punish* and *congratulate* were NP2 biasing. In order to account for these findings two different taxonomies were proposed. The first, was the Thematic Role Taxonomy, which was introduced by Brown & Fish (1983a) and was later revised by Au (1986). The second taxonomy was the Linguistic Category Model (Semin & Fiedler, 1988). Each of these will be discussed in turn, followed by a review of the relative merits and pitfalls of the taxonomies.

### 1.2.1. Thematic Role Taxonomy

Thematic roles relate the arguments of a verb to the meaning of that verb. As a construct, they have played a significant role in linguistics since the ‘60’s, exemplified by the work of Fillmore (1968) and Jackendoff (1972). Although there has been some debate about whether thematic roles are syntactic, semantic or conceptual entities (Carlson & Tanenhaus, 1988; Dowty, 1991), it is thematic roles at a semantic/conceptual level that are crucial to the implicit causality bias.

According to the Thematic Role Taxonomy, verbs are classified according to the thematic roles they have. There are four crucial thematic roles: Agent, Patient, Experiencer and Stimulus. An Agent is someone who instigates an action; a Patient is someone, or something, which is affected by an action; an Experiencer is someone having an experience; and finally a Stimulus is someone, or something, giving rise to an experience. Based on these thematic roles, Brown & Fish (1983a) divided implicit causality verbs into one of three categories. A verb was either an Agent-Patient verb, an Experiencer-Stimulus verb or a Stimulus-Experiencer verb. Agent-Patient verbs are those

where the sentence subject plays the role of the Agent and the sentence object the role of the Patient, for example, *telephone* and *apologise to*. Experiencer-Stimulus verbs have the Experiencer role in the subject position and the Stimulus in the object position (e.g. *fear*, *adore*); whereas Stimulus-Experiencer verbs have the Stimulus in subject position and the Experiencer in object position (e.g. *frighten*, *fascinate*). It is clear that this categorisation maintains the distinction between actions and states, with Agent-Patient verbs falling in the action category and Experiencer-Stimulus and Stimulus-Experiencer verbs falling in the state category.

The advantage of using this categorisation scheme is that it allows one to capture generalisations. There are regularities that can be expressed which go beyond the sentence position of the nouns in a sentence. Thematic roles provides a way of capturing these regularities (Dowty, 1991). For example, by using the distinction between Experiencer-Stimulus and Stimulus-Experiencer verbs one can investigate whether the implicit causality bias is a bias to attribute causality to a sentence position (NP1 or NP2) or for entities playing a particular semantic role (Stimulus or Experiencer).

It appears that the implicit causality bias is actually a preference for a thematic role, rather than sentence position. Experiencer-Stimulus and Stimulus-Experiencer verbs both bias towards the Stimulus. Hence, Experiencer-Stimulus verbs are NP2 biasing and Stimulus-Experiencer verbs are NP1 biasing. This shows that it is the thematic role rather than sentence position which is crucial in judgements of causality. Further, this preference for the Stimulus has proven to be a robust finding.

Early studies, such as those of Brown & Fish (1983a; 1983b) further argued that Agent-Patient verbs have a bias towards the Agent (i.e. NP1), in the same way that state verbs have a bias towards the Stimulus. However, it is not the case that all action verbs lead to NP1 attributions. In fact, a distinct sub-set of action verbs lead to NP2 attributions.

Au (1986) distinguished action verbs into two categories. As well as the original set of Agent-Patient verbs, Au distinguished a second set of verbs: Agent-Evocator verbs<sup>1</sup>. This group consists of actions where the sentence object has occasioned the event described by the verb. Examples of Agent-Evocator verbs are *scold*, *praise* and *punish*. In each case, the sentence object has done something, which has led to the sentence subject reacting in some way. For example, *Ted praised Mary* is the result of Mary

having performed some action which is good (or not performing a bad action), which led to Ted speaking to Mary in a particular way.

Agent-Evocator verbs seem to be tapping into a different notion of causality from that involved in Agent-Patient verbs. In the case of Agent-Evocator verbs, someone is held causally responsible when they have *elicited* a behaviour. This is distinct from the causation involved in Agent-Patient verbs where the person held causally responsible is the one who *initiated* the event. In this sense, Agent-Evocator verbs are more similar to state verbs, because for state verbs the Stimulus elicits the response from the Experiencer, rather than the Stimulus, or the Experiencer, initiating the state. This is clearest for the case of Stimulus-Experiencer verbs. For example, in the case of *Ted attracts Mary*, it seems that Ted has done something, or possesses some distinguishing characteristic which elicits attraction in Mary. This has obvious parallels with Agent-Evocator verbs, such as in *Ted praised Mary*, where Mary has done something to elicit praise from Ted.

In order to be able to distinguish Agent-Evocator from Agent-Patient verbs Au proposed three different criteria. The first criteria came from Fillmore (1971) and applied to “verbs of judging”. Fillmore argued that the person who is held causally responsible for an event is the one who is presupposed to have been responsible for the situation leading to the event. For example, in the sentence *Ted praised Mary* the speaker presupposes and expects the addressee to presuppose that (1) something good happened, and (2) Mary was responsible for it. The second criterion used is that of intention, and is taken from Heider (1958). Au argued that “The agent’s intention presupposed in action verbs may [therefore] predict their implicit causality” (Au, 1986, 107). So, according to Au, Agent-Patient verbs are intended, or self-initiated actions, whereas Agent-Evocator verbs are accidents or reactions to the situation. Finally, Au used the Initiating-Reacting distinction proposed by Osgood (1970). Greene & McKoon (1995) also discussed this criterion, where they proposed that the roles of initiator and reactor were different for the two types of verb. Whereas the initiator in Agent-Patient verbs is the Agent, who is also the sentence subject; for Agent-Evocator verbs, the initiator is the Evocator, which is also the sentence object. The Agent in Agent-Evocator verbs is the reactor.

The use of Agent-Evocator verbs has been criticised by some theorists. Semin & Fiedler (1988; 1989) argued that Agent-Evocator verbs are not representative of the general class of action verbs, and consequently do not encourage their use in the study

---

<sup>1</sup> Au (1986) labelled the two sets of action verbs Action-Agent verbs and Action-Patient verbs. However, due to the confusing nature of these names, Rudolph & Forsterling’s (1997) Agent-Patient and Agent-

of implicit causality. They argued that this class of verbs were better described as “speech acts”, or “a mental attitude”. Semin & Fiedler (1989) even suggested that such verbs were “metalinguistic devices” used to express attribution to the sentence object.

There are a number of objections to these claims. First, it is not the case that Agent-Evocator verbs are unrepresentative of interpersonal events, although they may be unrepresentative of action verbs. In a corpus study by Rudolph & Forsterling (1997), they found that out of 1,100 interpersonal verbs approximately 55% were Agent-Patient, 12% Agent-Evocator, 18% Experiencer-Stimulus and 15% Stimulus-Experiencer. It can be seen that the number of Agent-Evocator verbs is comparable to the number of state verbs. Hence, Agent-Evocator verbs are not unrepresentative of interpersonal events. Agent-Evocator verbs are as important as any of the other verb categories.

Second, Semin & Fiedler claimed that Agent-Evocator verbs are atypical because they are speech acts<sup>2</sup>. However, not all Agent-Evocator verbs are speech acts (*e.g. hire, stare at, hit, pick up, etc.*), and more importantly some speech acts are Agent-Patient verbs, (*e.g. flatter, greet, lie to, etc.*).

Third, Semin & Fiedler claimed that Agent-Evocator verbs were better described as a mental attitude. It is not quite clear what they meant by this, but if it is taken to mean something like the action is associated with cognition, then there are Agent-Evocator verbs which do not involve a mental attitude. For example, the Agent-Evocator verb *pick up* does not involve a mental attitude any more than the Agent-Patient verb *greet* does.

Finally, the claim that Agent-Evocator verbs are metalinguistic devices used to attribute causality is very unclear. Agent-Evocator verbs are no more metalinguistic devices than are Agent-Patient, Experiencer-Stimulus or Stimulus-Experiencer verbs. Consequently, the arguments proposed by Semin & Fiedler (1988; 1989) against the use of Agent-Evocator verbs are not well founded.

In their review paper, Rudolph & Forsterling (1997) advocate the distinction between Agent-Patient and Agent-Evocator verbs. However, with a caveat: Rudolph & Forsterling point out that the above criteria used to classify Agent-Evocator verbs are not as “sophisticated” as those used to classify thematic roles. This is certainly true. Although Fillmore’s “verbs of blame” are easily identified, other Agent-Evocator verbs,

---

Evocator labels will be used instead.

<sup>2</sup> Semin & Fiedler use “speech act” to mean something like “verb of communication”.



which are classified using notions such as intention and the initiator-reactor distinction, are not so easy to categorise.

Although the Thematic Role Taxonomy distinguishes the causal biases associated with verbs, there is a problem with using this classification. The use of thematic roles in order to predict the implicit causality in verbs is a circular enterprise. For example, according to the Thematic Role Taxonomy for Agent-Patient verbs the Agent is the cause of the event. So, for simple active sentence with the verbs, *hit*, *chase* or *telephone* cause is attributed to NP1. However, the roles of Agent and Patient are not neutral with respect to causality. One of the defining attributes of Agent is cause. That is, the thematic role of Agent is defined by establishing who the cause of the event is. Hence, what thematic roles show is that empirical demonstrations of causality parallel linguists' intuitions about the cause of events.

The fact that thematic roles are not independently predictive of the implicit causality associated with interpersonal verbs is further demonstrated by the introduction of Agent-Evocator verbs by Au (1986). For Agent-Evocator verbs, it is clear that the thematic roles of Agent and Patient are not sufficient to predict who the cause of the event is. Instead of the expected NP1 bias, a NP2 bias is observed. Hence, a new role is introduced, to deal with the exceptions. The classification of Agent-Evocator verbs is based on intuitions as to the cause of the event. What is missing are more independent tests of causal role. This argument concerning the circularity of the Thematic Role Taxonomy may become clearer by examining the Linguistic Category Model.

### **1.2.2. Linguistic Category Model**

An alternative to the Thematic Role Taxonomy was proposed by Semin & Fiedler (1988; 1989; 1992; Semin & Marsman, 1994). Like the Thematic Role Taxonomy, the Linguistic Category Model divides interpersonal verbs into one of four categories. The first category is Descriptive Action Verbs. A Descriptive Action Verb makes reference to a concrete action, which manifests itself in a particular sort of activity, and has a clear beginning and end. It is generally a neutral description of an action, that is, it does not carry an evaluative component. Examples of Descriptive Action Verbs are *call*, *meet*, *see*, and *visit*. Verbs in this category are weakly NP1 biasing.

The second category is Interpretive Action Verbs, which also make reference to an action, but in this case the verb refers to a general class of behaviours, rather than a

particular sort of behaviour. These behaviours do not have a clear beginning or end, and do have an evaluative component (*e.g. cheat, imitate, help, inhibit*). Verbs in this category are strongly NP1 biasing.

The third category is State Action Verbs. These are similar to Interpretive Action Verbs except that State Action Verbs refer to states (not actions) that are caused by the observable action of an Agent. The verb describes the resultant emotion of the Patient. Examples of such verbs are *fascinate, surprise, bore* and *thrill*. State Action Verbs are also strongly NP1 biasing.

The final category of verbs is State Verbs. Unlike State Action Verbs, State Verbs simply refer to mental or emotional states, for example, *admire, hate*, and *abhor*. On the other hand, like Interpretive Action Verbs, State Verbs, have no clear beginning or end and do include an evaluative component. State Verbs are strongly NP2 biasing.

Again, as in the Thematic Role Taxonomy a distinction is made between actions and states. Descriptive Action Verbs and Interpretive Action Verbs are both examples of action verbs, while State Action Verbs and State Verbs are of course state verbs. The Linguistic Category Model distinguishes action and state verbs using the tests introduced earlier, for example, the applicability of the imperative form, the progressive form *etc.*. Additionally, Semin & Marsman (1994) introduce a way of distinguishing State Action Verbs from State Verbs by means of the “but”-test (Johnson-Laird & Oatley, 1988). Although one can say of State Verbs ‘*I like Mary, but I don’t know why*’; one cannot do the same with State Action Verbs. That is, ‘*Mary surprises me, but I don’t know why*’ sounds distinctly awkward.

Semin and colleagues argue that the Linguistic Category Model is superior to the Thematic Role Taxonomy because the former classification scheme uses a number of independent criteria for classifying verbs. It does not rely on a notion of causality. For example, Descriptive Action Verbs are those that refer to a concrete action, have a clear beginning and end, and do not carry an evaluative component. None of these criteria relies on using the notion of cause. On the other hand, thematic roles presuppose causality, as was seen in the previous section. Therefore, the Linguistic Category Model avoids circularity.

Further, the Linguistic Category Model predicts a wide range of inferences associated with interpersonal verbs, as well as the implicit causality verb bias. The Linguistic Category Model has been applied to a number of phenomena in social cognition, including Actor-Observer differences (Semin & Fiedler, 1989), the

egocentric bias (Fiedler, Semin & Koppetsch, 1991), intergroup relations (Maass & Arcuri, 1992), and constructive memory effects (Semin & Fiedler, 1988).

The Linguistic Category Model has also predicted other types of inferences associated with verbs. Semin & Fiedler (1988; 1992) argued that verbs lie on a concrete-abstract dimension, with Descriptive Action Verbs representing the most concrete situation, and State Verbs the most abstract<sup>3</sup>. A number of inferences are associated with this dimension. Semin & Fiedler (1988) showed that as one moves from Descriptive Action Verbs to Interpretive Action Verbs through to State Action Verbs and finally to State Verbs, the properties associated with the verb change. Participants think the activity described by the verb is more enduring, gives more information about the subject, but less information about the situation. The verb is considered to be less verifiable and more highly disputable. The Linguistic Category Model has classified verbs so that it predicts not only the implicit causality bias, but also numerous other inferential properties of the verb.

The Linguistic Category Model has been criticised by Edwards & Potter (1993). They argue that the defining features used to categorise the verbs do not accurately reflect the properties of individual verbs. For example, they argue that some Interpretive Action Verbs do have a clear beginning and end (*e.g. dominate*). Further, there are Descriptive Action Verbs that do not have invariant physical features (*e.g. kiss, kick etc.* can be accomplished in a number of ways) and Descriptive Action Verbs that do include an evaluative component (*e.g. kiss* [positive], *hit* [negative]). Similarly, not all State Action Verbs tap an “implicit action frame”. An emotion, such as *fascinate*, can occur because of a characteristic of a person, not necessarily a preceding action. Ted can fascinate Mary because he is kind, as well as because he gives to charity.

However, these counter-examples are controversial. For example, it is not clear that *dominate* is an example of an Interpretive Action Verbs, a class of verbs which are not supposed to have a clear beginning and end. The verbs *kiss* and *kick* do appear to have invariant physical characteristics, in that *kiss* always involves the mouth and *kick* always involves the foot. On the other hand, Edwards & Potter’s (1993) arguments about the existence of an evaluative component in some Descriptive Action Verbs and the lack of an implicit action frame for some State Action Verbs do seem to be accurate.

---

<sup>3</sup> The Linguistic Category Model actually has a fifth category, which consists of adjectives. In actual fact, it is the adjectives that are the most abstract thing in this model.

Still, Semin & Fiedler (1992) show that their criteria for verb category membership gives 90% correct identification for verbs.

### 1.2.3. Discussion

In actual fact, the Linguistic Category Model is not incompatible with the Thematic Role Taxonomy. Both classifications distinguish action and state verbs. The subdivisions of the state verbs are the same in both classification schemes. NP1 biasing state verbs are called Stimulus-Experiencer in the Thematic Role Taxonomy but State Action Verbs in the Linguistic Category Model; while NP2 biasing state verbs are Experiencer-Stimulus verbs and State Verbs respectively.

The case of action verbs is more complex. The Linguistic Category Model distinguishes action verbs into Descriptive Action Verbs and Interpretive Action Verbs, while the Thematic Role Taxonomy distinguishes Agent-Patient and Agent-Evocator verbs. In this case, there is not a straightforward parallel between the classification schemes. The Linguistic Category Model distinguishes between those verbs which are strongly NP1 biasing and those which are weakly NP1 biasing, in contrast with the Thematic Role Taxonomy which distinguishes NP1 biasing action verbs from NP2 biasing ones.

In reality, it is not clear what is the appropriate strategy for distinguishing action verbs. Brown & Fish (1983a) originally predicted that all action verbs were NP1 biasing. However, a little later, Brown & Fish (1983b) reported a production study in Chinese where they found only five of twelve verbs they classified as Agent-Patient showed a NP1 bias. *Apologise to*, *cheat*, *compete with*, *defy*, and *flatter* were the only verbs that showed a NP1 bias, with only the first three verbs producing a reliable bias. Of the remaining seven verbs, two were reliably NP2 biasing (*criticise*, *protect*), and were in actual fact Agent-Evocator verbs. The remaining verbs did not produce a strong bias in either direction (*disobey*, *dominate*, *harm*, *help*, *slander*). If the Linguistic Category Model could independently predict weakly biasing verbs, in other words, if the weakly biasing verbs are Descriptive Action Verbs, then a strong case could be made for retaining the distinction between Descriptive Action Verbs and Interpretive Action Verbs in future research. However, in many studies including Brown & Fish (1983b) the verbs that do not produce a strong bias are not Descriptive Action Verbs, but rather Interpretive Action Verbs.

Similarly, Au (1986) found that some Agent-Patient verbs do not show a strong NP1 bias at all. If the Descriptive Action/Interpretive Action distinction is a real one then weakly biasing action verbs should be Descriptive Action Verbs; while strongly biasing verbs should be Interpretive Action Verbs. Au found that in the active voice, the following verbs did not show a reliable NP1 bias: *betray*, *cheat*, *harass*, *harm*, *hit*, *interrupt*, *order around* and *slander*. These verbs should all be Descriptive Action Verbs. Instead, only *hit* is a Descriptive Action Verb. From the strongly NP1 biasing verbs *corrupt*, *flatter*, *help* and *telephone*, one of them is in fact a Descriptive Action Verb (*telephone*), although they should all be Interpretive Action Verbs. Hence, it is unclear whether the distinction between Descriptive Action Verbs and Interpretive Action Verbs is a useful one.

In a meta-analysis of 292 action verbs employed in various experiments Rudolph & Forsterling (1997) did find a tendency for Interpretive Action Verbs to evoke NP1 attributions, and for Descriptive Action Verbs to produce intermediate attributions. In order to determine whether this distinction was useful, they conducted a 2-way ANOVA distinguishing Agent-Patient/Agent-Evocator verbs from Interpretive Action Verbs/Descriptive Action Verbs. They found a reliable effect of both distinctions on bias. However, the Agent-Patient/Agent-Evocator distinction accounted for 31% of the variance, whereas the Interpretive Action Verbs/Descriptive Action Verbs distinction only accounted for 2% of the variance.

Therefore, the effectiveness of retaining the distinction between action verbs made by the Linguistic Category Model is not altogether clear. The categories of Descriptive Action Verbs and Interpretive Action Verbs may be useful in resolving some of the inconsistent findings surrounding action verbs. However, there are still some inconsistencies remaining to be explained (the fact that some of Descriptive Action Verbs are strongly biasing and some Interpretive Action Verbs are weakly biasing).

To summarise, the Thematic Role Taxonomy and the Linguistic Category Model are compatible in many respects. For example, both agree on the distinction between action and state verbs. However, even the things that both schemes agree on have problems. Take the distinction between action verbs and state verbs. First, some interpersonal verbs have equally strong active and stative connotations, for example, *entertain* and *honour* (Hoffman & Tchir, 1990). It is interesting to note, however, that this equality is not always preserved under negation. So, for example, Hoffman & Tchir

(1990) argued that *surprise* has equally strong active and stative senses. However, *does not surprise* has a dominant stative sense.

Second, it is argued that state verbs describe an unobservable event, that is something that is “in the mind”, which is best described by terms such as TO EXPERIENCE, TO FEEL, TO BE and TO INSPIRE (*e.g.* Au, 1986). However, some unobservable events can be described better by the terms TO DO, TO ACT, TO BEHAVE, TO PERFORM or TO EXPRESS, which are typically taken to be properties of action verbs. Examples of such verbs are *evaluate* and *think about* (Hoffman & Tchir, 1990).

Hoffman & Tchir (1990) argued that the lexical tests involving the progressive tense and the imperative form only classify verbs correctly around 90% of cases. There are state verbs which can take the progressive form, such as *enjoy*, in *I am enjoying the view*; and others which take the imperative form, such as *trust* in *Trust me!*. It is not known what the implicit causality bias is for such verbs. So, there are problems with the distinction between actions and states.

Both the Thematic Role Taxonomy and the Linguistic Category Model also correspond in the distinction made for state verbs. There are strongly NP1 biasing verbs (Stimulus-Experiencer or State Action Verbs) and strongly biasing NP2 verbs (Experiencer-Stimulus or State Verbs). This distinction has held up well.

The major difference between the two schemes arises in the treatment of action verbs. Whereas the Linguistic Category Model distinguishes between action verbs which are strongly NP1 biasing from those which are weakly NP1 biasing (Interpretive Action Verbs from Descriptive Action Verbs), the Thematic Role Taxonomy distinguishes NP1 biasing verbs from NP2 biasing verbs (Agent-Patient from Agent-Evocator). Semin & Fiedler (1988; 1989) argued that Agent-Evocator verbs are not typical action verbs, and therefore should not be included in a taxonomy of interpersonal verbs. However, Rudolph & Forsterling (1997) showed that Agent-Evocator verbs are as common as other interpersonal verbs such as Stimulus-Experiencer and Experiencer-Stimulus verbs. Finally, distinguishing strongly and weakly biasing verbs may prove to be useful when discrepancies in findings occur (Rudolph & Forsterling, 1997).

### 1.3. Effects that moderate the implicit causality bias

Given that there is a bias to attribute causality based on the information given by the verb, the next question is what other factors influence the implicit causality bias. In this section, a number of factors which researchers have used to determine the breadth of the implicit causality bias will be described. First, presentation format will be discussed. This includes a discussion of the effects of presenting a sentence in the active versus passive voice, interrogative versus declarative form and the effect of explicitly negating a sentence. Second, semantic factors that effect the implicit causality bias will be discussed. Finally, contextual factors, such as the gender of the characters involved in the interaction will be reviewed.

#### 1.3.1. Manipulating presentation format

In this section, three different presentation formats will be considered. First, the active versus passive voice, second the interrogative versus declarative form and finally affirmative versus negative phrasing. In each case, the data will be evaluated in order to establish whether changes in presentation format lead to changes in causal attributions.

##### 1.3.1.1. *Active versus passive voice*

When researching the implicit causality bias, the standard presentation format is a simple active sentence. However, there are a few studies that have compared the active voice with the passive voice, that is whether there is any difference between *Ted VERBED Mary* and *Mary IS VERBED BY Ted*. The first to do this was Garvey, Caramazza & Yates (1976). They found that verbs which are strongly NP1 biasing in the active voice are weakly NP2 biasing in the passive form; weak NP1 biasing verbs are stronger NP2 biasing in the passive; and weak NP2 biasing verbs are weakly NP1 biasing in the passive. However, these findings cannot be considered generalisable as they considered such a small number of verbs. There were only three strongly biasing verbs, and two weakly biasing, one with a weak NP1 bias and one with a weak NP2 bias.

The first systematic study to investigate the effects of using the passive voice was reported in Brown & Fish (1983a). They found that even when a sentence is presented in the passive voice the underlying bias is retained. That is, sentence subject

attributions in the active voice become sentence object attributions in the passive, and vice versa. This can be interpreted as the implicit causality verb bias being a preference for a thematic role, rather than sentence position. This is consistent with other studies also.

Au (1986, Experiment 1) found a general preservation of attributions to thematic roles. Further, this preservation was strongest for state verbs. For sentences in the active and passive voice participants attributed cause to the Stimulus. However, action verbs did not show as clear a pattern. For example, Au found that for Agent-Evocator verbs not all verbs showed a reliable bias to the Evocator in the active or passive voice. This was even more apparent for Agent-Patient verbs. Only 4 out of 12 verbs showed a reliable bias in the active voice (*telephone, help, flatter, corrupt*), and 5 in the passive (*telephone, hit, harm, flatter, betray*). Crucially, only two Agent-Patient verbs preserved their bias when passivised (*telephone, flatter*). Overall, Au's results are consistent with Brown & Fish's claim that verbs retain causal bias towards a thematic role, when presented in the passive voice.

A fourth study, which investigated the effects of the passive form on attributions, is that of Kasof & Lee (1993). While Kasof & Lee's overall findings are the same as Brown & Fish and Au, in that there was an overall preference for a thematic role, regardless of voice, they found a systematic difference between the active and the passive voice. When NP1 biasing verbs in the active voice (Agent-Patient and Stimulus-Experiencer verbs) were presented in the passive voice there was a relative weakening of the bias to the thematic role (Agent and Stimulus respectively). On the other hand, when verbs that are NP2 biasing in the active voice (Experiencer-Stimulus verbs) are presented in the active voice there was a relative increase in the strength of the bias to the thematic role (i.e. Stimulus). These findings are also reported by Hoffman & Tchir (1990 Experiment 1), who report a stronger bias to the sentence subject than sentence object. (The interpretation of these findings are discussed further in Chapter 2.)<sup>4</sup>

Kasof & Lee's findings are also in accord with the results of Garvey *et al.* reported at the beginning of this section. The three strongly biasing NP1 verbs used by Garvey *et al.* are all Agent-Patient verbs. To recap, they found that strongly biasing NP1 verbs in the active voice became weakly biasing NP2 verbs in the passive voice.



This is exactly what Kasof & Lee found. When Agent-Patient verbs are passivised, there is a relative weakening of the bias.

Garvey *et al.* also found that the weakly NP2 biasing verb in the active voice (an Experiencer-Stimulus verb) became weakly NP1 biasing in the passive voice. A closer look at the data shows that the bias to choose the Stimulus is stronger in topic position. In the active voice, there are 29/52 NP2 completions while in the passive voice there are 37/50 NP1 completions. However, Garvey *et al.* found that a weakly biasing NP1 verb (an Agent-Patient verb) becomes a strongly NP2 biasing verb when passivised. Kasof & Lee found a weakening of the bias when the Agent was moved to object position. However, considering that Garvey *et al.* only present data from one verb, this shift could just be random fluctuation.

Rudolph & Forsterling (1997) conducted a meta-analysis using verb type and sentence form as independent variables. They found a significant interaction between these two factors. Closer examination of the results showed that there was no difference in the strength of the causal bias when Agent-Patient verbs were presented in the active voice or the passive voice. Stimulus-Experiencer verbs showed the opposite pattern of results from Kasof & Lee: when a Stimulus-Experiencer verb was presented in the passive voice the strength of the bias decreased. However, Rudolph & Forsterling found the opposite to be true. The strength of the bias increased when Stimulus-Experiencer verbs were presented in the passive voice. Finally, for Experiencer-Stimulus verbs Rudolph & Forsterling found that the strength of the bias did increase, as was the case in the Kasof & Lee study.

To summarise, studies investigating the effects of active and passive voice on the implicit causality verb bias have shown that regardless of the form of the sentence the causal bias is preserved. The thematic role that is the cause in the active voice is the cause in the passive voice. Kasof & Lee found that there was a relative increase in the bias when the thematic role associated with the cause was in subject position. However, based on a meta-analysis conducted by Rudolph & Forsterling this does not appear to be the case. Across verb types, there is no reliable increase in attributions associated with sentence position.

---

<sup>4</sup> Hoffman & Tchir's findings depart from those of Kasof & Lee in the case of Agent-Patient verbs. The latter found a decrease in the strength of the bias when Agent-Patient verbs were presented in the passive voice, while the former found no difference.

### 1.3.1.2. *Interrogative versus declarative form*

Another change that can be made to the presentation format is to present a sentence in the interrogative or declarative form. So, a sentence could be presented in the typical declarative form of *Ted VERBED Mary*, or in the interrogative form *Ted VERBED Mary?*.

Only one study to date has investigated whether this presentation format influences causal attributions. Garvey, Caramazza & Yates (1976) compared attributions given when sentences were presented in the interrogative form to those in the declarative form. They found no differences in causal attributions as a result of this manipulation, although it should be noted that only a small number of items were used.

### 1.3.1.3. *Affirmative versus negative phrasing*

A third manipulation of presentation style is whether the sentence was presented in an affirmative or negative form, *i.e.* whether there was an explicit negation in the sentence. In all of the studies reported in this section, participants were presented with either the affirmatively phrased *Ted VERBED Mary*, or the negatively framed *Ted DID NOT VERB Mary*, and then asked to make a causal attribution.

Garvey, Caramazza & Yates (1976) manipulated phrasing in this way, and found an attenuation of the implicit causality bias. However, as before, they only tested a limited number of verbs. In this the verbs were *kill*, *fear*, *blame* and *pick up*. All four verbs are NP2 biasing when in the affirmative form, and weakly NP2 biasing in the negative form.

Hoffman & Tchir (1990) also provided some evidence for an attenuation of the causality bias when a sentence was negatively phrased. However, this appeared to be the case for state verbs only. There was not a consistent weakening of the bias with Agent-Patient verbs. Similarly, Franco & Arcuri (1990) did not find an effect of explicit negation on attributions for any verb type.

Finally, Gordon & Hendrick (1997) found a consistent effect of phrasing on judgements of causality. In accord with the findings of Garvey *et al.*, and the findings of Hoffman & Tchir, Gordon & Hendrick found that when a sentence with an implicit causality verb was phrased negatively there was an attenuation of the implicit causality bias. Overall, the data supports the hypothesis that negatively phrased sentences lead to weaker causal attributions than positively phrased sentences.

### 1.3.2. Manipulating semantic properties of the verb

#### 1.3.2.1. *Semantic valence*

In this section, a review of the effects of the semantic valence on causal attributions will be provided. Semantic valence is used to refer to the evaluative dimension of verb meaning. This dimension can be traced back to the seminal work of Osgood, Suci & Tannenbaum (1957) where they investigated the connotative meaning of words. They found that connotative meaning could be classified on three dimensions: evaluation, potency and activity. The evaluative dimension of a word is the same as semantic valence, and indicates whether the word is positive or negative. That is, whether the word is viewed as a “good thing” or a “bad thing”.

Au (1986) demonstrated that implicit causality verbs include a strong evaluative dimension to their meaning. She conducted a sorting task, where subjects were given a pile of either state or action verbs. Subjects were then asked to sort the verbs into as many or as few groups as they liked, as long as the groups were meaningful. For both state and action verbs the first dimension, which accounted for the most variance, was an evaluative one. Au called this the Good-Bad dimension, as the dimension had pleasant events at one extreme and unpleasant events at the other.

Some researchers have attempted to determine whether semantic valence, or the evaluative component of verb meaning, influences attribution. The effect of semantic valence on causal attributions has been tested by firstly classifying verbs as positive, negative or neutral. This is done either by explicitly asking subjects to rate a verb on a scale, which has as values either positive-negative (Franco & Arcuri, 1990; Hoffman & Tchir, 1990; Mannetti & DeGrada, 1991), or agreeable-disagreeable (Brown & VanKleeck, 1989). Alternatively, a more complete set of differential scales are used, where the verb is rated on a number of variables (*e.g.* good-bad, happy-sad, beautiful-ugly) which are aggregated to give a score for semantic valence (Semin & Marsman, 1994).

The first study which investigated semantic valence was, once again, Garvey, Caramazza & Yates (1976). They found that the semantic valence of verbs did not influence causal attributions. However, as with all the other variables Garvey *et al.* tested these results were confounded by a number of factors, such as the small sample size. The remaining studies investigating semantic valence have found that verbs with negative valence produce higher levels of NP1 attributions than those with either

positive or neutral valence (e.g. Franco & Arcuri, 1990; Mannetti & DeGrada, 1991; Semin & Marsman, 1994). However, there are a number of qualifications.

Franco & Arcuri (1990) used the Linguistic Category Model in order to classify the Italian verbs they studied, specifically they were interested in Interpretive Action Verbs and State Verbs. In their study, they found an effect of semantic valence for Interpretive Action Verbs only. Negative Interpretive Action Verbs gave stronger NP1 attributions than positive ones. However, the interpretation of these findings are unclear, as there is a problem in their choice of positive Interpretive Action Verbs. Three of the four positive verbs (*thank*, *help*, *encourage*) chosen by Franco & Arcuri are in fact, Agent-Evocator verbs, which have been shown to bias to the second NP, not the first NP (Au, 1986). Therefore, a difference in the attribution pattern associated with negative and positive Interpretive Action Verbs in this study may in fact be due to the use of Agent-Evocator verbs in the positive category, and the use of Agent-Patient verbs in the negative category.

Mannetti & DeGrada (1991) also investigated the effect of semantic valence on causal attributions in Italian. Like the above study, they found that verbs with negative semantic valence are associated with higher NP1 attributions than positive or neutral verbs (although only the difference between negative and neutral verbs is reliable). They also suggested that the difference in attributions due to semantic valence decreases as the verb becomes more abstract. That is the largest difference due to semantic valence occurs for Descriptive Action Verbs (positive = 0.67; negative = 0.39), then Interpretive Action Verbs (positive = 0.79; negative = 0.54), followed by State Action Verbs (positive = 0.79; negative = 0.89), and State Verbs (positive = 0.20; negative = 0.10). However, although there was a trend in the data, the differences were not reliable.

Finally, Semin & Marsman (1994) also found that negative verbs have a stronger NP1 bias than positive verbs, this time in Dutch. However, this difference is only apparent in event instigation inferences, not in dispositional inferences.

Semin & Marsman argue that previous experiments investigating implicit causality have confounded two separate inference processes. They argue that there is confusion in the literature as to what “cause” means in “implicit causality”. Early work on implicit causality equated the cause of an event with the instigator of the event (Garvey & Caramazza, 1974; Garvey, Caramazza & Yates, 1976).

“The property itself was identified as ‘implicit causality’, a feature which selects for one or the other of the available candidate nouns *as primarily responsible for instigating the action or state denoted by the antecedent clause*” (Garvey, Caramazza & Yates, 1976, 240-241, [italics added]).

Further, Semin & Marsman argued that the Thematic Role Taxonomy actually defines Agency in terms of the character that instigates an event. This is exemplified by Dowty (1991), who argues that the role of Agency has been subdivided in many ways, but always with the inclusion of an instigative property.

Semin & Marsman argue that people make two types of causal judgements, when explaining the behaviour of people. One of these is indeed based on deciding who instigated an event; however, a second type of cue that can be used is dispositional information.<sup>5</sup> For example, Brown & Fish (1983a) state:

“For actions.... we attribute causality to dispositions in the Agent... Thus an instance of helping seems to us to arise out of a disposition to be helpful,... rather than a tendency to instigate or merit help in certain persons. For mental states,... we attribute causality primarily to dispositions in the Stimulus... Thus, an instance of someone being charmed seems to us to arise out of a disposition of a Stimulus to charm, to be charming... rather than a by a tendency to be easily charmed by certain persons.” (Brown & Fish, 1983a, 247).

In a series of experiments, Semin & Marsman attempted to separate out different cues to causality. When they asked event instigation questions (*i.e.* who brought about the event or state?), they found strong verb effects. Interpretive Action Verbs and State Action Verbs both gave strong NP1 attributions and State Verbs gave strong NP2 attributions. There was, also, an effect of semantic valence consistent with the previous studies when an event instigation question was asked. Negative verbs led to stronger inferences than positive verbs. However, when implicit causality was defined as a dispositional inference (*i.e.* did the event occur because of something about the sentence subject or the sentence object?) then there was an overall NP1 bias, with the strongest effects for Interpretive Action Verbs, followed by State Action Verbs, and finally State Verbs. The verb effect on inferences was much weaker than it was in the event instigation case. Furthermore, using a dispositional definition of causality, there was no effect of the semantic valence of the verb on attributions.

Although this seems like a plausible account of the differences found in previous experiments, these results conflict with those of Mannetti & DeGrada reported above.

---

<sup>5</sup> For a similar distinction see Hilton, Smith & Kim (1995).

To recap, Mannetti & DeGrada found an effect of semantic valence, even though they investigated dispositional inferences. This is in direct contrast to Semin & Marsman's findings.

Another attempt at differentiating types of causal judgements was made by Lafrance, Brownell & Hahn (1997). They distinguished between causality as initiation, similar to Semin & Marsman's causality as event instigation, from causality as elicitation. Whereas Agent causality is due to the Agent initiating an event, Recipient causality is due to the Patient eliciting or triggering the event. Using this distinction, they found that Agent causality gave more NP1 attributions for negative verbs, whereas Recipient causality gave more NP2 attributions for positive verbs. That is, Ted is more causal in *Ted hit Mary* than he is in *Ted kissed Mary*; whereas Mary is more causal in *Ted likes Mary* than she is in *Ted hates Mary*.

This difference was found when they used an event instigation question: '*Probably NP1 verbs (many.....few) other people*' (initiation) and '*Probably (many.....few) other people verb NP2*' (elicitation), as well as when they asked a dispositional question: '*X is the kind of person who abhors people*' (initiation) and '*X is the kind of person whom people abhor*' (elicitation). Whereas Semin & Marsman used a forced-choice procedure in order to determine the differences between event instigation and dispositional inferences, Lafrance *et al.* used two independent scales and asked participants how causal the characters were both as initiators and as elicitors.

In a meta-analysis of verb type and semantic valence (positive, negative and neutral), Rudolph & Forsterling (1997) found no main effect for semantic valence and no interaction of valence and verb type. However, it seems that this analysis may have masked some real differences, especially if Semin & Marsman and Lafrance *et al.* are correct and the effects of semantic valence are only observed in particular types of inferences.

To summarise, semantic valence appears to affect causal attribution with in certain circumstances. Negative verbs lead to stronger NP1 attributions than positive verbs – but only when a question about who initiated an event is asked. However, overall, there does not appear to be a reliable effect of semantic valence on causal attributions.

### 1.3.3. Manipulating contextual factors

A large part of the research on the implicit causality bias has focussed solely on the verb. However, there are some studies which have attempted to establish what other cues are used in making causal attributions. In this section, a review of contextual factors on the implicit causality bias will be provided. This will begin with a discussion of how the gender of the characters in a sentence and the gender of participants making the causal attributions affect the implicit causality bias. It will be followed with a look at whether the animacy of characters influences causal attributions. The section will conclude with a review of general context effects.

#### 1.3.3.1. *Gender of characters and participants and the implicit causality bias*

##### 1.3.3.1.1. *Are men more causal?*

In nearly all experiments investigating the implicit causality verb bias, researchers have used nouns that give minimal information. This allows one to control for extraneous variables that may influence causal attributions. Sentences have been presented with placeholders such as *A* and *B* (Van Kleeck, Hillger & Brown, 1988), or proper names (*e.g.* Brown & Fish, 1983a; Au, 1986 *etc.*). When using proper names, the effects of gender are assumed not to be a factor in the attribution process because they are usually male names for both the subject and object nouns of the sentence (*e.g.* Brown & Fish, 1983a; Hoffman & Tchir, 1990).

However, recently researchers have argued that the gender of the characters may be an important cue in making causal attributions. For example, Lafrance, Brownell & Hahn (1997) argued that when there is unequal social power between individuals, then the individual with greater power is also taken to be more causal. Lafrance *et al.* showed that there is a lot of evidence to suggest that gender acts as an indication of status, and thus of power. Typically, males are seen as having more social status than women have, and therefore more power. Based on this, Lafrance *et al.* argued that manipulating gender of characters in interpersonal sentences should lead to differences in causal attributions.

There are studies that have manipulated character's gender when investigating implicit causality, but which do not report any differences based on gender (*e.g.* Au, 1986). Other studies have manipulated gender and have found differences, but have not considered all gender combinations. For example, Franco & Arcuri (1990) only considered male pairs and female pairs and found no effect of gender.

However, Mannetti & DeGrada (1991) considered all gender combinations and measured both causal attributions and confidence ratings. They found no differences in attributions when gender was manipulated, but there was an effect on confidence ratings. Namely, causal attributions to Male-Male gender pairs resulted in the highest confidence ratings, with no consistent differences between any other gender combination.

Finally, Lafrance *et al.* investigated the effects of gender on causal attributions, using their distinction between Agent causality and Recipient causality. Agent causality is when one character initiates the event; Recipient causality is when a character elicits an event (see section 1.3.2.1). Lafrance *et al.* found that men were given higher ratings of Agent causality when they were in a Male-Female combination than when they were in a Male-Male combination; whereas females were given lower ratings when they were in a Female-Male combination than a Female-Female combination. That is, males are viewed as more causal when they are in a different sex interaction, while females are viewed as more causal when they are in a same sex interaction. A difference pattern emerged when considering Recipient causality. In this case, females were always considered more causal than men. That is, women were seen to elicit behaviour in others.

To summarise, there do appear to be differences in causal attributions when the gender of the characters in a sentence is varied. Men are viewed as the initiator of events, while women are viewed as elicitors of events.

#### 1.3.3.2. *Do men and women differ when making causal attributions?*

Just as there are subtle changes in casual attributions when the gender of characters is manipulated, there are changes in attributions when the gender of the participant is taken into consideration. Although many studies have not found any difference in the implicit causality bias based on the participants gender (*e.g.* Kasof & Lee, 1993), Mannetti & DeGrada (1991) and Lafrance *et al.* (1997) do report some differences.

Mannetti & DeGrada found that male participants made stronger NP1 attributions than female participants. However, this difference was confined to attributions for Descriptive Action Verbs. Furthermore, the effect was strongest for a particular subset of verbs, most notably relationship terms. For verbs such as *kiss*, *hold*, *dance*, *hug* and *telephone*, male participants assigned higher levels of causality to NP1 than females. Women did not show a preference for NP1 or NP2, but thought both parties were equally causal. So, while male respondents inferred that Ted was the cause in *Ted kissed Mary*, female respondents were more likely to infer that both parties



caused the kiss. Mannetti & DeGrada also found a difference in the confidence ratings associated with the attributions. Females were more confident about their causal attributions than men. So, although men made stronger attributions of causality, they were less confident about their attributions.

Lafrance *et al.* (1997, Experiment 1 and 2) also found that the gender of participants affected attributions. Males made stronger NP1 attributions for initiation inferences, as found by Mannetti & DeGrada. Similarly, men made stronger NP2 attributions for elicitation inferences. That is, men seem to make stronger attributions of causality than women do, regardless of the type of cause involved.

#### 1.3.4. *Is the implicit causality bias “interpersonal”?*

In the previous section, it was shown that the gender of a character could make a difference to the causal attributions that people make. This leads to the question of whether the implicit causality bias is restricted to interpersonal events. Is the implicit causality bias still apparent when the noun phrases refer to inanimate things?

In a series of experiments, Corrigan (1988; 1992; 1993) has investigated the effects of animacy on causal attributions. She argues that the implicit causality bias is not restricted to interpersonal events, but any interactional event. Corrigan (1988; 1992) paired verbs with different types of noun phrases: (1) the subject and object were both referred to by proper names (*i.e.* proper name, *e.g.* Ted-Paul). Obviously, in this case both noun phrases refer to animate things. (2) The subject and object were of different status, but both were animate (*i.e.* animate-animate, *e.g.* sheriff-bandit). (3) The subject was animate, but the object was inanimate (*i.e.* animate-inanimate, *e.g.* baker-cake). (4) The subject was inanimate, but the object was animate (*i.e.* inanimate-animate, *e.g.* sportscar-woman). And finally (5) the subject and object were both inanimate (*i.e.* inanimate-inanimate, *e.g.* bridge-train). Subjects were given a force-choice rating task, where they had to indicate how likely it was that the event was due to some property of the subject, some property the object, or some other reason.

Corrigan (1988) found that Stimulus-Experiencer verbs always led to Stimulus attributions, regardless of animacy. However, Corrigan (1992) found that when there were animate-animate, and inanimate-animate pairings the Stimulus and Experiencer were equally causal. That is, the NP1 attribution associated with Stimulus-Experiencer verbs was not apparent with certain noun phrases in the latter experiment.

Corrigan (1988) also found Stimulus attributions for Experiencer-Stimulus verbs, which were more pronounced when the Stimulus was inanimate. However, again

a slightly different pattern emerges in Corrigan (1992). In the latter study, Corrigan found that when the Stimulus was inanimate the Stimulus and Experiencer were attributed equal causality. This is exactly the opposite pattern of results found in Corrigan (1988).

For action verbs, the picture is even more complicated. Corrigan classifies action verbs as either Actor verbs or Non-actor verbs. Actor verbs, she states, have derived adjectives associated with the Agent. (These are verbs such as *helped*, which have derived adjectives, such as *helpful*, associated with the Agent; but do not have a derived adjective associated with the Patient.) On the other hand, Non-actor verbs either do not have an associated derived adjective, or have one for the Patient. (These are verbs such as *praised*, which have the derived adjective *praiseworthy* associated with the Patient, but not derived adjective associated with the Agent.) She goes on to claim that Actor verbs refer to unique actions, that is actions that “reside more within individuals”; whereas Non-actor verbs are “more generic” and are “performed by all humans” (Corrigan, 1988, 449). So, for example, *defy* is an Actor verb, because when *Mary defies Ted* then the defiance resides in Mary. However, *pick up* is a Non-actor verb as when *Mary picks up the box* this is not because of some unique ability Mary has, but rather is something all humans can do.

This latter criteria for distinguishing action verbs seems highly dubious, as it is unclear why *criticise* which Corrigan identifies as an Actor verb is due to some unique disposition, but *praise*, which is classified as a Non-actor verb is due to some generic human quality. Consequently, Actor and Non-Actor verbs are best regarded as “adjective referring to Agent” and “no adjective/adjective referring to Patient” categories.

For Actor verbs, Corrigan (1992) found a bias towards the Agent, *i.e.* NP1 attributions, regardless of the animacy of the noun phrases. This was also reported in Corrigan (1988), apart from in one condition where the noun phrases were animate-animate. In that condition there appeared to be equal causality assigned to the subject and object of the sentence.

Non-actor verbs showed an interaction of animacy and verb, in both studies. Corrigan (1988) found that Non-actor verbs received NP1 attributions when the noun phrases were proper names, and when the subject was inanimate (*i.e.* inanimate-inanimate and inanimate-animate). However, they showed NP2 attributions when the noun phrases were animate-animate and animate-inanimate. Corrigan (1992) found NP1 attributions for the proper name condition, animate-inanimate and inanimate-

inanimate pairs. The animate-animate case led to NP2 attributions as in Corrigan (1988). But the inanimate-animate condition produced the same number of attributions to NP1 and NP2. So, in the case of Non-actor verbs in one study the animate-inanimate pair led to NP2 attributions, whereas in the other it led to NP1 attributions. Similarly, in the first study, the inanimate-animate condition gave NP1 attributions, but in the other there was equal causality attributed to the two noun phrases.

Despite the divergent results between the two studies reported above, Corrigan (1992, 361) claimed that “results replicated the earlier study with the exception of sentences with non-actor verbs and AI (animate-inanimate) actor/patient pairings”. She argued that the difference in results between the two studies with respect to these Non-actor verbs could be explained by differences in the typicality of the sentences. Corrigan (1992) showed that there was a positive correlation between typicality and causality. That is, causal ratings were higher when an event was described as more typical than when it was less typical.<sup>6</sup>

Corrigan argued that some noun phrase and verb combinations are high in typicality, whereas others are low in typicality. So a typical Non-actor sentence given animate-inanimate noun phrases would be *Mary praised the movie*, whereas a non-typical sentence would be *Mary praised the tow-truck*. Typical Non-actor, animate-inanimate sentences evoke attributions to NP2, whereas non-typical ones attribute causality to NP1. Corrigan (1992) argued that the Non-actor, animate-inanimate sentences in the Corrigan (1988) study were high in typicality and therefore led to NP2 attributions; while the Non-actor, animate-inanimate sentences in Corrigan (1992) were low in typicality and hence led to NP1 attributions. In support of this hypothesis, one sentence in Corrigan (1992) was rated to be high in typicality and it did show a NP2 attribution as predicted.

The same explanation could be used for the other discrepancies between the two studies. That is, the differences in causal attributions could be due to differences in the typicality of events described by the sentences. This could be confirmed by taking typicality ratings as well as causality ratings and investigating the relationship between the two.

---

<sup>6</sup> It may seem that this is contrary to the findings that the less typical, or abnormal something is, the stronger the causal attributions made to it (*e.g.* Hilton & Slugoski, 1988). However, they argue that stronger attributions are made to the factor that is atypical within an event, whereas Corrigan is evaluating the typicality of the event itself (see Corrigan, 1992).

Overall, Corrigan found that the implicit causality bias for state verbs is less malleable with changes in the animacy of noun phrases than action verbs. Corrigan argues that the changes in animacy that are observed are the result of changes in typicality associated with the events. There is a correlation between these two factors.

#### 1.3.5. General context effects

As well as manipulating animacy in the above studies, Corrigan also manipulated the social status between individuals, by using role descriptions. Earlier, Lafrance *et al.* outlined that differences in social status can change the causal attributions people make (section 1.3.3.1.1). Lafrance *et al.* manipulated gender as a way of changing the social status of the characters, and found that this gender manipulation did affect causal attributions. They then manipulated social status explicitly (Lafrance *et al.*, 1997, Experiment 3) by using occupational roles. In their experiment, the female character could be of higher social status than the male (*e.g.* employer), equal status (*e.g.* co-worker), or lower status (*e.g.* secretary). As before, they found a main effect of gender, such that females were seen as less causal in initiating an event than males, and a main effect of social status, such that the higher the social status the higher the causal ratings attributed to the character. So, changes in role descriptions changed the implicit causality bias.

Garvey, Caramazza & Yates (1976), also, manipulated the power status between the subject and object, such that together with the verbs they presented a congruent or incongruent situation. A congruent sentence represents a situation such as *The father praised the son*, whereas an incongruent situation would be *The son praised the father*. They found that for both NP1 biasing verbs (*argue with*, *confess to* and *join*) and NP2 biasing verbs (*praise* and *criticise*) there was a significant effect of congruency. That is, in the incongruent situation the verb bias was reduced, or even eliminated.

Turnbull & Slugoski (1988) reported a similar study, where they used proper names and role descriptions. For example, they contrasted sentences such as *The headmaster hit Ted* with *Ted hit the headmaster*. Consistent with Garvey *et al.*, they found that changing the noun phrases led to changes in causal attributions.

The interesting thing to note about these studies is that Garvey *et al.* take this sensitivity to noun status as evidence against using a Thematic Role Taxonomy to predict causal bias:

“We propose that the sensitivity of the bias to semantic properties of the noun status in N-V-N collocations... rules out any direct relation between case structure and the implicit causality feature. In the congruent collocation, ‘The assistant argued with the boss’, a strong NP1 bias was obtained. In the non-congruent collocation, ‘The boss argued with the assistant’, a weak NP2 bias was obtained. The underlying case relations, however, would presumably be the same in both variants” (Garvey *et al.*, 1976, 241).

However, a problem with all three of the above studies is that none of them have used adequate experimental items. There are at most four experimental items in each condition. Consequently, the generalisability of these findings is in question.

Semin & Fiedler (1988, Experiment 2) investigated the effects of context on implicit causality judgements. They provided participants with a description of a person: the person was described as being, an introvert, extrovert or Machiavellian personality. The character described was placed in one of three hypothetical situations: at a seminar, at a party or at a business deal. Participants were asked to indicate how likely it was for each item that the character would display 30 different behaviours and dispositions (ten verbs from each of the following categories: Descriptive Action Verb, Interpretive Action Verb, State Verb, and 10 adjectives).

They found that Descriptive Action Verbs were highly context dependent, but insensitive to who was performing the action. That is, the ratings for Descriptive Action Verbs were highly sensitive to whether the person was said to be at a seminar, a party or a business deal, but were not influenced by whether the person was an introvert, extrovert or Machiavellian. On the other hand, adjectives were not context dependent, but were highly person dependent. Interpretive Action Verbs and State Verbs showed intermediate levels of situation and person dependence. This again shows the sensitivity of verbs to contextual manipulations.

To summarise, changing the general context in which a verb appears, for example, changing the relative status of noun phrases, influences causal attributions. This has implications for the type of explanation the implicit causality bias warrants.

#### **1.3.4. Summary**

In this section, some of the factors that moderate the implicit causality bias have been presented. These include the presentation format, semantic properties of the verbs, and the wider context, which includes the gender, animacy and social status of the noun phrases.

Overall, changing the presentation format that a sentence appears in does not appear to change the implicit causality bias. Regardless of whether a sentence is presented in the active or passive voice, or the imperative or declarative form, the implicit causality bias does not change. The one factor that makes a difference is whether a sentence is phrased negatively rather than affirmatively. When a sentence has an explicit negation, then there is a slight weakening of the implicit causality bias. On the other hand, when a verb has negative semantic valence then it has a stronger bias than when a verb has a positive semantic valence.

Finally, there are differences in causal attributions when the context is changed. Male characters are viewed as more causal than female characters; typical events are given stronger causal attributions than atypical events; and congruent events are given stronger attributions than incongruent ones.

## 1.4. Overview

This chapter provided an introduction to the implicit causality bias. The implicit causality bias occurs when people are asked a causal question after even the simplest of sentences. For example, when presented with *Ted loves Mary* and asked why this is the case, people respond by answering there's something about Mary: she's nice, she's pretty, she plays golf... Basically, verbs can bias towards either the sentence subject (NP1) or the sentence object (NP2). Examples of NP1 biasing verbs are *attract*, *frighten*, *flatter* and *telephone*; while examples of NP2 biasing verbs are *love*, *fear*, *praise* and *blame*.

As well as identifying verbs by their bias, verbs can be classified according to other properties. Two different classification schemes were introduced: the Thematic Role Taxonomy and the Linguistic Category Model. Although there are some differences between the two schemes, overall they are in agreement over many of the important categories. For the remainder of this thesis, verbs will be referred to by their thematic role category, for ease of reference. Introducing thematic roles is useful because it helps to capture some properties of the implicit causality bias. Although one can identify NP1 and NP2 biasing verbs, the bias is not for a particular sentence position, but rather is a bias towards a particular thematic role. For Agent-Patient verbs it is the Agent, Agent-Evocator verbs the Evocator, while for both types of state verbs, the implicit causality bias is for the Stimulus.

A number of factors influence causal attributions other than the verb. This includes the gender, animacy and the congruency of the noun phrases with that verb. However, the basic finding is that implicit causality bias is found in many different types of task and in many different languages. Asking why *Ted loves Mary* in any manner or language generates the answer “there’s something about Mary”.

## **CHAPTER 2:**

# **EXPLANATIONS OF IMPLICIT CAUSALITY**

### **2.1. Introduction**

In this chapter, explanations of the implicit causality bias are presented: namely, implicit covariation theory, implicit salience theory and the priming hypothesis. The first theory, and the one that has received the most attention, is the implicit covariation theory. According to the implicit covariation theory, general cues, which social psychologists have shown to be used in making attributions, have become lexicalised, and this lexical information leads to the implicit causality bias. The second account is implicit salience theory, which is also based on findings from social psychology. According to this theory, most causal weight is given to the most salient argument. And finally the priming hypothesis claims that the implicit causality bias is the result of spreading activation from the verb to other related words in an associative memory. Each of these theories will be presented in turn, with empirical evidence as well as the problems associated with each account.

The common assumption underlying the explanations of the implicit causality bias is that the lexical entry of a verb carries a causal weighting, which leads to the implicit causality bias effect. The theories differ in what they assume underlies the causal weighting: covariation, salience or priming, but all agree that the implicit causality bias is a verb-based phenomenon. Let's turn to the first of these theories, the implicit covariation theory.

### **2.2. Implicit causality is implicit covariation**

#### **2.2.1. Implicit covariation theory**

##### *2.2.1.1. Causal attributions in social psychology*

Within social psychology, researchers are interested in how people make causal inferences in real life situations. However, needless to say, manipulating and measuring independent and dependent variables in real life scenarios involving *kicking*, *blame*,



*loving*, or *fascination* can prove to be very difficult. Consequently, the majority of research on causal inferences within social psychology has used verbal descriptions of scenarios. By using vignettes of this sort, researchers can study the effects of various variables on causal inferences (*e.g.* Hilton & Slugoski, 1986; McArthur, 1972).

Researchers interested in the causal inferences underlying interpersonal events work within attribution theory. However, there is not just one attribution theory. Rather attribution theory is a collection of diverse theoretical and empirical contributions that share several common concerns (Fiske & Taylor, 1991). That is, attribution theory is concerned with what information people use, how it is combined and how people make causal inferences (Heider, 1958; Jones & Davis, 1965; Kelley, 1967; 1972). Early work in attribution theory very much emphasised the role of language. So, for example, Heider (1958) argued that the best way to arrive at a social psychology theory is by examining common sense psychology and the best way to examine common sense psychology is by studying the language used by people to describe their experiences.

Because work in social psychology has relied so heavily on language in order to investigate which factors influence causal inferences, there are a number of studies examining the effects of language itself on participants' responses (*e.g.* Turnbull & Slugoski, 1988; see Hilton 1995 for a review). These studies attempt to examine the factors in language that influence judgements of cause. The research on implicit causality can be viewed as another example of this type of research.

#### 2.2.1.2. *Covariation theory*

One theory of how people make causal attributions is Kelley's (1967) covariation theory (also known as the ANOVA model of attribution). Based on J.S. Mill's observation that causes are determined by ascertaining the co-occurrence of events, Kelley proposed that determining the cause of a particular event entails finding "that condition which is present when the effect is present and absent when the effect is absent" (Kelley, 1967, 154). That is, in order to deduce what the cause of an event is, one observes the variations in the effect when the possible causes are varied. In simple situations where there are two things interacting, there are three possible causes. The causes are not exactly the same as those used when discussing the implicit causality bias. Instead, they are: (1) the Person, *i.e.* the one "who produces the response"; (2) the Entity, *i.e.* the thing "toward which the response is directed"; and (3) the Circumstances, *i.e.* the time and/or modality of interaction (McArthur, 1972, 182). In addition, the effect can be attributed to some combination of these three causes.

For example, consider an event in which Ted telephoned Mary. The three things to consider are how other people act towards Mary, how Ted acts towards other people and whether Ted telephones Mary at other times. How other people react towards Mary is a variation of *Person*. Variations of *Person* manipulate the *Consensus* variable. If Ted is one of many people to telephone Mary, then there is high consensus. However, if Ted is one of few people to telephone Mary then there is low consensus.

Another thing which may vary with the effect is how Ted acts towards other Entities. Variations of the *Entity* manipulate the *Distinctiveness* variable. Hence, if Mary is one of the few people who Ted telephones, then there is something special about Mary that is causing Ted to telephone her. Mary (the Entity) is of high distinctiveness. However, if Mary is one of many people whom Ted telephones then Mary is not crucial for the action to occur. In this case, Mary is of low distinctiveness.

It is very important to note the terminology of distinctiveness and consensus. For *distinctiveness* information, the value *high* means there is little generalisation; and the value *low* that there is much generalisation. However, for *consensus* information, the value *high* means that there is a lot of generalisation; the value *low* that there is little (Brown & Fish, 1983a). This is illustrated in Table 1.

Finally, there can be variations in the *time or modality* of the interaction. Generalisations across *time or modality*<sup>7</sup> manipulate the *Consistency* variable. It is unclear whether time and modality are to be treated as a single factor, or two different ones. Sometimes Kelley (1967) treated them as two separate variables, where he talked about consistency over time and consistency over modality (*e.g.* Kelley, 1967, 197), but at other times he treated them as one, and talked only of consistency (*e.g.* Kelley, 1967, 195). Consistency over time means that each time the Entity is present the Person's reaction is the same. That is, every time Ted thinks of Mary, he telephones her. Consistency over modality means that the Person's reaction is the same, even though the mode of interaction with the Entity is different. That is, *Mary likes the play Romeo and Juliet* is consistent if she likes it when she sees it in the theatre, when she reads the play, when she watches the film *etc.*. It is not consistent if she only likes *Romeo and Juliet* when Leonardo diCaprio stars in it. If 'A' always and in every modality responds to 'B' in this manner, then there is high consistency. On the other hand, if 'A' responds to 'B' like this only some of the time and in some modalities, then there is low consistency.

VARIABLE	VALUE	GENERALISATION	ATTRIBUTION
CONSENSUS	high	Many people X Entity	Entity
	low	Few people X Entity	Person
DISTINCTIVENESS	high	Person X few entities	Entity
	low	Person X many entities	Person
CONSISTENCY	high	Person always X Entity	Person/Entity
	low	Person rarely X Entity	Circumstance

Table 1: Above is a table of covariation values and corresponding attributions. The X in the generalisation statements is a placeholder for the type of event occurring.

These three variables can be varied orthogonal to one another, in order to assess how they affect attributions of cause. Kelley (1967) predicted, and later research by McArthur (1972) confirmed, that a combination of high distinctiveness, high consensus and high consistency leads to Entity attributions; whereas a combination of low distinctiveness, low consensus and high consistency leads to Person attributions. Note, that high consistency is required for stable attributions to both the Person and the Entity. Low consistency leads to circumstance attributions.

2.2.1.3. *Implicit covariation*

This covariation theory of attribution was used to explain the implicit causality verb bias. Brown & Fish (1983a; 1983b; Brown, 1986; Rudolph, 1997; Rudolph & Forsterling, 1997; VanKleeck, Hillger & Brown, 1988) argued that the concepts from attribution theory, specifically Kelley’s covariation account, are so fundamental to human thought, that they give rise to schemata, which are represented in the lexical entry for verbs. These schemata produce the causal weighting which manifests itself as the implicit causality bias.

Brown & Fish (1983a; 1983b) argued that there are three basic schemata: (1) an Agent-Patient schema, in which the sentence subject causes or instigates an action and the sentence object suffers a resultant change of state; (2) an Experiencer-Stimulus schema, in which the sentence subject experiences a feeling, or mental state, and the

---

<sup>7</sup> For some events, this variable is not applicable, as in the kick case. Kicking can only be achieved in one modality, with a specific movement, *i.e.* contact with foot. This may explain why Kelley sometimes includes this variable and sometimes does not.

sentence object evokes that feeling or state; and (3) the Stimulus-Experiencer schema, where the sentence subject evokes a feeling, or a mental state, and the sentence object experiences that feeling or state. These schemata are associated with classes of verbs, which are identified by their thematic roles (see section 1.2.1).

According to Brown & Fish, all action verbs have associated with them an Agent-Patient schema, which is associated with low consensus and low distinctiveness information. That is, the Agent-Patient schema states that few people act in this way, but many entities are possible recipients of this behaviour:

“the Patient classes are very large, approximately coincident with the number of humans in the world. Agent classes would be smaller” (Brown, 1986, 270).

For example, the action verb *flatter*, as in *Ted flattered Mary* is associated with the information that few other people flatter Mary, but many people other than Mary can be the recipient of Ted’s flattery. This can be represented diagrammatically as in Figure 1.




Figure 1: In each of the diagrams above, the two red dots are the characters of interest, Ted and Mary, the arrows represent “flattering” relations. The left-hand diagram shows low consensus information, where Ted is one of few people who flatters Mary; while the right hand diagram shows low distinctiveness information, where many people other than Mary can be flattered by Ted.

Table 1 in the previous section, showed that low consensus and low distinctiveness information is associated with Person attributions, that is attributions to NP1. Therefore, according to the Agent-Patient schema, all action verbs should be NP1 biasing. However, not all action verbs are NP1 biasing, which the previous chapter demonstrated. There are a set of verbs which describe situations where the sentence

subject reacts to a situation provoked by the sentence object, for example, *blame, thank, punish, etc.*. Rudolph (1997) argued that these verbs require a fourth schema, which applies to NP2 biasing action verbs. This schema is called the Agent-Evocator schema and is associated with high consensus and high distinctiveness information. For these verbs, the class of Agents is larger than the class of Patients. So, for example, given the verb *praise*, as in *Ted praised Mary* the implicit covariation information is that many other people praise Mary and that Ted praises few people other than Mary (see Figure 2). Again, from Table 1, we know that high consensus and high distinctiveness information produces attributions to the Entity, *i.e.* NP2. So, this implicit information is consistent with the bias of the verb.




Figure 2: As in Figure 1, the two red dots are Ted and Mary, and the arrows represent “praising” relations. The left-hand diagram shows high consensus information, where Ted is one of many people who praise Mary; while the right hand diagram shows high distinctiveness information, where few people other than Mary are praised by Ted.

The explanation given to implicit causality action verbs is used for state verbs, also. In the original formulation of covariation theory, Kelley intended for his model to cover not just the causes of actions but also the causes of non-observable events, *i.e.* mental events. By assuming that effects in cause-effect chains can be responses *or* sensations, where the former can be viewed as actions, and the latter states, Kelley can extend his theory to just such phenomena. These non-observable events can be referred to by mental verbs, or psych-verbs, which include verbs of perception, cognition as well as those referring to emotions.<sup>8</sup>

<sup>8</sup> There are some problems with this assumption. For example, White (1992) argues that covariation theory requires that two things are regularly observed together, the cause is perceived and it co-occurs

Let us assume then that it is possible to “plug” non-observable things into covariation theory. Still the mapping of covariation roles to thematic roles is not a trivial process. To recap, in Kelley’s theory of attributions there is a Person who produces a response and an Entity towards whom the response is directed. For action verbs, the mapping between Agents and Patients to Persons and Entities is quite straightforward. It is clear that the Agent is the Person because by definition the Agent is the one that produces the response, just as the Patient is the Entity because that is to whom the response is directed. So, in *Ted kissed Mary*, Ted is the Agent and the Person, while Mary is the Patient and the Entity.

On the other hand, for state verbs the mapping is not quite so clear. This is because for state verbs it is not obvious that there is a ‘response’. Take the sentence *Ted loves Mary*. Is Ted the Person because he is the one who is experiencing the loving? Is that the response? Or, is Mary the Person because she is bringing about the state of loving?

This mapping of roles from social psychology onto linguistic roles is not merely a problem associated with Kelley’s covariation theory, but a more general problem with the terminology adopted by attribution theorists. So, for example, Jones & Nisbett (1972) have demonstrated a general bias in attribution, called the Actor-Observer bias. An Actor is someone who performs an action; an observer can be an Active Observer, *i.e.* the recipient of that action, or a Passive Observer, who observes the interaction between the Actor and the Active Observer. Jones & Nisbett found that when an Actor was asked to explain the cause of his behaviour he attributed it to the situation, yet, when an observer was asked to explain the same behaviour, he attributed it to some disposition in the actor. For example, for the event Ted helped Mary, Ted is the Actor because he is the person helping and Mary is the Active Observer, as she is the recipient of the act of helping. If Ted is asked to explain why he helped Mary then his explanations will be things like “*Mary needed the help, and I was available at the time*”, “*There was no-one else who could help*” *etc.*. That is, he will refer to the situation, which things external to him. However, if Mary is asked to explain why Ted helped her, then likely explanations are: “*Ted is a helpful person*” or “*Ted is a*

---

with the effect. However, in the case of states the effect is not observable, and in fact in many cases the cause is not observable either. For example, in cases such as *Ted likes Mary*, a plausible cause may be that Mary is nice. But, neither the effect, *i.e.* the liking, nor the cause, *i.e.* Mary being nice, is observable (although their effects might be, *i.e.* doing things that show that you like someone and doing things that show that you are nice).

*generous person*” etc.. Mary is more likely to refer to attributes of Ted. That is the Actor explains the event as being due to the situation, and the observer explains it as the result of a disposition of the actor.

In order to apply this theory to events depicted by state verbs, we face the same problem as encountered with Kelley’s theory. The Actor in Jones & Nisbett’s terminology is equivalent to the Person in Kelley’s theory; similarly, the Active Observer is equivalent to the Entity. So, exactly the same problems associated with mapping Kelley’s Person/Entity to the Stimulus/Experiencer, apply to mapping Jones & Kelley’s actor/Active Observer to the Stimulus/Experiencer.

To take a final example, Miller & Ross (1975) reported a Self-Serving bias in attribution, where people take credit for success, but deny responsibility for failure. In attribution terms, this means that when an Actor does well at a task, it is because of some internal attribute of his; but if he does badly at a task, then it is because of the environment. So, for example, if I pass a test it is because I am a very clever person; however, if I fail a test it is because the exam was very hard, the person behind me in the exam kept coughing and it distracted me so much I couldn’t concentrate etc.. Again, in order to generalise the Self-Serving bias to more complex situations, such as states, the Actor needs to be identified.

Three different mappings between Person/Entity and Stimulus/Experiencer have been proposed by researchers. Each of these will be discussed in turn:

(A) The Experiencer is the Person

Some theorists have mapped the Experiencer of state verbs to the Agent of action verbs, therefore making the Experiencer the Person and the Stimulus the Entity (e.g. Brown & Fish, 1983; Van Kleeck, Hillger & Brown, 1988).

Van Kleeck, Hillger & Brown (1988) begin with a distinction between “internal” and “external” locus of cause (Heider, 1958). Internal causes are due to the individual, and external causes are due to the environment. The environment can of course include other people. Van Kleeck *et al.* argue that the inner locus of an experience or an action designates the Person, in Kelley’s sense. So, a Person is the one having an experience or performing an action. In Figure 3, the grey character is the Person. In the left-hand box, an action verb *flatter* is depicted. The Person (Agent) is flattering by producing an observable behaviour, *i.e.* by saying, “You’re lovely”. In the right-hand box, is the state *love*. In this case, there is no observable action, just the Person (Experiencer) feeling a particular way.



The external locus is any other relevant aspect of the situation. In simple situations, such as those depicted in the box below, the external locus of cause is the recipient of the act (the Patient) or the feeling (the Stimulus).

Hence, according to the locus account, the Experiencer is the Person, and the Stimulus is the Entity. This mapping is also advocated by Cunningham, Starr & Kanouse (1979), with respect to the roles of Actor and Active Observer. They hold that the Experiencer is the Actor and the Stimulus is the Active Observer.

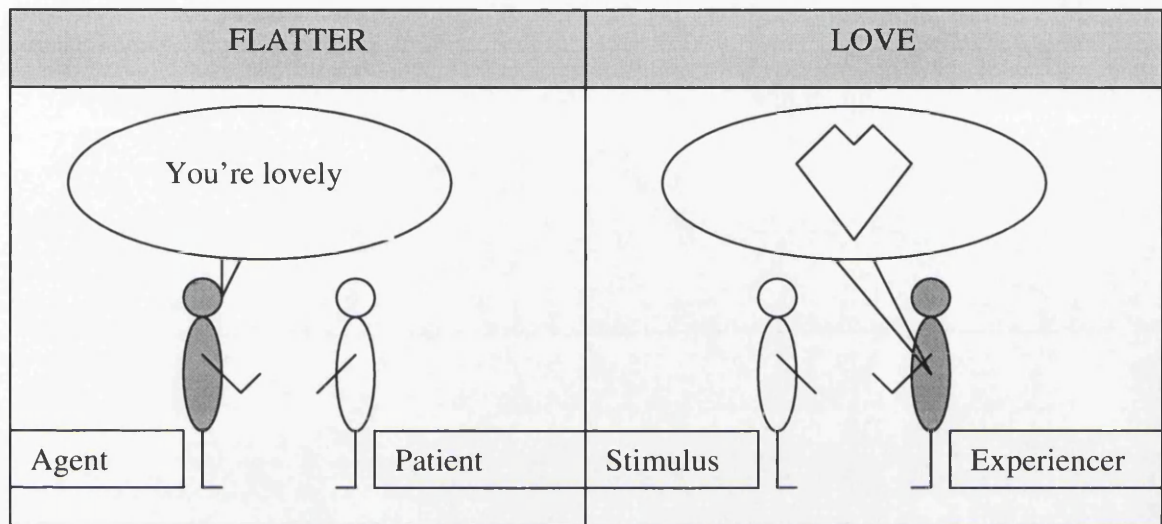


Figure 3: In the left-hand box is the event denoted by the Agent-Patient verb *flatter*. The grey character is the Person as he is the internal locus of the action. The right-hand box is the state denoted by the Experiencer-Stimulus verb *love*. Again, the grey character is the Person as he is the internal locus of the experience.

(B) The Stimulus is the Person

However, others have mapped the Stimulus to the Person, and the Experiencer to the Entity (e.g. Kasof & Lee, 1993; Gilovitch & Regan, 1986; Johnson-Laird & Oatley, 1988; see Figure 4). According to this mapping, overt behaviour is the key to identifying roles. For example, the event described by the verb *flatter* begins with the action of actually flattering, e.g. Ted telling Mary that she is lovely. Hence, the Agent is the Person, and the Patient is the Entity – exactly as was the case in mapping (A).

The mapping differs for state verbs. In Johnson-Laird & Oatley’s (1988) explanation of folk theories of emotion, emotions are also the result of an action. An action results in an emotion, which further produces an action, which then elicits another emotion, *ad infinitum*. So, if Ted loves Mary then this is because Mary has produced an action, e.g. Mary plays golf, and it is *this* action which led to Ted’s emotion.



Kasof & Lee (1993) also advocate mapping the Stimulus to the Person. When examining the Actor~Observer bias, Kasof & Lee map the Stimulus to the Actor, and the Experiencer to the Active Observer, because the Stimulus has produced an action, or a series of actions giving rise to the consequent emotion. Further, Gilovitch & Regan (1986) argue that the Stimulus has greater volitional control than the Experiencer does. That is, they also take the overt production, or control, of behaviour as characteristic of the Actor (Person). Hence, those that endorse mapping (B) argue that states are the result of a preceding action. This can be seen in the right-hand side of Figure 4. The state of love is produced by an action of the Stimulus: in this case playing golf. Therefore, the Stimulus is the Person and the Experiencer is the Entity.




Figure 4: In the left-hand box is the event denoted by the Agent-Patient verb *flatter*. Once again, the grey character is the Person, but this time because he produced the action of flattery. The right-hand box is the state denoted by the Experiencer-Stimulus verb *love*. Again, the grey character is the Person – but note, this time the Person is the Stimulus. The Stimulus produced an action (playing golf) which caused the Experiencer to feel love.

(C) The sentence subject is the Person

Yet another mapping is adopted by Rudolph (1997) and is also implicit in the work of Kelley (1967) and McArthur (1972). According to this mapping, the social roles are mapped onto sentence structure. The Person role is mapped to the sentence subject and the Entity role to the sentence object position. This means that while the Stimulus is the Person for Stimulus-Experiencer verbs, the Experiencer is the Person in Experiencer-Stimulus verbs. Conversely, the Experiencer is the Entity for Stimulus-Experiencer verbs and the Stimulus the Entity for Experiencer-Stimulus verbs.

To summarise, three different mappings between thematic roles and social psychology roles have been described. In order to appreciate the differences between these mappings, Table 2 provides a succinct overview.

VERB	THEMATIC ROLES	COVARIATION ROLES			ACTOR~OBSERVER ROLES		
		(A)	(B)	(C)	(A)	(B)	(C)
ES	Stimulus	Entity	Person	Entity	Observer	Actor	Observer
	Experiencer	Person	Entity	Person	Actor	Observer	Actor
SE	Stimulus	Entity	Person	Person	Observer	Actor	Actor
	Experiencer	Person	Entity	Entity	Actor	Observer	Observer

Table 2: The mapping of thematic roles to social psychology roles. The table provides three different mappings of how the thematic roles for state verbs relate to roles from covariation theory, as well as the Actor~Observer bias.

The issue of whether the Experiencer, the Stimulus, or the sentence subject is most “Agent-like” has been debated both by linguists (*e.g.* Dowty, 1991) and psycholinguists (*e.g.* Schlesinger, 1992). So, this plethora of mappings proposed by social psychologists are not due to the idiosyncrasies of implicit causality verbs; but is rather a more general problem. However, the assignment of roles does have additional implications for implicit covariation theory. The implicit covariation theory states that the implicit causality bias is the result of verbs having underlying covariation information. This information is based on generalisation of the Person and the Entity. However, if it is not clear who the Person and the Entity are for state verbs then there is a problem of what are the underlying generalisations.

In practice, researchers have agreed on the underlying generalisations, but changed the nomenclature associated with the information, depending on which of the above mappings they choose. Whereas Brown & Fish (1983a; 1983b) argue that both Stimulus-Experiencer and Experiencer-Stimulus verbs have underlying high consensus and high distinctiveness information; Rudolph (1997) argues that Experiencer-Stimulus verbs have high consensus and high distinctiveness information, but that Stimulus-Experiencer verbs have low consensus and low distinctiveness information. In both cases it is assumed that Experiencer-Stimulus verbs have the implicit information ‘*Many people verb X*’ and ‘*X verbs few people*’; whereas Stimulus-Experiencer verbs

have the information that, ‘*Few people verb X*’ and ‘*X verbs many people*’. This is illustrated in Figures 5 and 6.




Figure 5: Information associated with Stimulus-Experiencer verbs. If the arrows represented “fascinating” relations, then few people fascinate Ted, while Mary fascinates many people.




Figure 6: Information associated with Experiencer-Stimulus verbs. If the arrows represented “loving” relations, then many people love Mary, while Ted loves few people.

To summarise, different mappings associate different generalisations with different labels. Although, researchers have changed the labels they are using so that they are referring to the same generalisations, the different mappings have caused confusions with implicit salience theory, which will be discussed in later (section 2.2.2). Therefore, one mapping has to be chosen, for the work that will be presented in this thesis, in order to avoid confusion, and problems in interpretation.

Mapping (A) will be used. This is where the Experiencer is the Person and the Stimulus is the Entity. Brown & Fish (1983b) and Van Kleeck, Hillger & Brown (1988)

have a very convincing argument in favour of mapping (A). First of all, they argue against mapping (C), where the sentence subject is the Person. The basis of their argument is that attribution theory should be relevant to real-life situations and not just sentences. If mapping (C) is chosen then covariation generalisations are made in terms of grammatical relations and not real-life situations.

Van Kleeck *et al.* argue that the same real-life situation can be expressed in different ways. For example, the event of Mary laughing at a comedian can be expressed as *Mary laughed at the comedian*, or *The comedian was laughed at by Mary*, or even *The comedian amused Mary*. Under mapping (C), high consensus information for each of these utterances would be different. These would be *Many people laughed at the comedian*, *Many comedians were laughed at by Mary*, and *Many comedians amused Mary* respectively (see Van Kleeck *et al.*, 1988, 96). Based on covariation principles, whereas the comedian would be the cause in the first case, Mary would be the cause in the remaining cases. This is despite the fact that the same situation is depicted by the three target sentences. Therefore, mapping (C) is not satisfactory.

This leaves a choice between mapping (A) and mapping (B). Both of these would allow paraphrastic relations to be maintained. However, there is a serious problem with mapping (B). Proponents of mapping (B) argue that the Stimulus produces an action and therefore is the closest to the Agent. But it is quite conceivable that the Experiencer feels an emotion even if the Stimulus has not produced an action. Consider the Experiencer-Stimulus verb *love*. The Experiencer can feel love because of some property of the Stimulus; for example, *Ted loved Mary because she invited him to the high school prom*. More than this, the Experiencer can have an emotion that does not have a reason, for example, *Ted loved Mary for no reason at all*; or *Ted loved Mary, even though she did nothing to deserve it*. Finally, an action need not be involved as is demonstrated by the acceptability of using state words with inanimate things, *e.g. Ted loved the painting*.

The argument that emotions require a preceding action is actually based on a subset of state verbs, namely Stimulus-Experiencer verbs. Semin & Fiedler (1992) argued that Stimulus-Experiencer verbs (State Action Verbs) require a preceding action (see section 1.2.2). So, the cause of *Mary fascinated Ted* is some preceding action of Mary, for example, the fact that Mary can get a hole in one, when playing golf. But, once again, it is possible that Ted feels fascination, without Mary having produced an action to warrant this. The following sentences, where there is no preceding causal action, are not anomalous: *Mary fascinated Ted because she was nice*; *Mary fascinated*

*Ted for no reason; and Mary fascinated Ted even though he did not do anything to warrant it, and finally, The painting fascinated Ted because it was beautiful.*

To summarise, state verbs do not require a preceding action for the state to occur. Therefore, mapping (A), which is based on locus of cause, is the preferred mapping: the Experiencer is the Person and the Stimulus is the Entity. Table 3 presents the consensus and distinctiveness information associated with the two types of state verbs, according to implicit covariation theory.

VALUE	ES VERBS ( <i>e.g.</i> like)	SE VERBS ( <i>e.g.</i> fascinate)
HC	Many people like X	X fascinates many people
LC	Few people like X	X fascinates few people
HD	X likes few people	Few people fascinate X
LD	X likes many people	Many people fascinate X

Table 3: Consensus and distinctiveness information for state verbs (ES = Experiencer-Stimulus; SE = Stimulus-Experiencer)

To conclude this section, implicit covariation theory states that the implicit causality bias is the result of verbs carrying “causal weightings” as part of their meaning. It is argued that access to a verb entails access to:

“a lexical entry in memory, which includes information of syntactic class (verb); subclass (action or state); associated semantic roles or verb arguments (Agent-Patient or Stimulus-Experiencer); grammatical subject and object assignment of the roles; and relative causal weights; with redundancies eliminated” (Brown & Fish, 1983a, 244).

The “causal weights” are the result of generalisations based on consensus and distinctiveness and are derived from real world observations based on classes of behaviour. Verbs that bias towards NP1, Agent-Patient and Stimulus-Experiencer verbs carry with them the information that the subject class is much smaller than the object class. On the other hand, verbs that bias towards NP2, Agent-Evocator and Experiencer-Stimulus verbs, are associated with the information that the object class is much smaller than the subject class (see Figure 7). Covariation theory predicts that out of two entities in different sized sets, the entity in the smaller set is the cause of the event. There’s something about Mary because she belongs to a smaller set than Ted. This difference in set sizes is the basis of the implicit causality bias.

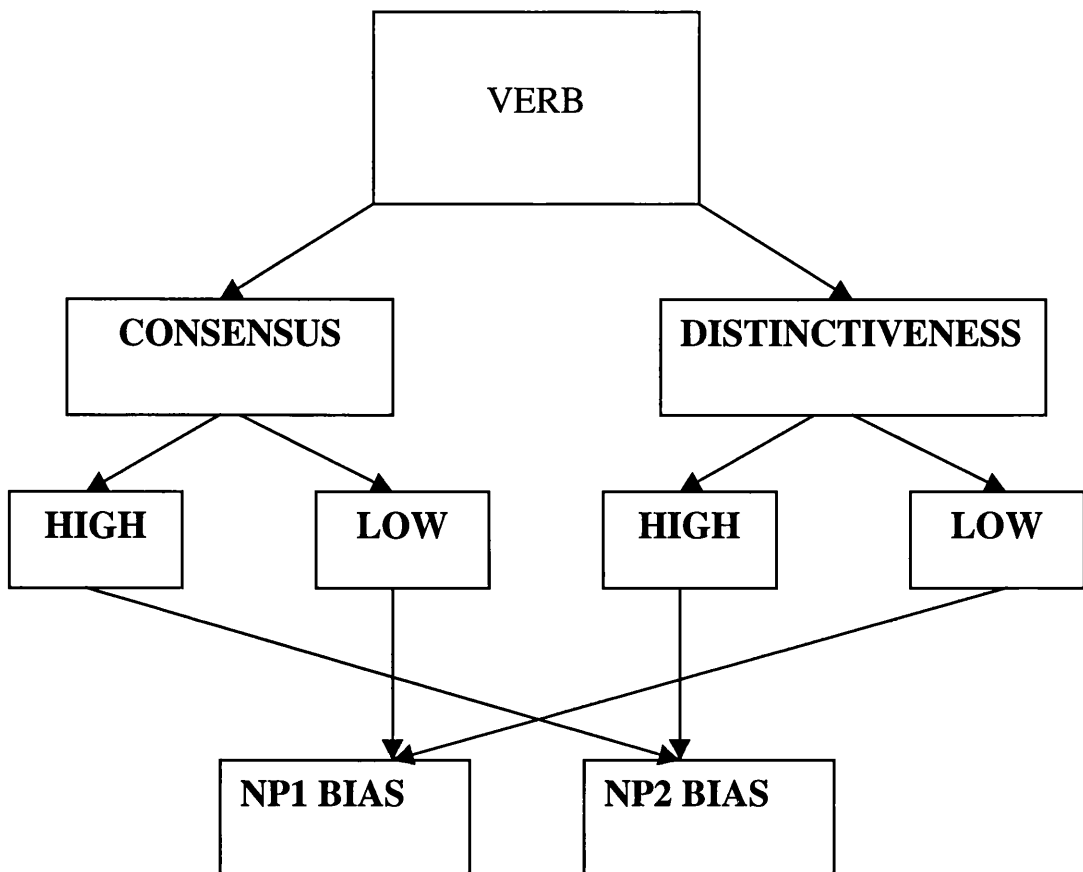


Figure 7: Implicit covariation information predicts that consensus and distinctiveness information is lexicalised with the entry for the verb, and this information gives rise to the implicit causality bias.

#### 2.2.1.4. *Tests of the implicit covariation theory*

In order to test whether verbs do indeed carry implicit covariation information, participants are presented with simple SVO sentences and asked to judge the generalisation of the verb both across people and across entities (Brown & Fish, 1983a, Hoffman & Tchir, 1990; Rudolph, 1997). Participants are presented with the target sentence and then are asked to rate how many other people the sentence subject would react to in that manner, and how many other people would respond in that way to the sentence object.<sup>9</sup> For example, for the sentence *Ted flatters Mary* participants would be presented with:

---

<sup>9</sup> This is equivalent to the tasks of Abelson & Kanouse (1966; Kanouse, 1972), where they ask people to make generalisations based on unquantified sentences. The covariation generalisations are an extension of the unquantified sentence generalisations.

Ted flattered Mary.

A. Probably [few 1 2 3 4 5 6 7 8 9 many] other people flattered Mary.

B. Probably Ted flattered [few 1 2 3 4 5 6 7 8 9 many] other people.

When participants are asked to make generalisations using this procedure, subject biasing verbs are associated with the information that a small number of people produce the action, but that a large number of objects are possible recipients; whereas object-biasing verbs are associated with the opposite information: that is, a large number of people produce the action, but there are small number of possible recipients. That is, NP1 biasing verbs have low consensus and low distinctiveness associated with them; while NP2 biasing verbs have high consensus and high distinctiveness associated with them.

Further, there is a significant correlation (with values of  $r$  around .60) between the implicit causality bias of the verb and implicit covariation information. For example, Hoffman & Tchir (1990) conducted two studies, one to obtain implicit causality judgements and a second for implicit covariation judgements. They found a strong correlation between attributions in the first study and covariation generalisations in the second study ( $r = .73$ ). When ratings for covariation were then introduced as a covariate to Experiment 1 the effect of implicit causality was reduced (although not eliminated).

Based on such data, Rudolph & Forsterling (1997) in a comprehensive review of the implicit causality literature argued strongly in favour of the covariation account. However, although these results are consistent with the hypothesis that implicit covariation mediates the implicit causality bias, they do not prove that this is the case. The results from the above experiments are only correlational. It could be that implicit covariation and implicit causality are mediated by some third variable. Or even that participants utilise the implicit causality bias to make generalisations about covariation information. That is, the bias gives rise to generalisation inferences, rather than covariation information giving rise to the bias. Given that the only evidence in favour of the implicit covariation theory is based on the generalisation inference task above, either of these alternative explanations of the data is as plausible as the original interpretation of the data.

There is a second problem with the implicit covariation theory, and this is the fact that it is a lexically based explanation. The implicit covariation information is



stored at the lexical level (see Figure 7). For the results from the generalisation inference task to be convincing evidence for implicit covariation theory, it has to be demonstrated that the task taps into the lexical representation of the verb alone, and is not the result of other processes. This does not appear to be the case. Generalisation inferences are affected by the type of noun phrase the verb takes as an argument.

Most of the research on the implicit causality bias has focussed on the verb alone, but as was seen in the previous chapter (sections 1.3.3.1; 1.3.4 and 1.3.5) changing the noun phrases in a sentence can change the implicit causality verb bias (Corrigan, 1988; 1992; Garvey, Caramazza & Yates, 1976; Turnbull & Slugoski, 1988). This finding poses a serious challenge to implicit covariation theory. According to this theory, covariation information is stored in the mental lexicon, specifically in the lexical entry for the verb. This being the case, explaining how presenting a role description changes the implicit causality bias is not trivial.

Not only do the causal attributions change when a role description is introduced, but so do the generalisation inferences participants make. Corrigan (1988) found that for some verbs, changing the animacy of the noun phrase shifted the causal bias, as well as the generalisations people made. For example, although participants assign causality to NP1 when presented with *Ted chased Mary*, they made much weaker attributions when given *Ted chased the bus*. As with the causal ratings, participants also changed the generalisations they made. So, whereas for *Ted chased Mary* people made low consensus and low distinctiveness generalisations, for *Ted chased the bus* they made high consensus and high distinctiveness generalisations.

Because generalisation patterns change when the same verb is presented with different arguments, it is unlikely that generalisation inferences are tapping into the lexical entry for the verb alone. Since, the generalisation data is the only evidence for implicit covariation theory, proponents of the theory need to maintain this evidence. They can do this in one of two ways.

First, it could be argued that noun phrases affect causal attributions and generalisation patterns because they help select a sense of the verb. If this is the case then the lexical explanation of implicit covariation theory could be retained. For example, the verb *attacked* means something slightly different in *The soldiers attacked the castle* and *Ted attacked Pat*. In the first example, it is the literal sense of attack that is being used; while in the second example it is the figurative. Hence, the verb *attack* has slightly different senses in the two examples. Therefore, when different causal attributions and generalisations are obtained, it could be because different senses of a



word are being activated, and the different senses have different biases. Although this may be true, this explanation cannot explain all the cases where the causal bias changes when it's argument differs. So, for example, *chase* in the two examples above does not have different senses, but the causal bias and the generalisations differ just the same.

A second explanation a proponent of the implicit covariation theory could give is that when a verb is presented with role-descriptions, further processing occurs. That is, when role-description information is introduced this causes modification of the causal attribution at a later stage. The verb bias is activated when the verb is activated, it is just over-ridden by the contextual information after that. Just as the causal attributions are over-ridden at a later stage, the generalisations produced by participants are over-ridden at a later stage, also. The participants would then be using the information from the noun phrases to infer appropriate generalisations. However, if this is indeed the case, and participants infer generalisation patterns when noun-phrases are not just proper names, it is very hard to maintain that generalisations are not inferred even when the noun phrases are proper names.

It is more plausible to assume that generalisation inferences are constructed from the sentence, rather than stored for each verb. For example, when presented with *The father praised the son*, a participant *computes* how many *fathers* praise and how many *sons* are praised. This is the account that Corrigan (1988) endorses. It is important to note that once the content of the noun phrases is taken into consideration the covariation account is no longer a lexical explanation, but a scenario based one.

The covariation account can easily include content information from the noun phrase because covariation theory says nothing about how the comparison groups are identified (Semin, 1980). In fact, in the above example, the relevant target group to identify consensus information could be a more general group than just “fathers” such as “parents”, “adults”, “human beings” *ad infinitum*. Or it could even be a more specific group than “fathers”. For example, it could be fathers from a particular location such as “fathers from Britain”, “fathers from Scotland”, or “fathers from Glasgow” *etc.*.

Although it could be argued that covariation information is stored for every possible scenario, this appears very unlikely given that there is an abundance of comparison groups that can be used, and infinitely many scenarios if complex noun phrases are used as arguments. This is obvious from the preceding discussion. Because consensus and distinctiveness information would vary as a function of comparison group, scenario level covariation information cannot be stored for every possibility.

Therefore, it appears that norm information associated with scenarios is constructed as needed, rather than stored in memory (Miller & Kahneman, 1986).

To summarise, proponents of the implicit covariation theory cite patterns of generalisation inferences as evidence for the theory. In order to show that the implicit causality bias is the result of covariation information being stored in the lexicon generalisation inferences are cited. However, this evidence is flawed in a number of ways. First, it was argued that the evidence is only correlational, and therefore does not prove that the generalisation inferences give rise to the causal bias. And second, it is more plausible that generalisation inferences are constructed based on scenario level information. Therefore, the implicit covariation theory is not supported.

## **2.3. Implicit causality is implicit salience**

### **2.3.1. Implicit salience theory**

In the previous section, a specific attribution theory, namely covariation theory, was applied to the implicit causality bias. Yet another determinant of causal attributions is salience. Salience is one of the basic cues<sup>10</sup> used in making attributions of cause, alongside cues such as temporal and spatial contiguity (Einhorn & Hogarth, 1986). Attributions of causality are made more strongly to the most salient stimuli in the perceptual field.

Salience is defined as “features that catch attention” (Fiske & Taylor, 1991). Objects can have salient perceptual features based on Gestalt properties, such as brightness, complexity, change, movement, *etc.* (McArthur & Post, 1977). People can have salient characteristics in their immediate context (*i.e.* to others who are present) by being of a different race, sex, hair colour; by wearing bright things; or by moving. Further, people can be salient by behaving in an unusual manner, by trying to achieve a goal relevant to them, as well as by dominating the visual field (Taylor & Fiske, 1978). Overall, various factors can make a thing more or less salient.

Kasof & Lee (1993) explain the implicit causality bias through salience. They argue that the implicit causality bias is the result of the arguments of a verb differing in salience, which they define as “relative degree of attention” (Kasof & Lee, 1993, 878,

---

<sup>10</sup> It is basic in the sense that it appears at an early stage in development.

see Figure 8). Because people attribute greater causality to the more salient stimuli, more attributions are made to one argument than to the other.

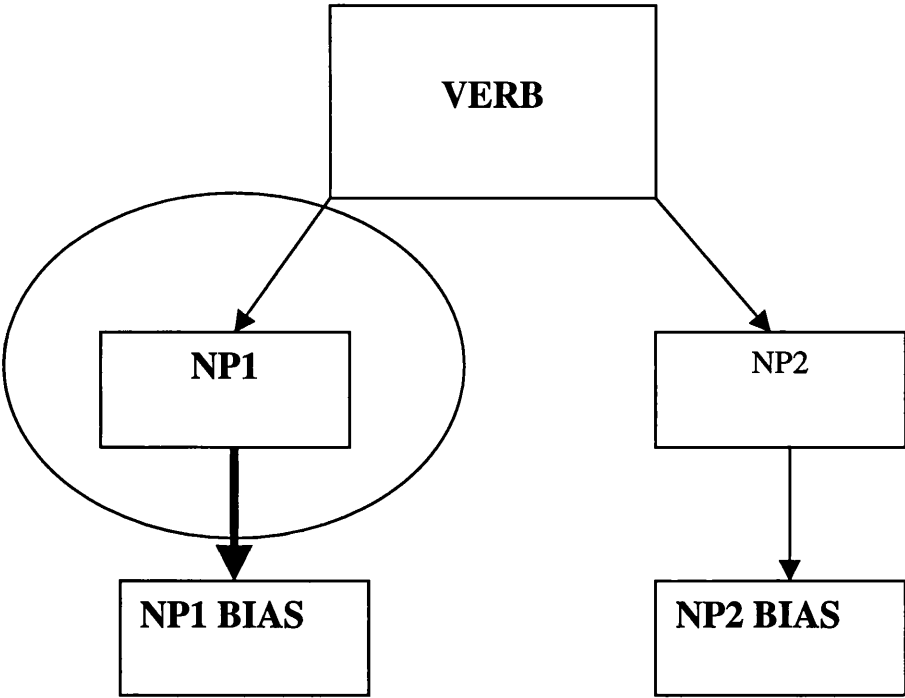


Figure 8: Implicit salience theory states that in the mental representation of a verb, one argument of the verb is more salient than the other. This is represented by the circle of attention over the NP1 argument. This salience leads to more weight being given to that argument, resulting in the implicit causality bias.

Salience of an argument is determined by the some of the things described above, such as change, complexity and movement. Kasof & Lee emphasise action as a cue to causality, as performing an action includes many of the cues used in determining salience. For Agent-Patient verbs, the Agent is more salient because the Agent is the actor, whereas the Patient is just part of the situation. Similarly, for state verbs, Kasof & Lee argue that the Stimulus is more salient because the Stimulus performs actions, whereas the Experiencer just perceives the actions. Hence, according to this account, the arguments of a verb carry a causal weight, due to one argument being more salient: implicit causality verbs bias towards the Agent and the Stimulus. In fact, Kasof & Lee show a strong correlation between ratings of salience and attributions of causality.

### 2.3.2. Tests of the implicit salience theory

Based on implicit salience theory, Kasof & Lee make a number of predictions. Each of these predictions and the relevant data are discussed in turn. To begin with,

Kasof & Lee argue that because verb arguments are more salient in the surface subject position than the surface object position (Crawley & Stevenson, 1990; Gernsbacher & Hargreaves, 1988), the Agent in Agent-Patient verbs and the Stimulus in Stimulus-Experiencer verbs should be more causal in the active voice than in the passive voice. Further, the Stimulus should be more causal in the passive voice than in the active voice for Experiencer-Stimulus verbs. Although Kasof & Lee (1993, Experiment 1) found strong evidence in favour of this hypothesis, in a meta-analysis conducted by Rudolph & Forsterling (1997) this was only partially supported (see section 1.3.1.1). Consistent with Kasof & Lee, Rudolph & Forsterling found the strength of the bias associated with Experiencer-Stimulus verbs did increase when it was presented in the passive voice instead of the active voice. However, they found no difference for Agent-Patient verbs, and the opposite pattern for Stimulus-Experiencer verbs, where the bias actually increased when presented in the passive voice, instead of decreasing as predicted. Therefore, Kasof & Lee's predictions about the effects of passivisation are only partially supported.

Secondly, Kasof & Lee predict a difference in causal attributions between verb classes, namely between Stimulus-Experiencer verbs and Experiencer-Stimulus verbs. They argue that the Stimulus is more salient in the former class of verbs because Stimulus-Experiencer verbs refer to events with specific, discrete actions involving two people close in space and time (see The Linguistic Category Model, section 1.2.2); whereas Experiencer-Stimulus verbs refer to events that are diffuse. Indeed, Kasof & Lee found that the Stimulus is more causal for Stimulus-Experiencer verbs (Kasof & Lee, 1993, Experiments 1 and 2; see also, Brown & Fish, 1983a; Brown & van Kleeck, 1989; Corrigan, 1988, Experiment 1; Corrigan, 1992; Hoffman & Tchir, 1990; Semin & Fiedler, 1992).

However, the interpretation of this cannot be the one that Kasof & Lee advocate, as it is unclear how Ted and Mary are in closer temporally and spatially in *Ted fascinates Mary* (Stimulus-Experiencer verb) than in *Ted loves Mary* (Experiencer-Stimulus verb). Furthermore, note that the arguments Kasof & Lee use to show the Stimulus is more salient actually entail that participants are using the cues of temporal and spatial contiguity, not necessarily salience. Temporal and spatial contiguity are a further two cues that people use when making judgements of causality (Einhorn & Hogarth, 1986; Kummer, 1995). When making certain types of causal judgements, people are more likely to attribute an effect to a cause that is close in time and space to the effect. Michotte's (1946/1963) research is a very good example of this (see section

0.1). Box A moves in a straight line until it reaches Box B, stops moving, and then is followed by Box B moving along the same path. Box A is perceived as causing B's motion because the cause (Box B moving) was in close temporal and spatial contiguity to the effect (Box A moving). Similarly, Kasof & Lee argue that the Stimulus is more causal in Stimulus-Experiencer verbs than in Experiencer-Stimulus verbs, because in the former case there is closer temporal and spatial contiguity between the two characters. This argument does not rely on a difference in salience at all. Therefore, the observed difference in attributions between these two classes of verbs could be due to salience or contiguity.

Thirdly, Kasof & Lee (1993, Experiment 3) predict that by manipulating the salience of arguments, the implicit causality bias should be modified. They examined causal attributions to inanimate objects by presenting sentences that used different types of collision verbs, *e.g. bang into, smash into, tap, etc.*. Participants were asked to do two tasks: (1) estimate the speed of the cars involved; (2) make causal attributions about which car caused the accident. Kasof & Lee argue that in sentences such as *Car A bangs into Car B*, Car A (the Agent car) should be more salient. Further, because faster moving entities are more salient (McArthur & Post, 1977), there should be a strong correlation between estimates of speed and attributions of causality. This was indeed the case. They found that Agent cars were given higher ratings of causality than Patient cars. Moreover, there was a strong correlation between the speed of the car and the strength of the causal attribution, as predicted. The relationship between speed and causality is a nice demonstration of the effect of salience on the implicit causality bias.

However, the above evidence has all been of a correlational nature. As with the implicit covariation theory, salience could be causing the implicit causality bias, or the implicit causality bias could be driving judgements of causality, or some third variable may be the cause.

Kasof & Lee (1993, Experiment 2) did attempt to test their theory directly. This was done by using the Actor-Observer bias, which was introduced earlier (section 2.2.1.1.3). To recap, in a situation with an actor, an Active Observer (the recipient of the action) and a Passive Observer (observer of the event), both observers attribute causality to the actor; whereas the Actor attributes causality to the Active Observer. One explanation for this finding is that the Actor is the most salient thing in the environment from the perspective of the observers; whereas the Active Observer is the most salient thing for the actor. The Actor is most salient for the observers because the Actor is goal-driven. The Active Observer is most salient from the actors perspective

because he or she is goal relevant to the actor. Kasof & Lee argue that if the Actor perspective is introduced into experimental items then the causal bias should be reduced or even eliminated.

Under the usual circumstances, the Agent and the Stimulus are seen as causal. According to Kasof & Lee, this is because the Agent and the Stimulus are the actors, the Patient and the Experiencer the Active Observers, and the reader is the Passive Observer. However, if the reader is given a sentence where the self is the Agent/Stimulus then attributions ought to change. Since the Actor attributes causality to the Active Observer, sentences where the self is the Agent/Stimulus should lead to more attributions to the Patient/Experiencer than where the Agent/Stimulus is an observer. That is, one can manipulate the salience of the arguments by changing the perspective of the reader. Kasof & Lee tested causal attributions for three different perspectives (see Table 4).

Kasof & Lee’s hypotheses were confirmed, in that fewer causal attributions were made to the Agent/Stimulus than the Patient/Experiencer when the Actor perspective was introduced, than when the observer perspectives were used. That is, there appears to be a reversal of the implicit causality bias effect when a change in perspective is introduced.

View	ES VERBS	BIAS	View	SE VERBS	BIAS
Actor	Ted admires me	E (subject)	actor	I fascinate Ted	E (object)
AO	I admire Mary	S (object)	AO	Mary fascinates me	S (subject)
PO	Ted admires Mary	S (object)	PO	Mary fascinates Ted	S (subject)

Table 4: Experimental items in three perspectives (*i.e.* actor, AO-Active Observer, and PO-Passive Observer) and the predicted causal attribution (E = Experiencer; S = Stimulus).

The results from this study appear to support the implicit salience theory. Further, implicit salience theory makes predictions that are not be made by implicit covariation theory.<sup>11</sup> However, there is a major problem with the above study. In order

---

<sup>11</sup> Rudolph (1997) argues that the implicit salience theory can be reduced to the implicit covariation theory, and therefore salience can be reduced to covariation. However, the results from Kasof & Lee’s

to derive predictions in accordance with implicit salience theory, Kasof & Lee have to map the role of the Stimulus to that of the Actor. However, in section 2.2.1.1.3, it was argued that it is more appropriate to map the *Experiencer* to the Actor. In fact, Kasof & Lee present a re-analysis of a study by Cunningham, Starr & Kanouse (1979) who had taken the Experiencer as the Actor. By changing the mapping from Stimulus to Experiencer, the predictions about the causal attributions are reversed. This can be seen in Table 5.

View	ES VERBS	BIAS	View	SE VERBS	BIAS
Actor	I admire Mary	E (subject)	actor	Mary fascinates me	E (object)
AO	Ted admires me	S (object)	AO	I fascinate Ted	S (subject)
PO	Ted admires Mary	S (object)	PO	Mary fascinates Ted	S (subject)

Table 5: Experimental items when the Experiencer is mapped onto the role of Actor (*i.e.* actor, AO-Active Observer, and PO-Passive Observer) and the predicted causal attribution (E = Experiencer; S = Stimulus).

For example, according to Kasof & Lee the sentence *Dave admires me* makes the Stimulus take on the Actor perspective, but according to the second mapping, the same sentence actually makes the Stimulus take the Active Observer perspective. So while, in accordance with the implicit salience theory, Kasof & Lee show a decrease in attributions to the Stimulus when the Actor perspective is introduced; what they actually show is a decrease in Stimulus attributions when the Active Observer perspective is introduced.

Earlier in this chapter, it was argued that the majority of the evidence supports a mapping between the Experiencer and the Actor. Although Kasof & Lee defend the Stimulus-Actor mapping on the grounds that the Stimulus produces an action, this is not necessarily the case as is shown by the acceptability of statements such as *Ted loves the painting*. On the other hand, there are a number of positive arguments that support the Experiencer-Actor mapping. If the Experiencer is treated as the Actor then the above

---

perspective experiment show that this is not true. The implicit covariation theory states that covariation is generalised from experience, and this is stored with the verb. It is content blind. Therefore, it cannot account for differences in causal judgements based on perspective change.

results do not support the implicit salience theory. Instead, what the results show is an increase in causal attributions to the Stimulus when the Actor perspective is taken. This is not compatible with implicit salience theory. It is not obvious what these results do demonstrate.

Further, the interpretation is even more unclear when the pronouns used by Kasof & Lee are considered. Under normal circumstances, the pronouns *I* and *me* refer to the writer. However, Kasof & Lee assume that participants are accepting the sentence as self-referential. This is not at all obvious. Therefore, the results of this experiment cannot be taken as support for implicit covariation theory.

Furthermore, Brown & van Kleeck (1989) and Gilovich & Regan (1986) also report data where the Actors perspective is included in sentences with implicit causality verbs. However, neither of these studies found that change of perspective influenced causal attributions. Nonetheless, a caveat is required here, in that Brown & van Kleeck and Gilovich & Regan used real-life scenarios where participants were asked to describe situations that had occurred to them. As Brown & Van Kleeck stated, the causal attributions in their studies may have more to do with the processes of encoding and recall than the verb always showing the same bias

To summarise, Kasof & Lee argue that the implicit causality bias is due to the arguments of a verb carrying uneven causal weights. They argue that the Agent and the Stimulus carry more causal weight than the Patient or the Experiencer. This is due to the fact the Agent/Stimulus are Actors, and all other things being equal, an Actor is most salient. Kasof & Lee discuss a number of findings that cannot be accounted for by covariation theory. However, most of these findings are of a correlational nature and do not prove that the implicit salience theory is correct. Further, the one experiment reported by Kasof & Lee, which would produce convincing evidence in favour of the salience theory, is difficult to interpret and therefore does not clearly support the implicit salience theory.

## **2.4 Implicit causality and the priming hypothesis**

### **2.4.1. The priming hypothesis**

The final explanation of the implicit causality bias is the priming hypothesis. According to the priming hypothesis, the implicit causality bias is the result of the mental lexicon's structure. This theory is based on an observation by Brown & Fish (1983a;



1983b), which was later pursued by Hoffman & Tchir (1990), that derived adjectives<sup>12</sup> corresponding to the Agent and Stimulus are more common than those referring to the Patient or Experiencer are. Hence, for the Agent-Patient verb *help* there exists an adjective *helpful* referring to the Agent, but no adjective *helpable* referring to the Patient. Similarly, for the Experiencer-Stimulus verb *notice*, there is an adjective *noticeable* referring to the Stimulus, but no corresponding adjective *noticeful* referring to the Experiencer. Finally, Stimulus-Experiencer verbs, such as *compete with*, have an associated adjective *competitive* referring to the Stimulus, but not *competeful*, referring to the Experiencer. Thus, the priming hypothesis states that the mental lexicon is structured as an associative network. When a verb is accessed, the corresponding derived adjective is also activated in the mental lexicon. The argument that this adjective refers to is therefore put into focus, and so, causal attributions are made to that argument (see Figure 9).

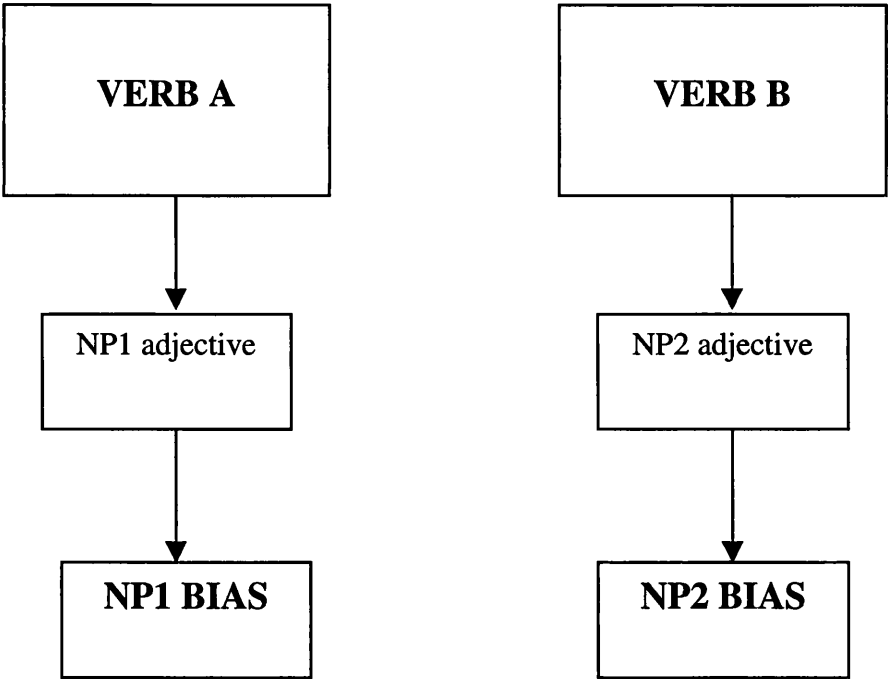


Figure 9: The priming hypothesis states that if a verb has a derived adjective then activation will pass from the verb, to the adjective, which will lead to an implicit causality bias. The left-hand diagram shows that for Verb A, which has a derived adjective referring to NP1, an NP1 bias is observed. Similarly, on the right-hand side: for Verb B, which has a derived adjective referring to NP2, a NP2 bias is observed.

<sup>12</sup> Hoffman & Tchir include derived nouns as well as derived adjectives in their model.

## 2.4.2. Tests of the priming hypothesis

Based on corpus analyses a number of researchers have shown that there is a strong correlation between the derived adjective of a verb and the causal bias of that verb. Just as derived forms for the Agent and Stimulus are more common, Agent and Stimulus causal biases are more common (*e.g.* Brown & Fish, 1983a; Gilovich & Regan, 1986; Hoffman & Tchir, 1990).

Further, in a series of experiments Hoffman & Tchir (1990) demonstrated the effect of derived adjectives on attributions. They compiled a list of verbs that had a derived adjective referring to one argument. So, for example, the Agent-Patient verb *combated* has the derived adjective *combative* referring to the Agent, while *identified* has the derived adjective *identifiable* referring to the Patient. They found that derived adjectives increased causal attributions to the associated argument. Furthermore, this effect of derived adjective was still apparent when the effect of covariation information was covaried out. Finally, participants were presented with nonsense words and a corresponding derived adjective (*e.g.* for the verb *gelf* participants were presented with either *gelfing* referring to the sentence subject, or *gelfable* referring to the sentence object). Once more, causal attributions were influenced by the adjective. Similarly, Greene & McKoon (1995) varied the derived adjective and verb orthogonally and again found evidence in favour of the priming hypothesis.

Holtgraves & Raymond (1995) also reported a series of experiments where they find evidence in support of the priming hypothesis. The priming hypothesis states that the derived adjectives associated with the Agent and the Stimulus are activated in memory when the associated verb is activated. So, Holtgraves & Raymond predicted that when a verb is presented in a sentence, the names associated with the Agent and Stimulus would be recalled more than the names associated with the Patient or Experiencer. They tested their hypothesis using a cued recall paradigm. And indeed, the results showed a recall superiority for Agent and Stimulus names over Patient and Experiencer ones.

Consequently, there is converging evidence from a number of sources supporting the priming hypothesis. However, there are a number of problems with this account, both theoretically and empirically. First, consider the empirical evidence. Although the above studies show an effect of derived adjectives on causal attributions, other studies find no, or weak effects of adjective reference on causal attributions (*e.g.* Corrigan, 1988; 1992; 1993; Semin & Marsman, 1994, Experiment 1; see also Rudolph & Forsterling, 1997).

Even in the cases where an effect of derivational form is found, the evidence is not sufficient to prove the priming hypothesis. If the derived adjective alone led to the implicit causality bias, then when the derived adjective and verb type are varied orthogonally, as in Hoffman & Tchir and Greene & McKoon, a main effect of derivational form would be expected. Instead, what these studies actually showed was a derived adjective by verb type interaction, such that when the derived adjective referred to the argument that was the causal bias, then causal attributions were even stronger. On the other hand, when it did not refer to the causal bias, the causal attributions were less pronounced. So, rather than the derived adjective causing the implicit causality bias, it appears to modify the strength of an existing bias. Furthermore, this modification of the implicit causality bias does not appear to hold across all tasks assessing the implicit causality bias.

For example, Semin & Marsman (1994) investigated two different operationalisations of implicit causality: implicit causality as event instigation and implicit causality as a dispositional inference (see section 1.3.2.1). When implicit causality was defined as an event instigation inference then participants were asked to state who brought about the event; whereas when it was defined as a dispositional inference participants were asked to say whether the event occurred because of some property of the sentence subject or some property of the sentence object. The previous chapter showed that the two operationalisations resulted in different patterns of attributions. Semin & Marsman found that the derivational adjective did not affect attributions when an event instigation question was asked, but did affect attributions to dispositional questions. Dispositional inferences to the sentence subject were weakened in the presence of a derived adjective referring to the sentence object.

Similarly, although Greene & McKoon (1995) found an effect of derivational form on attributions, this effect was relatively small when participants completed sentence fragments with a causal connective, as opposed to when they performed a causal rating task. The former task is open-ended as to whether participants refer to event instigation or dispositions, while the latter asks only a dispositional question. Hence, Greene & McKoon's results confirmed the findings of Semin & Marsman, that the derived form of an adjective only affects some sorts of causal questions, and even then it only acts so as to weaken the existing bias, rather than actually mediating the bias.

Secondly, consider the Holtgraves & Raymond findings that participants remember the names of those characters that are in the argument slot of the derived

adjective better than those that are not. There is a serious flaw in this study. Holgraves & Raymond did not use verbs where the derived adjective and the verb type were varied orthogonally. They only ever presented participants with verbs where the derived form was the same as the causal bias, that is the adjective always referred to the Agent or the Stimulus. Therefore, better recall could be the result of the causal bias itself, and not the adjective.

Finally, a number of converging sources show that causal biases are still observed when there is no supporting derivational morphology. For example, even Agent-Patient verbs that do not have an associated derived adjective show a reliable causal bias (Corrigan, 1988). Further, it has been shown that children between the ages of 18 months to 6 years do not have knowledge of derivational morphology (Jaeger, 1992). But, Corrigan & Stevenson (1994) have demonstrated that three to four-year-olds display an implicit causality bias. Therefore, it is highly unlikely that derivational morphology underlies the causal bias displayed by these children. And even more convincingly, languages that do not have verb-based derived adjectives, such as Chinese and Japanese, show an implicit causality bias (Brown, 1986). This cannot possibly be due to priming.

So, on empirical grounds alone, it appears that derived adjectives do not mediate the implicit causality bias. To this, a number of theoretical arguments can be added. Two of these are considered.

First, those that support the priming hypothesis show that derived adjectives referring to the causal bias are more common than those referring to the other argument. That is, adjectives referring to the Agent and Stimulus are most common. However, the English language is not biased to create adjectives that preferentially refer to the Agent and the Stimulus. There do exist suffixes that can be used to produce adjectives referring to the Patient and Experiencer. For example, by attaching *-ful* and *-able*, new adjectives can be created that refer to the Patient and Experiencer, such as, *helpable*, *noticeful* and *competeful* which were introduced above. The fact that these adjectives do not exist suggests that:

“derivational forms of adjectives of interpersonal interaction are in themselves a phenomenon in need of explanation and do not constitute a satisfactory explanation of implicit causality” (Rudolph & Forsterling, 1997, 215).

And secondly, if attributions are the result of a derivational adjective being accessed, then one would expect that the same attributions would be made regardless of

the type of question asked. It is assumed that when a verb is accessed the associated adjective is accessed. This is in focus and is used to answer any question, which in the case of the implicit causality bias is a causal question. However, one would expect to get the same attributions when a consequential question is asked. So, the priming hypothesis predicts that if instead of presenting participants with *Ted helped Mary because* they were presented with *Ted helped Mary so* attributions would be to Ted in both cases because the derived adjective refers to Ted in both cases. This is not the case. When a consequential question is asked of Agent-Patient verbs participants preferentially refer to the Patient, in this case Mary (e.g. Au, 1986; Fiedler & Semin, 1988; Stewart, Pickering & Sanford, 1998). The priming hypothesis cannot explain this result.

To summarise, it has been argued that the implicit causality bias is the result of priming. When a verb is accessed, the associated derived adjective is accessed, which then puts focus on the associated argument. This argument is taken to be the cause of the target event. However, there is a huge body of evidence against this proposal. First, the empirical evidence is weak. The strongest claim that can be made from the data is that when the derivational form of a verb is inconsistent with the bias of the verb, then there is a slight weakening of the bias. Secondly, the priming hypothesis cannot be correct because even in situations where there is no derivational morphology, the implicit causality bias is still apparent. This is the case for some words in English, for children without knowledge of derivational morphology as well as in languages that do not have any derivational morphology. Finally, the priming hypothesis predicts that the same argument will be focussed whenever the verb is accessed. This is not the case as can be seen when participants are asked a consequential question instead of a causal one.

## 2.5. Conclusion

In this chapter, the explanations of the implicit causality bias were reviewed. There are three major theories: implicit covariation theory, implicit salience theory and the priming hypothesis. According to implicit covariation theory, generalisations about how people in the real world behave are lexicalised as part of the verbs meaning. When a verb has low consensus and low distinctiveness information associated with it then it has a NP1 bias; but when it has high consensus and high distinctiveness information associated with it then it has a NP2 bias. When participants are explicitly asked to make generalisation inferences, there is a reliable correlation between these inferences and the

implicit causality bias. However, there are a number of problems with this account. First, the evidence cited is only correlational. It does not prove that implicit covariation causes implicit causality. Second, there are a number of factors that influence the implicit causality bias, which cannot be reduced to covariation of entities in the real world, for example, the difference in attributions between active and passive voice; the correlation between estimates of speed and implicit causality judgements (Kasof & Lee, 1993, section 2.2.2.2), and the effect of derived adjective on the bias (Hoffman & Tchir, 1990, section 2.2.3.2). Finally, the type of noun phrase that is introduced into a sentence with an implicit causality bias verb influences causal attributions (see sections 1.3.5 and 2.2.1.2). Potentially, a modified implicit covariation account could explain these findings, but this account would have to be well-specified (see below).

Implicit salience theory explains the implicit causality bias by assuming a difference in salience between arguments. Agent and Stimulus arguments are more salient than Patient and Experiencer arguments because the Agent and the Stimulus produce an action, and actors are more salient than non-actors. However, the experimental evidence for the implicit covariation account is only partial, and in some cases very difficult to interpret. Furthermore, the theory is circular, as there is no independent mean of presupposing causality to identify the salient argument. This is illustrated by the example of Agent-Evocator verbs. Agent-Evocator verbs are NP2 biasing. Implicit salience theory would predict the Evocator to be more salient. But, surely a strong determinant of salience would be the person who caused the event. Therefore, a clear, independent test of salience is required, as well as a controlled test of the relationship between salience and implicit causality.

Finally, the priming hypothesis was reviewed. The priming hypothesis states that derived adjectives associated with the Agent and the Stimulus are more common than those associated with the Patient or Experiencer. When a verb is activated, the derived adjective also receives activation, which makes the associated argument focal. When a participant is then asked to make a causal attribution, it is made to the focussed argument. There are clear reasons to reject this theory. First, there is very weak evidence for this theory; while there is convincing evidence against the priming hypothesis. And second, English morphology does not preclude the existence of derived adjectives associated with the Patient or the Experiencer. The fact that such words do not exist, or are very infrequent, is in need of explanation – and is not an explanation of the implicit causality bias.

So, it seems that none of the explanations of the implicit causality bias are adequate. Apart from the individual problems associated with each theory there is a general problem in common to all three: the assumption that the implicit causality bias requires a lexical explanation. There are good reasons to believe this assumption is not correct.

First, there is a criticism based on discursive theory. Edwards & Potter (1993) criticise most of the experiments investigating the implicit causality bias. They state that presenting participants with “decontextualised sentences, devoid of stake and interest is not interesting to attribution theory, it “simply confirm(s) intrasentential semantics” (Edwards & Potter, 1993, 26). Their critique is actually a broader critique on the view of language held by cognitive psychologists, where language is viewed as providing a window on the mind. Edwards & Potter believe it is more appropriate to think of language as a type of social action. So, for example, given the sentence *Ted scolded Mary* there are three things happening: (1) what Mary did; (2) what Ted did in scolding Mary and (3) what the current speaker is doing by calling the action scolding. They argue that the social act (3) is the starting point in investigating attributions, while studies investigating the implicit causality bias are interested in (1) (see also Clark, 1996).

However, although Edwards & Potter state that the above studies are merely “empirical verifications of linguists presuppositions” they do not deny the existence of the implicit causality bias. In fact, they themselves argue that the empirical findings are useful in order to determine the causal processes involved in using verbs. Therefore, although fundamentally Edwards & Potter are arguing for a more ecologically valid look at attributions, this does not entail that the implicit causality bias is not a lexically based phenomenon. But it does alert researchers to a wider context than that of just a single verb.

One consistent thing that has been shown to affect attributions other than the verb is the type of noun phrase the verb is presented with (Garnham, Traxler, Oakhill & Gernsbacher, 1996; see also Garvey & Caramazza, 1974). Garnham, Traxler, Oakhill & Gernsbacher (1996) argued that the implicit causality bias is an event based phenomenon, rather than a lexical one. When a participant reads a text, they build a mental model. The mental model consists of representations of people, things, events *etc.*. This representation is the result of the text plus background knowledge. Garnham *et al.* proposed that the implicit causality bias operated at this level – the scenario level. This type of explanation, one based on scenarios rather than lexical semantics, can account for the influence of the noun phrases on attributions. The noun phrase affects the mental

model as much as the verb. The scenarios associated with *Ted chasing Mary* and *Ted chasing a bus* as well as *Ted attacked Pat* and *The soldiers attacked the castle* are determined as much by who was doing what to whom, as well as what they were doing. The challenge a mental model account of the implicit causality bias has is to specify exactly what is in the mental model that determines attributions of causality. Indeed, Garnham *et al.* did not go into any detail about their mental model account or even specify how the mental model account differs from a lexical account. A mental model account could be one that is a modified version of covariation or salience, but that has still to be determined.

After over 25 years of research, we do not know why there's something about Mary.



## **CHAPTER 3:**

# **IS IMPLICIT CAUSALITY A FOCUSING PROPERTY OF VERBS?**

### **3.0 Introduction**

In the previous chapter, various explanations of the implicit causality bias were reviewed. According to these explanations, the causal bias gives privileged status to one argument of a verb, based on the semantics of that verb. That is, one argument of a verb is in focus in the mental representation. Researchers may disagree about the mechanisms or processes giving rise to this focussing, whether it is covariation information, salience information or the structure of semantic memory, but they do agree that causal information is coded at the lexical level, and that the causal information makes one argument focused. This can be labelled the causal focus account. However, the issue of whether verbs have a causal bias is orthogonal to the issue of whether verbs have a focussed argument. So, when examining the implicit causality bias two distinct questions can be asked: (1) do verbs have a causal bias; and (2) do verbs have a bias to one noun phrase over another, independent of content?

There have been many studies investigating the implicit causality bias (see section 1.1.), which have demonstrated a preference to refer to a particular argument of a verb when a causal question is asked. This has been taken as evidence that there is a causal focus and that verbs have a preferred argument, which is in focus. On the other hand, there is evidence that questions both of these claims. For example, Ehrlich (1980) argued that verbs do not have a causal bias or a preferred argument. Rather, she argues that the implicit causality bias is the result of asking a causal question. Furthermore, Stevenson, Crawley & Kleinman (1994; Stevenson, Knott, Oberlander & McDonald, 2000; Wilson & Stevenson, 1998) have argued that verbs do not have a causal focus but rather a consequential one. In this chapter, the questions of whether verbs have a causal focus and whether they have a consistent preference to refer to a particular argument in production will be examined.

As outlined in the previous two chapters (Chapters 1 and 2), there are a number of studies that have examined the implicit causality bias. The implicit causality bias is

observed using a number of different methods, such as rating scales, disambiguation tasks and completion tasks (see section 1.1). Regardless of the exact nature of the task, these studies presented participants with a simple scenario and then asked a causal question, either explicitly (by asking participants to make an attribution to one of the characters in the scenario) or implicitly (by asking participants to read a sentence where an explicit cause is either congruent or incongruent with the implicit cause).

These studies demonstrated that there was a preference for to refer to particular types of arguments (see section 1.2). That is, for both Stimulus-Experiencer and Experiencer-Stimulus verbs there was a preference to attribute cause to the Stimulus argument (*e.g.* Brown & Fish, 1983a; 1983b; Au, 1986 *etc.*). However, for action verbs, the preference for an argument was not quite so clear cut. Although the early studies by Brown & Fish (1983a; 1983b) seemed to suggest that action verbs demonstrated a preference for the Agent, subsequent research showed a more diffuse pattern of results. To date, the evidence suggests that there is a subset of action verbs that show a consistent bias. Agent-Evocator verbs show a reliable bias towards the Evocator role (*e.g.* Au, 1986). However, for the remainder of action verbs, namely Agent-Patient verbs, some studies suggest a strong preference for the Agent (*e.g.* Brown & Fish, 1983a; 1983b), while others a weak preference (*e.g.* Semin & Fiedler, 1988; 1989; 1992; Semin & Marsman, 1994), and still others a preference to the Patient (*e.g.* Corrigan, 1988; 1992; Hoffman & Tchir, 1990). These differences could be because Agent-Patient verbs do not have a preference to attribute cause to one argument, or that the group is made up of a heterogeneous sample of verbs.

Overall, the literature from the implicit causality bias demonstrates that, at least for some verbs, there is a preference to attribute cause to one argument of a verb. However, surprisingly few experiments have questioned whether there is in fact a *causal* bias. That is, do participants demonstrate a preference for causes even when no causal question is asked?

Ehrlich (1980) found that, in comprehension, when a verb was presented without a causal question it did not have a focussed argument. She presented participants with sentences that featured an implicit causality verb. The sentence was presented with one of three connectives, either *because*, *but* or *and*. Although, there was a preference for a resolve an anaphor consistently to particular antecedents when the connective *because* was used, the opposite preference emerged for *but*, whereas *and* showed a more evenly divided choice of antecedent. Based on this data, Ehrlich argued that implicit causality was not just a feature of the main verb. Instead, readers used linguistic knowledge

(from the verb) in combination with knowledge of event relations (from the connective) to compute the cause of an event. Ehrlich used the evidence from this study against the claim that there is an implicit causality bias. However, this need not be the case. Ehrlich actually demonstrated that the preference for a verb argument changes when the connective is changed. This is a different issue from whether verbs display a causal bias. Once again, preference to refer to an argument is confounded with preference for a particular type of relation.

Recently, Stevenson, *et al.* (1994; Stevenson *et al.*, 2000; Wilson & Stevenson, 1998) have argued that verbs do not have a causal focus. Instead of verbs having a causal focus, they argue that verbs have a consequential focus. Hence, according to Stevenson *et al.*, people do not infer causes when encountering a verb, rather they infer possible consequences.

One way of thinking of an event is as a series of sub-events, such as the initial conditions, the event itself, and the consequences following the event (*e.g.* Moens & Steedman, 1988; see Figure 10). As described above, many researchers have argued that a causal bias is part of the lexical entry of a verb. Meanwhile, other researchers have demonstrated a consequential bias associated with verbs, when connectives such as *so* are used (Au, 1986; Fiedler & Semin, 1988; Stewart, Pickering & Sanford, 1998). For example, Au (1986, Experiment 1) presented participants with sentence fragments with either no connective, *because* or *so*, in conjunction with one of the following verb types: Agent-Patient, Agent-Evocator, Stimulus-Experiencer or Experiencer-Stimulus. She found that the Stimulus was the preferred referent for both types of state verbs when a causal connective was used, while the Experiencer was the preferred referent when a consequential connective was used. On the other hand, for Agent-Evocator verbs the Evocator was preferred with a causal connective as well as with a consequential connective. Finally, although Agent-Patient verbs did not show a clear preference for a thematic role when a causal connective was used the Patient was the preferred referent with the consequential connective. When there was not a connective, an overall preference for the sentence object for active sentences was found. Similar findings are reported by Fiedler & Semin (1988) and Stewart, Pickering & Sanford (1998). These studies demonstrated a preference for a thematic role when a consequential question was asked as well as when a causal question was asked.




Figure 10: The rectangular boxes represent the sub-events associated with an event. An event consists of the initiating conditions, the event itself, and the consequences. Implicit causality theorists argue that the initiating conditions are focussed in the lexical representation of a verb; while Stevenson *et al.* argue that it is the consequences which are focussed.

Stevenson *et al.* (1994) argue that rather than the initiating conditions of the verb being focussed, as claimed by lexical theories of the implicit causality bias, it is the sub-events associated with the consequences that are focussed. This can be labelled the consequential focus account. That is, the causal focus account assumes that causal relations are focussed in the mental representation, and it is the causal information that gives privileged access to one argument of a verb. On the other hand, Stevenson *et al.* argue that consequential relations are focussed in the mental representation, and it is this information about consequences which gives privileged access to an argument.

Stevenson *et al.* claim that some arguments of a verb are more readily associated with the initiating conditions and others with the consequences of an event. Reference to one argument over another shows a preference, or a focussing, in the mental model, for the corresponding sub-event.

“In the goal-source sentences... the goal is more readily associated with the consequences of the described event, while the source is most readily associated with the initiating conditions. In agent-patient sentences where the patient is the preferred antecedent, the patient is most readily associated with consequences of the event, the agent with its initiating conditions.” (Stevenson, Crawley & Kleinman, 1994, 534)

To find which sub-event is preferred for different types of verbs, Stevenson *et al.* conducted a series of production studies. Participants were presented with a

sentence fragment, that either had no connective or the connective *because*, *so* or *and*. The presented fragment contained two names that were stereotypically the same gender, and a verb, from one of the following categories: Agent-Patient, Experiencer-Stimulus, Goal-Source or Agent-Goal/Source. The thematic roles of Agent, Patient, Experiencer and Stimulus are defined as previously (see section 1.2.1, see Table 6 below). In addition, a Goal was defined as “someone or something towards which something moves” and Source as “someone or something from which something moves”. Examples of Goal-Source verbs are *seize* and *pass*, while examples of Agent-Goal and Agent-Source verbs are *ran towards* and *ran away*, respectively.

Stevenson *et al.* claim that for action verbs the thematic roles of Agents and Sources are associated with the initiating conditions, while Patients and Goals are associated with the consequences. Therefore, in a production study, if participants preferentially refer to Agents and Sources then the cause is focussed, whereas if they refer to the Patient or the Goal then the consequence is focussed. On the other hand, they claim that state verbs, by definition, are not associated with either initiating conditions or consequences. That is, state verbs do not describe an event and hence cannot be broken down into a series of sub-events. Therefore, they argue, continuation patterns associated with Experiencer-Stimulus and Stimulus-Experiencer verbs are determined by extra-linguistic factors. Indeed, Stevenson *et al.* find evidence in favour of their hypotheses.

In Experiment 1 Stevenson *et al.* presented sentences such as *Ken admired Geoff. ...*, to which participants had to write a continuation. It was found that participants referred to Goals over Sources, Patients over Agents and the Stimulus over Experiencers, in their continuations. Further, Experiment 2 found the same preference when the sentence fragment was changed to include the connective *and* e.g. *Ken admired Geoff and....* Again, Goals and Patients were preferred (although now the Experiencer was preferred over the Stimulus). The preference for Patients and Goals is exactly what Stevenson *et al.* predict. Therefore, they take this as evidence in favour of consequences having a privileged status in the mental representation. Also, because there is a shift in the preference of referent for state verbs, they take this as evidence for state verb completions being determined by extra-linguistic factors.

When participants were presented with a connective, such as *so* and *because*, the pattern of continuations became consistent with the need for a consequential completion (*so*) or a causal continuation (*because*). Stevenson *et al.* agree with Ehrlich that a connective requires use of event relations, which are used in conjunction with the

information from the verb to compute an appropriate referent. Nevertheless, because they find a preference for a referent when there is no connective they claim that there is a default preference for a particular verb argument. Further, the preferred argument is one associated with consequences.

“The focusing properties of the verb direct attention to the endpoint or consequence of the described event.” (Stevenson, Knott, Oberlander & McDonald, 2000, p227)

However, as outlined at the beginning of this chapter the issue of whether there is a preference for a particular verb argument is distinct from the issue of whether there is a causal or consequential bias. The fact that Stevenson *et al.* demonstrate a preference for certain verb arguments does not address the issue of whether people make a causal or consequential inference. Therefore, within this chapter, three experiments are reported which try to establish whether verbs have a causal or consequential bias.

In the first two experiments, participants were presented with either a sentence fragment (Experiment 1) or a complete sentence (Experiment 2) to which they had to write a continuation. Because there was no connective, reference to a particular type of relation (cause/consequence) indicated what is focussed in the mental representation.

After the participant had produced their completions, these were then coded for which argument they referred to, and also for the content of the completion. Coding for the content of the completions is quite crucial. Although previous studies have included no connective conditions in experiments, they have only looked at the referent of continuations, but not at the content of continuations. For example, Au (1986) found that although there was an overall preference for the sentence object in active sentences without a connective, participants did introduce various connectives spontaneously such as *because*, *since*, *so* and *therefore*. It could be the case that one type of connective occurs more than others but that this is being obscured because researchers are ignoring the content of continuations.

In the studies reported here, continuations were coded for whether they referred to causes or consequences, as well as other types of continuations. The referent of the continuation and the type of continuation were coded separately. Both Experiments 1 and 2 provided evidence in favour of the causal focussing account. Therefore, in Experiment 3 participants were asked to complete a sentence fragment with a causal connective in order to establish whether the causal bias, which is evident when a causal

question was asked (Experiment 3), was the same as the one when no causal question was asked (Experiments 1 and 2).

### 3.1 Experiment 1:

In this experiment, participants were asked to complete a sentence fragment, which consisted of just a target verb and two proper names. The target verb came from one of four categories. It was an Agent-Patient, Agent-Evocator, Stimulus-Experiencer or Experiencer-Stimulus verb. The continuations that people produced were taken as indicative of what is focussed in the mental representation. Two types of information were coded from the continuations: first, what the content of their continuations was; and second, the choice of referent. Separate predictions can be made about the content and about the choice of argument.

The causal focus account, based on evidence from the implicit causality bias, predicted that participants would complete sentences with a continuation that referred to the cause of the target event, regardless of the type of verb involved. On the other hand, the consequential focus account, based on the work of Stevenson and colleagues, made two sets of predictions based on the type of verb involved. For action verbs, *Stevenson et al.* predicted a preference for consequential completions. However, for state verbs no preference for a particular type of continuation was predicted, as state verbs are not events and only events can have causes or consequences.

Based on the verb type and connective chosen (or content type *e.g.* causal or consequential), predictions were made about which arguments would be referred to. Stevenson *et al.* predicted that if a causal completion were made then participants would refer to the Agent for Agent-Patient verbs and the Stimulus for both Stimulus-Experiencer and Experiencer-Stimulus verbs. On the other hand, if a consequential completion were made then the Patient would be the preferred referent for Agent-Patient verbs, while the Experiencer would be preferred for state verbs. They make no predictions about Agent-Evocator verbs for causes or consequences.

To reiterate, predictions about the content of continuations were independent from the choice of the referent for a continuation. The crucial difference between the causal and consequential focus accounts was whether completions produced by participants were causal or consequential.

3.1.1. Method

3.1.1.1 Participants

Forty-eight University of Glasgow undergraduates participated in this experiment. All were native speakers of English.

3.1.1.2 Materials

A hundred verbs were randomly selected from a corpus consisting of interpersonal verbs. This corpus consisted of verbs previously used in the literature plus an additional set of such verbs taken from Levin (1993). The verbs were classified according to the Thematic Role Taxonomy (Table 6, section 1.2.1, see Appendix 1 for a full list of the verbs). That is, each verb was categorised as Agent-Patient, Agent-Evocator, Stimulus-Experiencer or Experiencer-Stimulus, using the definitions given in the revised Thematic Role Taxonomy. Two independent judges coded the verbs. There was 95% agreement on the coding. Disagreements in coding were resolved by a third judge. There were 44 Agent-Patient verbs, 13 Agent-Evocator verbs, 20 Stimulus-Experiencer verbs and 23 Experiencer-Stimulus verbs. This was comparable to the actual frequency of such verbs: 55% Agent-Patient, 12% Agent-Evocator, 15% Stimulus-Experiencer and 18% Experiencer-Stimulus (Rudolph & Forsterling, 1997).

The verbs were pseudo-randomly separated into two lists. Each list was presented to one subject. Every verb was paired with proper names, one stereotypically male and one stereotypically female. For example, participants were presented with:

Ted loved Mary.....  
.....

The ordering of the names was reversed for half the booklets, so that if Subject 1 saw *Ted loved Mary...* then Subject 2 saw *Mary loved Ted....* This created two different booklets for each list. For each of these booklets, two random orders of presentation were used, creating four booklets for each verb list. A further four booklets were created by reversing the order of the items. So, there were eight different booklets in total.



<b>THEMATIC ROLE</b>	<b>DEFINITION (and example)</b>
Agent	The instigator of an action <i>e.g. Ted in Ted defied Mary</i>
Patient	Someone, or something, which is affected by an action <i>e.g. Mary in Ted defied Mary</i>
Evocator	Someone responsible for the action, and/or who intends the action, and/or who initiates the action <i>e.g. Mary in Ted praised Mary</i>
Stimulus	Someone, or something, giving rise to an experience <i>e.g. Ted in Ted frightened Mary , and Mary in Ted feared Mary</i>
Experiencer	Someone having an experience <i>e.g. Mary in Ted frightened Mary , and Ted in Ted feared Mary</i>

Table 6: The thematic roles and definitions of roles used in order to classify the target verbs. Each definition is followed by a typical example of a verb from that category.

### 3.1.1.3 Procedure

Participants were asked to read a sentence fragment of the form *NP1 VERBED NP2*, and then to complete the sentence, so that it made sense. They were told there was no right or wrong answer.

### 3.1.2 Results

Two judges coded the continuations before any analyses were conducted. The coding criteria and the reliability between the judges are described below. This is followed by the data analyses.

#### 3.1.2.1 Coding

##### 3.1.2.1.1 Coding criteria

Completions were coded into one of following categories by two independent judges. The exact instructions given to coders are shown in Appendix 2.

First, responses were coded as causal. If the continuation began with the connective *because* then it was coded as causal. When the connective *because* was not included then the criterion of “necessity in the circumstances” was used (Mackie, 1980; see also Einhorn & Hogarth, 1986). Trabasso & Sperry (1985; Trabasso & van den Broek, 1985) used this criterion in order to identify causal relationships in narratives. They restated this criterion by means of a counterfactual, such that “event A is said to be necessary to event B if it is the case that had A not occurred then, in the circumstances, B would not have occurred” (Trabasso & Sperry, 1985, 598).

In this study, the judges coded a continuation as a cause if under the circumstances had the continuation not happened then the event in the main clause would not have happened. Once it had been established that the continuation was a causal one, the judge had to code whether the cause referred to the first noun phrase (NP1), the second noun phrase (NP2) or was ambiguous between the two<sup>13</sup>. The cause was taken to be the person which was referred to as subject in the continuation. This was the case for the other types of continuation, also. For example, if the main clause consisted of *Ted spoke to Mary*, and the participant continued with *because he was bored*, then the continuation would be coded as a causal one referring to NP1.

The second continuation type had a consequential relation to the main clause. First, if the participant began a continuation with *so* then it was scored as consequential. If it did not then the criterion for identifying causal relations was reversed. Again, after Trabasso & Sperry (1985) it was assumed that “the consequence is dependent in some manner on the cause or that the cause determines the consequence” (Trabasso & Sperry, 1985, 598). Hence, a continuation was defined as consequential when, if under the circumstances had the event in the main clause not happened then the continuation would not have happened. As before, once the continuation was identified as a consequential one, it was coded for whether the consequence was for NP1, NP2 or whether it was ambiguous. For example, if a participant continued *Ted spoke to Mary* with *and so she answered him back*, then it was coded as a consequential relation referring to NP2 (NP2 is the subject of the continuation).

The third category consisted of continuations that referred to the method by which the event was performed. The definition of method was taken from Norman, Rumelhart & LNR Research Group (1975). Basically, it describes how the event was

---

<sup>13</sup> The ambiguous category also included responses that referred to someone/thing other than NP1 or NP2 as the cause.

brought about. The method was further subdivided into two categories depending on whether NP1 or NP2 performed it. There was no ambiguous category in this case, as by definition the method included who was performing it. For example, if *Ted spoke to Mary* was followed by the continuation *by using a blackboard and a piece of chalk*, then this was coded as method referring to NP1. (Here, it was assumed that PRO is in subject position of *using*.)

The fourth category was a rather broad one that included continuations where the participant referred to some property of either NP1 or NP2, but where the type of relationship it held to the main clause was inherently ambiguous, or even non-existent. Hence, given *Ted spoke to Mary* the participant could respond with *shoe*, giving the sentence *Ted spoke to Mary's shoe*. In this example, the continuation would be coded as referring to Mary.

The final three categories modified the verb in some way. This could be by specifying the location *e.g. Ted spoke to Mary in the garden*, or time *e.g. Ted spoke to Mary in the morning*, or the manner of the event *e.g. Ted spoke to Mary loudly*. If the continuation could not be coded into any of the above categories then it was coded as 'other'.

### 3.1.2.1.2 Reliability

The two judges were given the coding criteria, and then independently coded the 2400 continuations. In order to test overall agreement between the two judges, Cohen's kappa statistic was used. The kappa statistic is a chance-corrected measure of agreement. It was preferred over a straightforward count of coincidence because the former eliminated the overestimate of agreement, which occurs when the data are not evenly distributed over categories, which was the case here. Hence, kappa is a more conservative measure of reliability than, for example, the contingency coefficient or Cramer's Phi (Howell, 1992).

The measure of agreement between the two judges was .90 ( $SE = .007$ )<sup>14</sup>.

### 3.1.2.2 Data analyses

Before analysing the data, the proportion of continuations in each of the categories was computed. A breakdown of the responses by content can be seen in Table 7.

---

<sup>14</sup> Data from Judge A was chosen randomly, by the toss of a coin, for all analyses of this experiment.

VERB TYPE	CONTINUATION TYPE							
	cause	consequence	method	property	Location	time	Manner	other
AP	39.1	9.2	16.9	5.9	11.4	4.4	6.3	6.9
AE	51.9	3.8	19.6	5.8	3.8	4.2	5.4	5.4
SE	22.9	7.3	54.6	18.7	4.2	2.0	2.3	3.3
ES	48.9	5.3	5.8	3.3	7.8	3.4	6.2	3.9

Table 7: Percentage of completions to each category, by verb type (AP = Agent-Patient; AE = Agent-Evocator; SE = Stimulus-Experiencer; ES = Experiencer-Stimulus)

In order to test the predictions of the causal and consequential focussing accounts, 2 (content [cause vs. consequence]) by 2 (referent [NP1 vs. NP2]) ANOVAs were conducted for each verb type, treating participants ( $F_1$ ) and items ( $F_2$ ) as random variables. (A full breakdown of the data for each verb is shown in Appendix 1.) To recap on the predictions, the causal focus account predicts a main effect for content, with more continuations referring to cause than consequence, for all verb types. On the other hand, the consequential focus account predicts a main effect for action verbs (Agent-Patient and Agent-Evocator), such that there are more continuations to consequence than cause. The consequential focus account predicts no preference for content for state verbs.

Specific hypotheses are made for the preferred referents for each verb. Stevenson et al predicted that for Agent-Patient verbs there would be a preference for the Agent for causal completions; however, for consequential continuations there should be a preference for the Patient. For Agent-Evocator verbs, no predictions were made. Finally, for both Stimulus-Experiencer and Experiencer-Stimulus verbs the Stimulus should be preferred for causal continuations but the Experiencer for consequential completions.

### 3.1.2.2.1 Agent-Patient verbs

There was a main effect of content [ $F_1$  (1, 47) = 149.97,  $MS_e$  = .55,  $p$  < .0001;  $F_2$  (1, 43) = 34.59,  $MS_e$  = .50,  $p$  < .0001], with more causal ( $M$  = 15.0) than consequential ( $M$  = 4.3) completions. This is consistent with the causal focus account, but inconsistent with the consequential account.

There was also a main effect of referent [ $F_1(1, 47) = 82.18$ ,  $MS_e = .21$ ,  $p < .0001$ ;  $F_2(1, 43) = 10.67$ ,  $MS_e = .19$ ,  $p < .002$ ]. This was due to more of the completions referring to NP2, *i.e.* the Patient ( $M = 13.0$ ) than to NP1, *i.e.* the Agent ( $M = 6.34$ ).




Figure 11: Percentage of continuations to content and referent for Agent-Patient verbs (error bars show  $MS_e$ ).

Finally, there was a reliable interaction between content and referent [ $F_1(1, 47) = 14.57$ ,  $MS_e = .09$ ,  $p < .0001$ ;  $F_2(1, 43) = 4.29$ ,  $MS_e = .08$ ,  $p < .04$ ]. There were more completions to NP2 when it was a causal completion [ $t_1(1, 47) = 6.66$ ,  $p < .0001$ ;  $t_2(1, 43) = 2.77$ ,  $p < .008$ ], as well as when it was a consequential completion [ $t_1(1, 47) = 2.28$ ,  $p < .03$ ;  $t_2(1, 43) = 1.78$ ,  $p < .08$ ] (see Figure 11). This interaction is most likely the result of the small number of observations for consequential continuations, which makes the difference between NP1 and NP2 referents for consequential completions much smaller than the difference for causal completions.

Overall, participants preferred to talk about the cause of an Agent-Patient verbs than the consequence. As predicted by Stevenson *et al.* there were more continuations referring to the Patient when there was a consequential completion. There were also more completions to the Patient when there was a causal completion, which is inconsistent with Stevenson *et al.*'s predictions, but this is consistent with the findings of, for example, Corrigan (1988; 1992; Hoffman & Tchir, 1990). Therefore,

participants preferred talking about causes and they preferred talking about the Patient – regardless of whether it was a cause or a consequence.

3.1.2.2.2 Agent-Evocator verbs

Once again, consistent with the causal focus account, there was a main effect of content [ $F_1(1, 47) = 150.88$ ,  $MS_e = 2.40$ ,  $p < .0001$ ;  $F_2(1, 12) = 62.63$ ,  $MS_e = .51$ ,  $p < .0001$ ] such that there were more causal ( $M = 24.0$ ) than consequential ( $M = 1.76$ ) completions.




Figure 12: Percentage of continuations to content and referent for Agent-Evocator verbs (error bars show  $MS_e$ ).

Also, as predicted, there was a main effect of referent [ $F_1(1, 47) = 104.03$ ,  $MS_e = 1.56$ ,  $p < .0001$ ;  $F_2(1, 12) = 24.31$ ,  $MS_e = .30$ ,  $p < .0001$ ], with more completions to NP2, *i.e.* the Evocator ( $M = 22.0$ ) than NP1, *i.e.* the Agent ( $M = 3.85$ ).

Finally, there was a reliable interaction between content and referent [ $F_1(1, 47) = 80.59$ ,  $MS_e = 1.30$ ,  $p < .0001$ ;  $F_2(1, 12) = 28.89$ ,  $MS_e = .25$ ,  $p < .0001$ ]. There were more completions to NP2 when there was a causal completion [ $t_1(1, 47) = 9.91$ ,  $p < .0001$ ;  $t_2(1, 12) = 5.18$ ,  $p < .0001$ ]. However, there was not a reliable difference between NP1 and NP2 referent when there was a consequential completion, although the means were in the predicted direction [ $t_1(1, 47) = 1.66$ ,  $p < .10$ ;  $t_2(1, 12) = 1.44$ ,  $p > .1$ ] (see Figure 12). As with the Agent-Patient verbs, this was most likely due to floor effects, *i.e.* the fact that there were such few data points for consequential completions.



### 3.1.2.2.3 Stimulus-Experiencer verbs

For state verbs, the consequential account predicts no difference between causal and consequential completions, while the causal account predicts a preference for cause. As before, there was a main effect of content [ $F_1(1, 47) = 28.14$ ,  $MS_e = .28$ ,  $p < .0001$ ;  $F_2(1, 19) = 19.39$ ,  $MS_e = .18$ ,  $p < .0001$ ], with more causal ( $M = 12.7$ ) than consequential ( $M = 3.13$ ) completions. This is consistent with the causal focus account.

Stevenson *et al.* predicted that there would be more continuations to the Stimulus when there was a causal continuation but more continuations to the Experiencer when there was a consequential continuation. There was a main effect of referent [ $F_1(1, 47) = 15.19$ ,  $MS_e = .10$ ,  $p < .0001$ ;  $F_2(1, 19) = 1.51$ ,  $MS_e = .02$ ,  $p > .1$ ]. This showed that overall there were more completions to NP1, *i.e.* the Stimulus ( $M = 9.48$ ) than NP2, *i.e.* the Experiencer ( $M = 6.35$ ). Finally, there was a reliable interaction between content and referent [ $F_1(1, 47) = 41.99$ ,  $MS_e = .36$ ,  $p < .0001$ ;  $F_2(1, 19) = 7.40$ ,  $MS_e = .11$ ,  $p < .01$ ]. As predicted, there were more completions to NP1 when there was a causal completion [ $t_1(1, 47) = 5.79$ ,  $p < .0001$ ;  $t_2(1, 19) = 2.06$ ,  $p < .05$ ]. On the other hand, there were more completions to NP2 when there was a consequential completion [ $t_1(1, 47) = 3.91$ ,  $p < .0001$ ;  $t_2(1, 19) = 3.16$ ,  $p < .005$ ] (see Figure 13).




Figure 13: Percentage of continuations to content and referent for Stimulus-Experiencer verbs (error bars show  $MS_e$ ).

### 3.1.2.2.3 *Experiencer- Stimulus verbs*

Similar to Stimulus-Experiencer verbs, there was a main effect of content [ $F_1$  (1, 47) = 184.76,  $MS_e = 2.12$ ,  $p < .0001$ ;  $F_2$  (1, 22) = 67.61,  $MS_e = 1.01$ ,  $p < .0001$ ], with more causal ( $M = 23.3$ ) than consequential ( $M = 2.35$ ) completions. This is more evidence consistent with the causal focus account.

As with Stimulus-Experiencer verbs, it was predicted that there would be more continuations to the Stimulus when there was a causal continuation but more continuations to the Experiencer when there was a consequential continuation. As before, there was a main effect of referent [ $F_1$  (1, 47) = 139.56,  $MS_e = 1.69$ ,  $p < .0001$ ;  $F_2$  (1, 22) = 58.89,  $MS_e = .81$ ,  $p < .0001$ ]. There more completions to the NP2, *i.e.* the Stimulus ( $M = 3.44$ ) than to NP1, *i.e.* the Experiencer ( $M = 22.2$ ). There was a reliable interaction between content and referent [ $F_1$  (1, 47) = 209.44,  $MS_e = 1.76$ ,  $p < .0001$ ;  $F_2$  (1, 22) = 61.25,  $MS_e = .84$ ,  $p < .0001$ ]. Consistent with the predictions, there were more completions to NP2 when there was a causal completion [ $t_1$  (1, 47) = 13.6,  $p < .0001$ ;  $t_2$  (1, 22) = 7.92,  $p < .05$ ]. However, there was no reliable difference between referents for consequential completions [ $t_1$  (1, 47) = .41,  $p > .1$ ;  $t_2$  (1, 22) = .36,  $p > .1$ ] (see Figure 13). Once more, this may well be due to floor effects.




Figure 14: Percentage of continuations to content and referent for Experiencer-Stimulus verbs (error bars show  $MS_e$ ).



### 3.1.3 Discussion

According to the causal focus and the consequential focus accounts, there would be a preference for certain types of continuations. The causal focus account predicted that there would be more completions referring to causes than consequences. On the other hand, the consequential focus account predicted that, overall, there would be more completions referring to consequences. It was found that regardless of the type of verb, causes are favoured over consequences. This was the case for state verbs too, which according to the consequential focus account would show no preference for either type of continuation. Overall, given a minimal event description, participants preferred to talk about causes rather than consequences. This is strong evidence in favour of the causal focus account.

Stevenson *et al.* also made predictions about the choice of referent, based on both the verb type and the connective type. Each of the verb types will be discussed in turn. First, for Agent-Patient verbs Agents would be the preferred referent when there was a causal continuation. On the other hand, for consequential completions a preference for the Patient was predicted. For both causal and consequential completions, a preference for the Patient was found. The pattern of causal referents is consistent with the findings of Corrigan (1988; 1992) and Hoffman & Tchir (1990). Therefore, for this set of Agent-Patient verbs it appears that there is an implicit causality bias towards the Patient.

Stevenson *et al.* do not make specific predictions for Agent-Evocator verbs. There was an overall preference for the Evocator. However, the comparison between Agent and Evocator continuations for consequential completions was not reliable. This finding will be discussed shortly.

For both types of state verbs, it was predicted that there would be a preference for the Stimulus when there was a causal completion but a preference for the Experiencer when there was a consequential completion. This was found for Stimulus-Experiencer verbs. Although, there was a preference for the Stimulus when making a causal completion for Experiencer-Stimulus verbs, no preference was demonstrated for consequential completions.


Both Agent-Evocator verbs and Experiencer-Stimulus verbs failed to show reliable differences for consequential completions. The most parsimonious explanation of this null result is that there was a floor effect. That is, there were not enough observations in the category of consequential completions to be able to detect an effect,

had there been one. Therefore, the lack of a reliable effect should be interpreted cautiously.


To summarise, this experiment demonstrated an overwhelming preference for causal completions over consequential completions for all verb types. Furthermore, preference for a particular thematic role, or argument, appeared to be the result of both the type of verb and the content of the continuation. For both types of action verbs considered, it appeared that the sentence object is the preferred referent for both causal and consequential completions. On the other hand, for state verbs, there is evidence consistent with the hypothesis that the Stimulus is the preferred referent of causal completions, while Experiencers are the preferred referent of consequential completions.

At this point, it is appropriate to consider the materials used in this experiment and the possible effect of the presentation style on the continuations produced. Participants were presented with a sentence fragment, for which they had to write a continuation. However, in Stevenson *et al.* (1994) participants were not presented with a sentence fragment in the no connective condition, but a whole sentence. Participants in that experiment had to continue with a new sentence. It could be that the differences in presentation style could explain choice of what to talk about next.

Consider the diagram below (Figure 15). This shows part of Figure 10, which depicts the breakdown of an event into a series of sub-events. Different presentation styles for materials could focus attention on different parts of the mental representation of an event. When presented with a sentence fragment, attention may be focussed on the preceding sub-events, as this puts the event itself into context, and helps create a fuller representation of the event. This is shown in the top figure. In contrast, when presented with a sentence attention may be focussed on the consequences, as the sentence indicates an end to things leading up to, and including, the event. This is shown in the bottom figure. Therefore, in Experiment 2 the presentation style was changed in order to test whether consequential completions are more focussed under certain conditions.



When given a sentence fragment attention is on the event and the initiating conditions.



When given a sentence, the full stop signals the end of the event, and attention is on the consequences.

Figure 15: The difference in attention given a sentence fragment or a sentence.

### 3.2 Experiment 2:

In this experiment, participants were presented with sentences to which they had to write a continuation. The causal and consequential focus accounts made the same predictions as in the previous experiment. The crucial issue was whether there were more causal or consequential completions.

#### 3.2.1 Method

##### 3.2.1.1 Participants

Forty-eight University of Glasgow undergraduates participated in this experiment. All were native speakers of English. None of the participants had taken part in the previous experiment.

3.2.1.2 *Materials and Procedure*

The materials and procedure were the same as the previous experiment, with the exception of the target fragment, which was presented as a sentence. The sentence was presented on one line, with dots for the participant’s response underneath each sentence. This presentation was chosen to emphasise to the participants that they should treat the target as a complete sentence. For example:

Ted loved Mary.  
.....  
.....

Participants were asked to read the sentence, and then to write a continuation to the sentence in the space provided underneath.

3.2.2 **Results**

3.2.2.1 *Coding*

The responses were coded and analysed as in Experiment 1. Once again, two independent judges coded responses according to the coding criteria. The measure of agreement between the two judges was .89 (*SE* = .007)<sup>15</sup>.

3.2.2.2 *Data analyses*

On examining the participants’ responses, some continuations were clearly not a new sentence, but continued with the previous sentence. This was obvious due to the fact that participants had used connectives and began their continuation without capital letters. 58% of the data was of this type. Therefore, all the analyses were performed both with and without these data points. Because the analyses did not differ, the results for the full data set are reported here.

Before analysing the data, the proportion of continuations to each of the categories was computed. A breakdown of the responses by content can be seen in Table 8. A full breakdown of the causal and consequential completions, by verb, is presented in Appendix 1.

---

<sup>15</sup> Data from Judge A was chosen randomly, by the toss of a coin, for all analyses.

VERB TYPE	CONTINUATION TYPE							
	cause	consequence	method	Property	location	time	manner	other
AP	56.6	18.8	13.8	1.1	4.5	1.3	2.8	0.9
AE	62.2	13.1	12.5	2.2	3.2	1.6	3.2	1.9
SE	55.8	10.4	30.2	0.6	0.0	0.6	1.0	1.3
ES	67.0	16.7	3.3	5.1	4.2	1.9	0.9	0.9

Table 8: Percentage of completions to each category, by verb type (AP = Agent-Patient; AE = Agent-Evocator; SE = Stimulus-Experiencer; ES = Experiencer-Stimulus)

3.2.2.2.1 Agent-Patient verbs

As in Experiment 1, there was a main effect of content, with more causal (M = 25.0) than consequential (M =8.76) completions, [ $F_1$  (1, 47) = 128.97,  $MS_e$  = 1.29,  $p$  < .0001;  $F_2$  (1, 43) = 64.37,  $MS_e$  = 1.17,  $p$  < .0001]. This preference for referring to causes was consistent with the causal focus account.

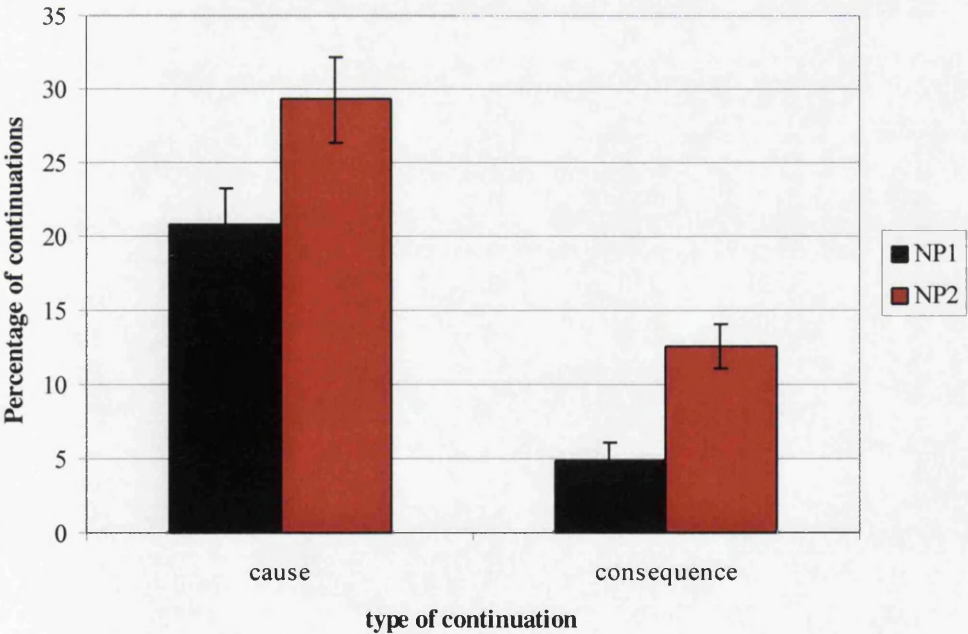


Figure 16: Percentage of continuations to content and referent for Agent-Patient verbs (error bars show  $MS_e$ ).

There was, also, a main effect of referent, with more completions to NP2, *i.e.* the Patient ( $M = 20.9$ ) than NP1, *i.e.* the Agent ( $M = 12.9$ ) [ $F_1(1, 47) = 20.49$ ,  $MS_e = .29$ ,  $p < .0001$ ;  $F_2(1, 43) = 9.62$ ,  $MS_e = .29$ ,  $p < .003$ ]. However, unlike the previous study, there was not a reliable interaction between content and referent [ $F_1(1, 47) = .008$ ,  $MS_e = .00007$ ,  $p > .1$ ;  $F_2(1, 43) = .026$ ,  $MS_e = .0006$ ,  $p > .1$ ] (see Figure 16). Overall then, this study, also, showed a preference to refer to the Patient, regardless of whether the content of the completion referred to causes or consequences.

### 3.2.2.2.2 Agent-Evocator verbs

Again, the results for Agent-Evocator verbs were the same as in Experiment 1. There was a main effect of content, with more causal ( $M = 29.2$ ) than consequential ( $M = 5.93$ ) completions [ $F_1(1, 47) = 140.94$ ,  $MS_e = 2.77$ ,  $p < .0001$ ;  $F_2(1, 12) = 66.08$ ,  $MS_e = .58$ ,  $p < .0001$ ]. This was consistent with the causal focus account.




Figure 17: Percentage of continuations to content and referent for Agent-Evocator verbs (error bars show  $MS_e$ ).

There was, also, a main effect of referent, with more completions to NP2, *i.e.* the Evocator ( $M = 24.4$ ) than NP1, *i.e.* the Agent ( $M = 10.7$ ) [ $F_1(1, 47) = 79.08$ ,  $MS_e = 1.25$ ,  $p < .0001$ ;  $F_2(1, 12) = 12.94$ ,  $MS_e = .22$ ,  $p < .004$ ]. Finally, there was a reliable interaction between content and referent [ $F_1(1, 47) = 70.35$ ,  $MS_e = 1.06$ ,  $p < .0001$ ;  $F_2(1, 12) = 6.54$ ,  $MS_e = .13$ ,  $p < .03$ ]. There were more completions to NP2 when there

was a causal completion [ $t_1(1, 47) = 9.75, p < .0001$ ;  $t_2(1, 12) = 3.14, p < .008$ ]. However, there was not a reliable difference between NP1 and NP2 referent when there was a consequential completion [ $t_1(1, 47) = .78, p < .10$ ;  $t_2(1, 12) = .64, p > .1$ ] (see Figure 17). This was exactly the same pattern of results as found in the previous experiment.

### 3.2.2.2.3 Stimulus-Experiencer verbs

As predicted by the causal focus account, there were more causal ( $M = 12.7$ ) than consequential ( $M = 3.13$ ) completions [ $F_1(1, 47) = 129.02, MS_e = 2.29, p < .0001$ ;  $F_2(1, 19) = 75.95, MS_e = 1.06, p < .0001$ ].

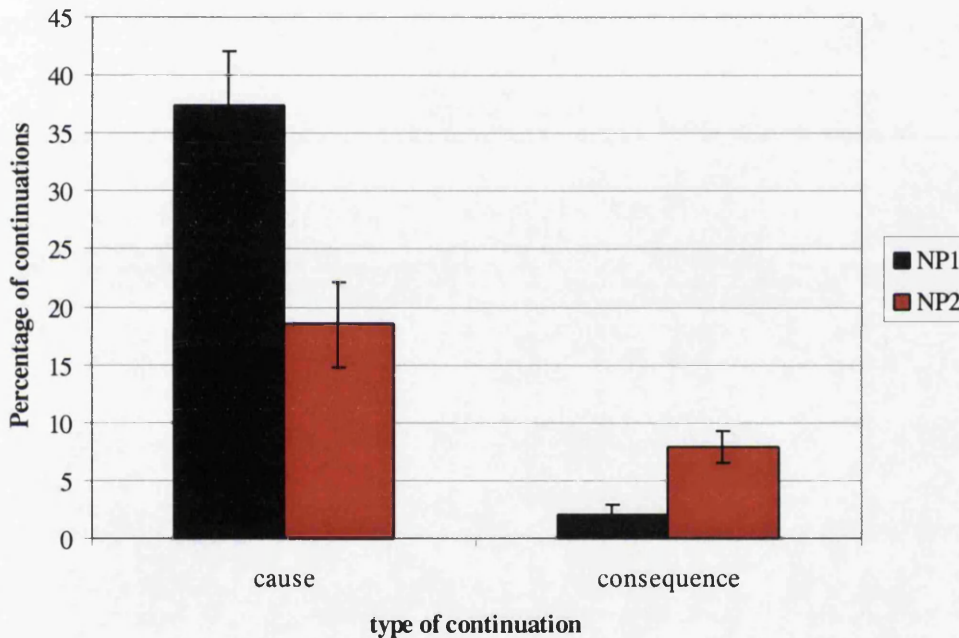


Figure 18: Percentage of continuations to content and referent for Stimulus-Experiencer verbs (error bars show  $MS_e$ ).

Furthermore, there was a main effect of referent, with more completions to NP1, *i.e.* the Stimulus ( $M = 9.48$ ) than NP2, *i.e.* the Experiencer ( $M = 6.35$ ) [ $F_1(1, 47) = 12.18, MS_e = .29, p < .001$ ;  $F_2(1, 19) = 3.84, MS_e = .08, p < .07$ ].

Finally, there was a reliable interaction between content and referent [ $F_1(1, 47) = 44.94, MS_e = .93, p < .0001$ ;  $F_2(1, 19) = 9.86, MS_e = .31, p < .005$ ]. There were more completions to NP1 when there was a causal completion [ $t_1(1, 47) = 5.37, p < .0001$ ;  $t_2(1, 19) = 2.65, p < .02$ ]. On the other hand, there were more completions to NP2 when there was a consequential completion [ $t_1(1, 47) = 4.19, p < .0001$ ;  $t_2(1, 19) = 3.69, p < .001$ ].



.002] (see Figure 18). Once again, this was the same pattern of results as found in the previous study.

3.2.2.2.4 *Experiencer- Stimulus verbs*

As predicted by the causal focus account, there was a main effect of content, with more causal ( $M = 23.3$ ) than consequential ( $M = 2.35$ ) completions [ $F_1(1, 47) = 188.55, MS_e = 2.92, p < .0001; F_2(1, 22) = 45.51, MS_e = 1.37, p < .0001$ ].

There was a main effect of referent, with more completions to the NP2, *i.e.* the Stimulus ( $M = 3.44$ ) than NP1, *i.e.* the Experiencer ( $M = 22.2$ ) [ $F_1(1, 47) = 184.41, MS_e = 2.58, p < .0001; F_2(1, 22) = 66.46, MS_e = 1.23, p < .0001$ ]. There was, also, a reliable interaction between content and referent [ $F_1(1, 47) = 174.93, MS_e = 2.86, p < .0001; F_2(1, 22) = 87.93, MS_e = 1.35, p < .0001$ ]. There were more completions to NP2 when there were causal completion completions [ $t_1(1, 47) = 14.73, p < .0001; t_2(1, 22) = 9.49, p < .0001$ ]. However, there was no reliable difference between referents for consequential completions [ $t_1(1, 47) = .82, p > .1; t_2(1, 22) = .51, p > .1$ ] (see Figure 19). Yet again, this is identical to the pattern of results found in the previous study.




Figure 19: Percentage of continuations to content and referent for Experiencer-Stimulus verbs (error bars show  $MS_e$ ).



### 3.2.3 Discussion

The results from this study were almost identical to the results from the previous study. There was an overwhelming preference for referring to causes rather than consequences, regardless of the type of verb. As for referent choice, once again, the Patient was preferred for Agent-Patient verbs, when making a causal or a consequential completion. This was the case for Agent-Evocator verbs, too. The Evocator was the preferred referent for causal completions; however, there was not a reliable preference for either referent when making a consequential completion.

For Stimulus-Experiencer verbs, there were more completions to the Stimulus when there was a causal completion, but more to the Experiencer when there was a consequential completion. Consistent with this, there were more completions to the Stimulus for causal completions in the case of Experiencer-Stimulus verbs, also. However, there was not a reliable preference for the Experiencer in the case of consequential completions. Once more, this lack of preference when making consequential completions for Agent-Evocator verbs and Experiencer-Stimulus verbs could be due to floor effects.

It was hypothesised that the addition of a full stop might change the focus of attention, so that focus was on the consequences more than when a sentence fragment was presented (as in the previous study). In order to test this hypothesis directly, a 2 (content [cause vs. consequence]) by 2 (experiment [fragment vs. full stop]) ANOVA was conducted. If the addition of a full stop focused attention on the consequences of an event, then there would be an interaction between the content of the continuation and the experiment: there would be an increase in the number of consequential completions in the full stop condition in comparison to the fragment condition; on the other hand, there would be a decrease in causal completions the full stop condition compared to the fragment condition.

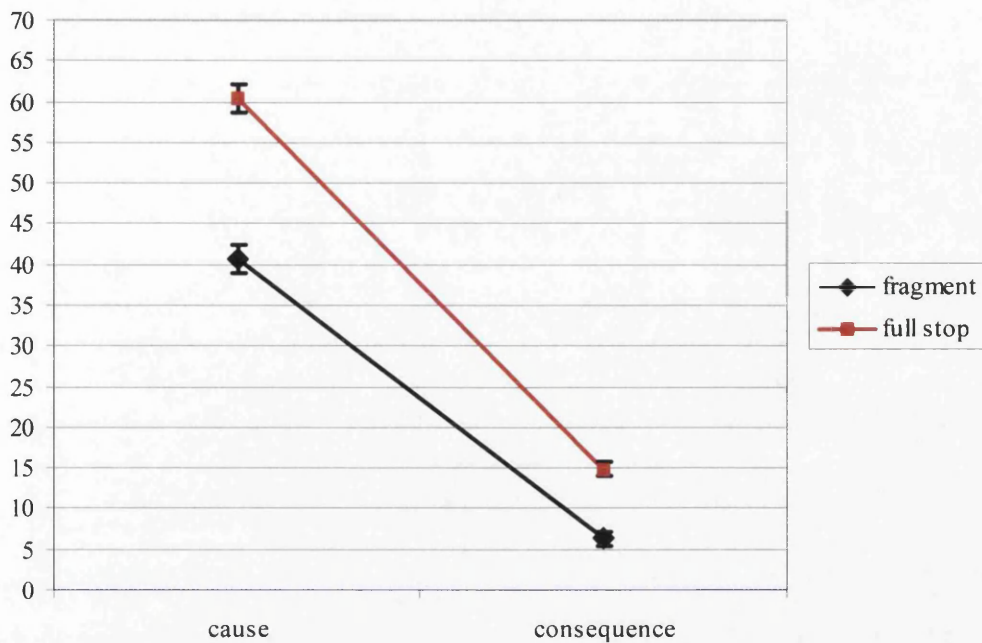


Figure 20: Comparison of continuations to cause and consequence, when presented with a sentence fragment or a complete sentence, with a full stop (error bars show  $MS_e$ ).

There was a reliable interaction between content and experiment [ $F_1(1, 94) = 14.33$ ,  $MS_e = .15$ ,  $p < .0001$ ;  $F_2(1, 198) = 11.69$ ,  $MS_e = .41$ ,  $p < .001$ ] (see Figure 20). As predicted there were more continuations to consequence when a sentence was presented in the full stop condition than when it was presented as a sentence fragment [ $t_1(1, 94) = 6.77$ ,  $p < .0001$ ;  $t_2(1, 198) = 6.16$ ,  $p < .0001$ ]. However, there were also more continuations to cause when a sentence was presented in comparison to when a sentence fragment was presented [ $t_1(1, 94) = 8.09$ ,  $p < .0001$ ;  $t_2(1, 198) = 6.97$ ,  $p < .0001$ ]. In fact, opposite to the prediction, the interaction was the result of a bigger increase in the number of causal completions than the consequential completions.

That is, there was an overall increase in the number of causal and consequential completions in Experiment 2, when a full stop is introduced. This could be explained by a decrease in other responses. For example, in the sentence fragment condition participants could continue the sentence with adjuncts. However, this was not possible in the full stop condition.

Therefore, regardless of the presentation, whole sentence or sentence fragment, there was strong evidence in favour of a causal focus account. There were far more causal completions than consequential completions as predicted by the causal focus account. Hence, verbs did indeed demonstrate an implicit *causality* bias. In order to determine whether this implicit causality bias maintains preference for a particular

argument, which was what previous studies demonstrated, a final study was conducted. In it, participants were given a sentence fragment with a causal connective. According to the causal focus account there would be a positive correlation across verbs between the referent of the causal continuations in Experiment 1 and Experiment 3, as well as between Experiment 2 and Experiment 3.

### 3.3 EXPERIMENT 3:

In this experiment, participants were asked to produce completions for a sentence fragment with the causal connective *because*. This provided a baseline to compare the results from Experiments 1 and 2. According to the causal focus account, the causal bias of a verb is inherent in the lexical entry. Therefore, the same argument would be referred to when a causal question was asked (Experiment 3) and when no question was asked (Experiments 1 and 2). Therefore, two predictions were made. First, there would be a positive correlation across verbs between the causal completions referring to NP1 in this experiment and the causal completions referring to NP1 in Experiments 1, and between this experiment and Experiment 2. The same was predicted for causal completions referring to NP2. A positive correlation would show that the causal bias was maintained, even when no causal question is asked.

#### 3.3.1 Method

##### 3.3.1.1 Participants

Forty-eight University of Glasgow undergraduates participated in this experiment. All were native speakers of English. None had participated in either of the previous experiments.<sup>16</sup>

##### 3.3.1.2 Materials and Procedure

The materials and procedure were identical to the previous two experiments, except for the addition of the causal connective *because*. Participants were presented with fragments such as *Ted loved Mary because....* They were told to complete the sentence so that it made sense.

---

<sup>16</sup> One participant did not complete a booklet, and so was replaced.

### 3.3.2 Results

#### 3.3.2.1 Coding criteria

Participants' answers were coded into one of five categories. Reference to an argument by means of a repeat name was coded separately from those using a pronoun. This was done because this experiment attempted to establish which argument was in focus in the participants' mental representation. Full noun phrases, or proper names can be used to refer to an antecedent that is no longer in focus in the mental model; whereas pronouns preferentially refer to an antecedent that is in focus (Gordon, Grosz & Gilliom, 1993; Gordon & Hendrick, 1998; Sanford, Moar & Garrod, 1988).

If the first pronoun mentioned by the participant referred to the subject of the main clause then the response was coded as NP1; if it referred to the object of the main clause then it was coded as NP2. (A full breakdown of the causal completions by verb is presented in Appendix 1.) Those continuations that referred to the main clause characters by repeating the noun phrase were coded separately as PN1 if they referred to the first character and PN2 if referring to the second. All other responses were coded as 'other' responses. NP1 and NP2 continuations were converted into proportions of the total continuations, for both participants and items. The following analyses were based on these converted scores. Overall, there were 0.33% responses to the sentence subject using a repeat name, 1.58% to the sentence object, and 5.08% other responses.

#### 3.3.2.2. Data analyses

It was hypothesised that there would be a positive correlation between the causal completions produced in this experiment and those produced in Experiments 1 and 2. A Pearson correlation showed that there was a reliable positive correlation between causal completions referring to NP1 when participants were presented with a sentence fragment without a connective and when they were presented with a sentence fragment with a causal connective [ $r(1, 99) = .44, p < .0001$ ]. This was found for NP2 completions, also [ $r(1, 99) = .65, p < .0001$ ].

Exactly the same pattern was seen for the causal completions in Experiment 2 and the present study. There was a reliable positive correlation between for both NP1 completions [ $r(1, 99) = .56, p < .0001$ ] and NP2 completions [ $r(1, 99) = .67, p < .0001$ ]. This further supported the causal focus account.

The data was also analysed in order to test whether preference for referent was the same as in the previous two experiments. That is, was there a preference for NP2

for Agent-Patient, Agent-Evocator and Experiencer-Stimulus verbs, and a preference for NP1 in the case of Stimulus-Experiencer verbs?

Agent-Patient verbs did not lead to reliably more NP2 ( $M = 43.6$ ) than NP1 ( $M = 47.7$ ) completions [ $t_1(1, 47) = 1.07, p > .1$ ;  $t_2(1, 43) = .88, p > .1$ ]. This differed from the results of the previous two experiments, where there was a reliable preference for the Patient.

However, the other three verbs showed the same preference as previously. There were more NP2 ( $M = 32.4$ ) than NP1 ( $M = 57.1$ ) completions [ $t_1(1, 47) = 3.78, p < .0001$ ;  $t_2(1, 12) = 3.16, p < .0001$ ] for Agent-Evocator verbs. This was true of Experiencer-Stimulus verbs also [ $t_1(1, 47) = 17.19, p < .0001$ ;  $t_2(1, 22) = 11.16, p < .0001$ ] (NP1  $M = 15.2$ ; NP2  $M = 80.4$ ).

Stimulus-Experiencer verbs also behaved in the predicted manner, that is they led to more NP1 ( $M = 76.0$ ) than NP2 ( $M = 20.0$ ) completions [ $t_1(1, 47) = 12.25, p < .0001$ ;  $t_2(1, 19) = 10.40, p < .0001$ ] (see Figure 21).




Figure 21: Percentage of responses to NP1 and NP2 for each verb (error bars show  $MS_e$ ).

A closer examination of each verb category revealed some interesting points. First, for Agent-Patient verbs only half of the means were in favour of the Agent. Out of 44 verbs, 18 had more Patient than Agent completions. In fact, Agent-Patient verbs showed a whole range of biases, from NP1 biasing to neutral through to NP2 biasing verbs. Secondly, Agent-Evocator verbs, which were coded according to the criteria of the sentence object having performed some action which led to the action of the sentence subject, did not all lead to more NP2 completions. In fact, two verbs *heal* and *judge* led to more NP1 completions.

This does call into question the validity of the concept of Evocator. There are a number of Agent-Patient verbs that received more NP2 than NP1 completions. However, there were still Agent-Evocator verbs that were not NP2 biasing as ought to be the case. This range of biases associated with action verbs is quite typical of what is found in the literature (Au, 1986; Semin & Fiedler, 1988; 1989; 1992; Semin & Marsman, 1994 *etc.*, see Chapter 1).

For state verbs, the individual item means were in the predicted direction. That is, almost all Stimulus-Experiencer verbs had more NP1 completions and Experiencer-Stimulus verbs had more NP2 completions. There was one exception. *Gape at* was categorised as an Experiencer-Stimulus verb, but led to more NP1 completions. This could be due to the fact that participants were interpreting the verb as an active rather than stative verb.

From these results, it can be concluded that the preference to refer to a particular thematic role was the same in the two experiments – with one exception. Although Patients were the preferred referent for Agent-Patient verbs in the previous two experiments, there was no preference in the present experiment. Overall, verbs maintain the causal bias they have when they are presented without a causal connective as well as when they are presented with one. There was a positive correlation between the present experiment and the previous two experiments. This was consistent with the causal focus account.

### 3.4 General discussion

According to the causal focus account, there is a preference for causes and particular thematic roles. However, this view has been challenged by Stevenson *et al.*, who argue causes are not primary; rather, they argued that consequences are more important. This is the consequential focus account. Like the causal focus account, they

agreed that particular thematic roles are focussed in the mental representations. However, they argued that this is because there is a preference to think about consequences. The three experiments in this chapter showed that there is a strong preference to refer to the cause of an event, even when there is no causal connective present.

It could be argued that some verbs have a causal bias and others a consequential bias. To check if this was the case, the individual means for each verb were examined. In Experiment 1, the sentence fragment experiment 90 out of 100 verbs showed more causal completions; while in Experiment 2, 95 out of 100 verbs did. The verbs that appeared to have more completions that are consequential were not the same in the two experiments. That is, in Experiment 1 *cheated, greeted, harmed, liked, pushed, saw, slandered, stopped by, trusted* and *loved* had more consequential completions; while in Experiment 2 *commanded, distracted, flattered, kicked* and *overexcited* had more. Therefore, this increase in consequential completions is probably noise, rather than a bias associated with particular verbs.

Also, there appeared to be a general preference to refer to particular thematic roles. In Experiments 1 and 2, participants referred to NP2 for action verbs, regardless of whether they continued with a causal or consequential completion. However, there was a slightly different pattern with state verbs. For Stimulus-Experiencer verbs, the choice of referent interacted with the content of the continuation. While participants referred to the Stimulus when talking about causes, they referred to the Experiencer when talking about consequences. This is consistent with what has been found in previous studies. Action verbs which bias causality towards NP2 are associated with consequential completions to NP2; while state verbs that bias causality to NP1/NP2 bias consequentiality to NP2/NP1 respectively (Au, 1986; Stewart, Pickering & Sanford, 1998). Therefore, the referent preference for particular types of relation was the same, with or without a connective.

Although the evidence was consistent with the causal focus account, it has to be emphasised that the data does not shed any light on whether the preference for causes is a product of the lexical representation of verbs. An equally likely explanation is that the preference for causal completions found for verbs is a general preference of the cognitive system, rather than being explicitly coded in the lexicon. Given that people are constantly searching for causal relations in order to make sense of the world, and that causality is central in learning, perception, development, social cognition, categorisation, to name but a few areas (see section 0.1.), it is parsimonious to assume

that the cognitive system is tuned to seek causes. Therefore, when a person is presented with a minimal scenario the cognitive system is going to ask for an explanation of that event so that it can produce an effective response. If the system already knows the cause of an event then other sources of information become more important. It could be in cases such as this where consequential relations are crucial to the system.

This chapter showed that the implicit causality bias is not just the result of asking a causal question. Rather, there is a cognitive bias towards causes as well as a bias to refer to a particular argument. In the next chapter, an examination of the strength of the implicit causality bias was tested, by pitting causal information in the verb against explicit cues to causality.



## CHAPTER 4:

# COVARIATION INFORMATION IN DISCOURSE

### 4.0. Introduction

People make causal attributions from minimal information. In the previous chapter, it was shown that participants made causal attributions when given a sentence fragment with just a verb and two names – even when they have not been asked a causal question. In this chapter, another cue to causality was investigated, and the interaction between implicit causal information in a verb, and explicit causal information was examined.

In Chapter 2, the relationship between the implicit causality verb bias and the covariation theory of attribution was outlined. Covariation theory has been used as an explanation of the implicit causality bias. Implicit covariation information is represented in the lexicon with other information about the verb (see section 2.1). To recap, covariation theory states that the cause of an event is determined by what covaries with what. If we are trying to determine why *Ted loves Mary*, then we can look at how other people feel about Mary (consensus information); how Ted feels about other people (distinctiveness information); and how Ted feels about Mary on other occasions (consistency information). If many people love Mary, and Ted loves few other people, and he loves Mary all the time (high consensus, high distinctiveness and high consistency) then there's something about Mary which causes Ted to love her. However, if few people love Mary, and Ted loves many other people, and he loves Mary all the time (low consensus, low distinctiveness and high consistency) then Ted is the cause.

There have been attempts to study the effect of explicit covariation information on implicit causality verbs. Three questions can be posed regarding explicit covariation and the implicit causality bias. First, do explicit covariation information and implicit causality both influence attributions when presented together? If so, which is the stronger cue when making causal attributions? And finally, is covariation information a strong enough cue to over-ride the implicit causality bias?

One of the first papers to shed some light on these questions is McArthur (1972). Although she was not directly answering the above questions, she was one of the first researchers not only to empirically test the covariation theory, but also first to examine it with respect to different classes of verbs.

McArthur divided verbs into four classes: emotions (*e.g. afraid*), accomplishments (*e.g. translate*), opinions (*e.g. think*) and actions (*e.g. contribute*). She then presented participants with covariation information. So, for example, the statement *John laughs at the comedian* was followed by either high or low consensus information: *Almost everyone who hears the comedian laughs at him*; or *Hardly anyone who hears the comedian laughs at him*. This was followed by high or low distinctiveness information: *John does not laugh at almost any other comedian*; or *John also laughs at almost every other comedian*. Finally, high or low consistency information was presented, such as *In the past, John has almost always laughed at the same comedian*; or *In the past, John has almost never laughed at the same comedian*. Participants were presented with all three pieces of covariation information, *i.e.* consensus, distinctiveness and consistency information. A baseline condition where no information was presented was also included.<sup>17</sup> Following each experimental item, participants were asked to choose among four different causes: the Person (*i.e.* John), the Entity (*i.e.* the comedian), the particular circumstances (*i.e.* that particular occasion) or some combination of the three previous causes.

It was found that participants did indeed make more Entity attributions when presented with high consensus, high distinctiveness and high consistency information than in the baseline condition, while low consensus, low distinctiveness and high consistency information led to more Person attributions. There was also evidence of an effect of verb type influencing attributions. Hence, when making causal attributions participants used both explicit and implicit information. McArthur also found that the effect of verb type was relatively small compared to the effect of explicit covariation. Consequently, she concluded that covariation information is given more weight in determining causal attributions than the verb's causality bias.

Unfortunately, although McArthur demonstrated the use of implicit and explicit cues when making attributions, there were a number of problems with her experiments. For example, the verbs were not representative of implicit causality verbs, as they included non-interpersonal verbs such as *translate*. Further, the arguments of the verb

---

<sup>17</sup> The baseline task is equivalent to a test of the implicit causality of the verb.

were not always equally salient. McArthur used materials such as *John laughed at the comedian*. It has been shown that use of such role descriptions can affect causal attributions (see section 1.3.5). Therefore, there were a number of confounds in the results McArthur presented. Nonetheless, Van Kleeck, Hillger & Brown (1988) and Rudolph (1997) replicated the relative impact of verb type and explicit covariation with a more representative set of implicit causality verbs.

Van Kleeck *et al.* and Rudolph argued that implicit causality information and explicit covariation information are “pitted against each other”, that is, there is competition between the two sources of information when making causal attributions. Both sets of researchers presented explicit consensus and distinctiveness information with a target sentence containing an implicit causality verb. Van Kleeck *et al.* investigated state verbs, whereas Rudolph used action verbs, also. Participants had to indicate the cause of the target event using a rating scale. Both papers presented similar results. There was a strong influence of explicit covariation information and a weaker effect of verb bias. Hence, explicit information was given more weight when making causal attributions. This was the case, regardless of verb type: whether it was a NP1 or NP2 biasing verb; or whether it was an action or a state verb.

Furthermore, there did appear to be evidence to suggest that explicit and implicit cues were in competition with one another. Rudolph found that although there was still an effect of the implicit causality verb bias when *either* consensus *or* distinctiveness was presented, that effect disappeared when both types of covariation information were presented together. When they were presented together, the context appeared to completely over-ride the implicit causality verb bias. That is, when covariation information and implicit causality information were incongruent with one another, participants made their causal attributions based on the covariation information and *not* on the verb information. As for the effect of the verb bias, he found that as the strength of the implicit causality bias increased, the relative impact of the context information grew weaker. That is, participants gave the verb bias more weight when it had a stronger bias. Both pieces of evidence suggest that there is competition between the two sources of information.

To summarise, previous experiments showed that contextual information and verb bias both influence causal attributions, with the former always being a stronger determinant of attributions than the latter. When the contextual information is incongruent with the verb bias, it over-rides, and in fact, reverses the bias. It must be noted that the data from these experiments do not shed any light on the question of

whether covariation information underlies the implicit causality verb bias (see section 2.2.). The fact that two sources of information (covariation information and verb bias) are used in order to make causal attributions does not mean that one source of information (verb bias) can be reduced to the other (covariation information). The data shows that participants can use multiple sources of information and then reason with these sources to come to some solution about the likely cause.

In this chapter, Experiments 4 and 5 examined the effect of covariation information on implicit causality verbs in a discourse situation. As McArthur (1972; Van Kleeck, Hillger & Brown, 1988; Rudolph, 1997) showed covariation information influenced causal attributions in situations where participants were explicitly asked to make a causal attribution. The present studies investigated the effect of covariation information in a discourse setting, in contrast to the above studies which used rating tasks. Social psychologists use rating tasks as a way of testing the kinds of inferences people make about social situations. This is partly because presenting real-life events would be too time-consuming and expensive. However, not all inferences that people make are based on the visual perception of the social scenario. Instead, a lot of information we make inferences about are things that are presented to us in discourse. Therefore, a more appropriate measure of the inferences that people make is to use discourse. It is important to see if people do actually use covariation information when it is presented in discourse.

The effect of covariation information was examined separately for Agent-Patient verbs and state verbs. As shown in Chapters 1 and 3, Agent-Patient verbs tend to have a weaker and less robust bias. Therefore, it was hypothesised that if covariation information was to have an effect it might be stronger for just such verbs as they should be more malleable to contextual information. State verbs were examined for precisely the opposite reasons: in order to investigate whether contextual information can over-ride strong and robust biases.

Two hypotheses were tested: (1) does covariation information influence causal attributions in production and (2) does contextual information over-ride the implicit causality bias?

#### **4.1. Experiment 4:**

In this experiment, the effect of context on the implicit causality bias associated with Agent-Patient verbs was examined. Participants were presented with one of four

pieces of covariation information, low consensus (*Few people verbed Mary*), high consensus (*Many people verbed Mary*), low distinctiveness (*Ted verbed many people*), or high distinctiveness (*Ted verbed few people*). Each contextual sentence was followed by the target sentence fragment (*Ted verbed Mary because...*). Participants were asked to continue the target sentence with a few words, so that it made sense, given the information presented.

Covariation theory predicted that low consensus and low distinctiveness information would lead to Person attributions. This means that there ought to be more NP1 continuations and less NP2 continuations. However, high consensus and high distinctiveness information would lead to the opposite pattern of results. There ought to be more NP2 continuations and less NP1 continuations. Secondly, it was hypothesised that there would be a main effect of verb type, such that verbs displayed the same bias they had in the no context condition. Based on the previous studies it was hypothesised that covariation information would be a stronger cue when making attributions than the implicit causality bias. Finally, it was hypothesised that covariation information would over-ride the implicit causality bias. This means, that when the covariation information and implicit causality verb information were incongruent with one another, attributions should be based on the covariation information rather than the verb information.

#### **4.1.1. Method**

##### *4.1.1.1. Participants*

Thirty-two undergraduates from the University of Glasgow participated in the experiment. All were native English speakers. None of the participants had taken part in any of the previous studies.

##### *4.1.1.2. Materials*

Thirty-six Agent-Patient verbs were chosen from Chapter 3, Experiment 3. The verbs were chosen so that they represented a full range of biases, from strongly biasing NP1 verbs through to neutral and strongly biasing NP2 verbs. The full set of verbs used can be seen in Appendix 3. The verbs were classified based on the difference scores between the number of NP1 continuations and the number of NP2 continuations for each verb. The difference scores were actually continuous, but for practical reasons they were divided three discrete categories: NP1 biasing verbs ( $M = 9.42$ ); neutral verbs ( $M = 0.67$ ) and NP2 biasing verbs ( $M = -6.75$ ).

In Experiment 3 (previous chapter), the verbs were presented with two proper names and the causal connective *because*. The data from that experiment was used as the baseline in the following experiments. The responses to the context plus target sentences were compared to the responses to just the target sentence (the baseline) in order to assess the extent of the influence of the covariation information on responses.

Each experimental item consisted of a context sentence, which was presented in the past tense. This provided the covariation information. The context sentence presented one of four pieces of information. So, for example for the verb *defied* the context sentence would be *Few people defied Mary* (low consensus information), *Many people defied Mary* (high consensus information), *Ted defied many people* (low distinctiveness information), or *Ted defied few people* (high distinctiveness information).

The context sentences were made up of four different quantifiers, *few* and *very few* denoting low consensus and high distinctiveness, and *many* and *most* denoting high consensus and low distinctiveness. For example, low consensus information was presented as either *Few people defied Mary* or *Very few people defied Mary*. This allowed some variation of the context sentences to encourage participants to pay attention to the materials. The quantifiers were paired with each context type, so that within each list each quantifier appeared equally often.

The context sentence was followed by the target sentence fragment, which consisted of a male and female name, a verb and the connective *because*. Each verb was randomly assigned a male name and then a female name. For those booklets in which the male name appeared in the subject position, a parallel booklet was created in which the female name appeared in the subject position. Within each booklet, the male and female names appeared in the subject position equally often.

A Latin-square design was employed, such that the four context sentences were crossed over the three types of verb. Four random orders of items were constructed for the booklets, with a further four using the reverse random order.

#### 4.1.1.3. Procedure

Participants were randomly assigned a booklet. They were instructed verbally and given written instructions to read each of the two-line passages carefully and then complete the target sentence so that it made sense.

#### 4.1.2. Results

Each participant's response was scored as either a NP1 or NP2 continuation based on whether the first pronoun in the continuation referred to the subject (NP1) or object (NP2) of the main clause. Continuations referring to the main clause by a repeat name were coded separately. A total of 1.54% of the continuations referred to NP1 with a repeat name and 2.24% referred to NP2 in such a manner. Any remaining continuations were scored as 'other', and these comprised 3.78% of the total responses. Only continuations referring to the main clause of the target fragment with a pronoun were used for further analyses.

##### 4.2.1. Covariation and the implicit causality bias

The data was analysed in order to test whether covariation information and the implicit causality bias influenced attributions. The responses were analysed using 2 (covariation [consensus vs. distinctiveness]) by 2 (level [high vs. low]) by 3 (verb type [NP1 vs. neutral vs. NP2]) ANOVAs, treating both participants ( $F_1$ ) and items ( $F_2$ ) as random variables. Only the results for NP1 continuations are presented, as NP2 analyses showed exactly the same pattern. This is the case for the remainder of the thesis.

To recap on the predictions, there should be a main effect of level, such that low levels of covariation information give rise to more NP1 continuations than high levels of information. There should also be a main effect of verb type, such that NP1 biasing verbs give the most NP1 continuations followed by neutral verbs and finally NP2 biasing verbs.

There was a main effect of level as predicted by covariation theory [ $F_1$  (1, 31) = 42.70,  $MS_e$  = .59,  $p$  < .0001,  $F_2$  (1, 33) = 38.66,  $MS_e$  = 1.66,  $p$  < .0001], such that low levels of consensus and distinctiveness ( $M$  = 4.5) led to more NP1 continuations than high levels ( $M$  = 3.2). That is, low consensus and low distinctiveness led to more NP1 continuations in comparison to high consensus and high distinctiveness information (see Figure 22).

There was evidence of a main effect of verb type, but it was not reliable by participants, [ $F_1$  (2, 62) = 2.31,  $MS_e$  = 1.02,  $p$  > .1,  $F_2$  (2, 33) = 4.29,  $MS_e$  = 3.23,  $p$  < .02]. Pair-wise comparisons showed that the difference between NP1 and NP2 verbs was reliable by both participants and items, [ $t_1$  (1, 31) = 2.05,  $p$  < .04;  $t_2$  (1, 12) = 2.15,

$p < .03$ ]. The difference between NP1 verbs and neutral verbs was reliable by items only [ $t_1(1, 31) = .39, p > .1$ ;  $t_2(1, 12) = 1.87, p < .05$ ]; whereas the difference between neutral verbs and NP2 verbs was reliable by participants only [ $t_1(1, 31) = 1.75, p < .04$ ;  $t_2(1, 12) = 1.15, p > .1$ ]. Overall, there was a trend showing less continuations to NP1 when the verb was NP2 biasing ( $M=3.27$ ), than when it was neutral ( $M=4.06$ ), or NP1 biasing ( $M=4.29$ ). There were no other main effects or interactions.

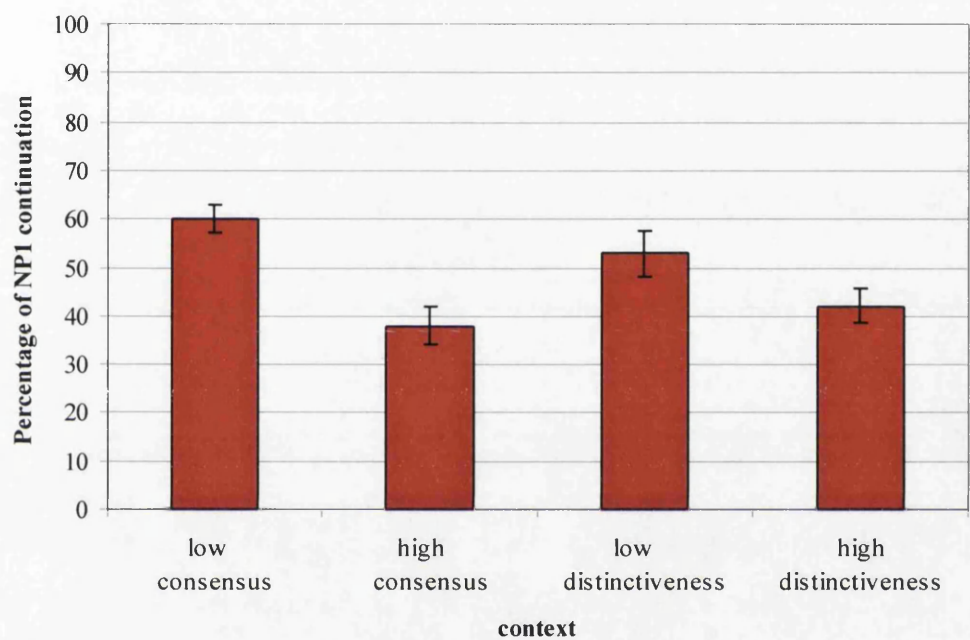


Figure 22: NP1 continuations to the four different types of covariation information (error bars show  $MS_e$ )

To summarise, the results show that covariation information does indeed influence causal attributions. Low levels of consensus and distinctiveness information led to more NP1 attributions than did high levels of consensus and distinctiveness. This is consistent with covariation theory. Despite the presence of a context sentence, the implicit causality verb bias still affected causal attributions, with NP1 biasing verbs receiving more NP1 and less NP2 continuations than either neutral or NP2 biasing verbs.

#### 4.2.2. *Is covariation strong enough to over-ride the implicit causality bias?*

The above results show covariation information influences causal attributions in production. However, it does not show the effect of covariation information on the implicit causality bias. That is, it does not indicate whether covariation information is a



strong enough cue to shift the implicit causality bias. The context could modify the bias in two ways. It could strengthen it by providing information consistent with the bias, or it could weaken, even reverse, the bias by providing information inconsistent with the verb bias.

In order to determine whether covariation information modifies the implicit causality bias, continuations for each of the four context types (Experiment 4) were compared to the continuations obtained with no context (Experiment 3). Covariation theory predicts that low levels of information will increase NP1 continuations relative to the baseline; whereas, the converse is predicted for high levels of information, which should decrease NP1 continuations relative to the baseline. Therefore, continuations to low consensus, low distinctiveness, high consensus, and high distinctiveness information were compared to continuations made when the target fragment was presented alone. Analyses were based on the proportion of responses referring to NP1. Because multiple comparisons were made alpha was set to less than .01. Only those comparisons that reached significance are reported.

If there is a reliable difference between the context and no-context situations, then it can be concluded that covariation information does modify the implicit causality bias. Further, if continuations are shifted over the 50% mark then this is taken as evidence that the contextual information caused a reversal of the bias.<sup>18</sup>

#### 4.2.2.1. *NP1 biasing verbs*

Low consensus and low distinctiveness information led to NP1 attributions. Therefore, low levels of information are consistent with NP1 biasing verbs. According to covariation theory, such information should increase the likelihood of NP1 continuations. On the other hand, high consensus and high distinctiveness should decrease NP1 continuations, as these provide information inconsistent with the bias.

Low levels of information did not increase NP1 continuations from the baseline as is predicted by covariation theory. However, as was predicted, high levels of information did produce reliably less NP1 continuations (Figure 23). This was the case for both high consensus information [ $t_1(1, 54) = 4.02, p < .0001$ ;  $t_2(1, 11) = 5.04, p < .0001$ ] and high distinctiveness information [ $t_1(1, 54) = 3.05, p < .003$ ;  $t_2(1, 11) = 3.49, p < .005$ ].

---

<sup>18</sup> Note that it is not possible to test neutral verbs for reversal, only for modification.

In fact, it could be said that the bias was reversed, as in both cases less than 50% of the continuations referred to NP1, when the verb was presented with context. 35% and 45% of continuations referred to NP1 given high consensus and high distinctiveness information respectively compared to 66% in the baseline condition. This is closer to the number of NP1 continuations observed with NP2 biasing verbs (where there are 32% NP1 continuations). Therefore, it appears that causal attributions are being made on the basis of the covariation information rather than verb information.




Figure 23: Continuations for NP1 biasing verbs to both the target fragment with no context (baseline) and when presented with covariation information (context). The baseline is shown as a line only for convenience. There is a decrease in NP1 continuations when given high consensus and high distinctiveness information, but not an increase when given low consensus and low distinctiveness information.

4.2.2.2. *Neutral verbs*

No reliable differences emerged in any of the comparisons between the four context sentences and the baseline (Figure 24). However, all the means were in the direction predicted by covariation theory: low consensus (59%) and low distinctiveness (56%) produced more NP1 continuations, whereas high consensus (44%) and high distinctiveness (40%) produced less NP1 continuations than the baseline (46%).




Figure 24: Continuations for neutral verbs to both the target fragment with no context (baseline) and when presented with covariation information (context). The baseline is shown as a line only for convenience. There is a slight increase in NP1 continuations given low consensus and low distinctiveness information and a slight decrease given high consensus and high distinctiveness information.

#### 4.2.2.3. NP2 biasing verbs

For NP2 biasing verbs, low consensus and low distinctiveness information are inconsistent with the bias, and therefore should increase NP1 continuations from the baseline. On the other hand, high consensus and high distinctiveness information are consistent with the bias and so should decrease NP1 continuations.

Low consensus information shifted attributions reliably from the baseline [ $t_1(1, 54) = 3.84, p < .0001$ ;  $t_2(1, 11) = 2.38, p < .04$ ]. There was evidence of a shift with low distinctiveness information, but this was reliable only by items [ $t_1(1, 54) = 1.53, p > .1$ ;  $t_2(1, 11) = 3.00, p < .01$ ] (Figure 4c). That is, low levels of information increased NP1 continuations, as predicted, and further modified the implicit causality bias. There was even evidence of a reversal of the bias, when low consensus information was provided. Low consensus caused 54% NP1 continuations, which was closer to the baseline for NP1 biasing verbs (66%) than it was to the baseline for NP2 biasing verbs (32%). Once again, causal attributions are made on the basis of covariation information rather than verb information.

There were no reliable differences between context that provided high levels of information and the baseline. Neither high consensus nor high distinctiveness decreased NP1 continuations (see Figure 25).

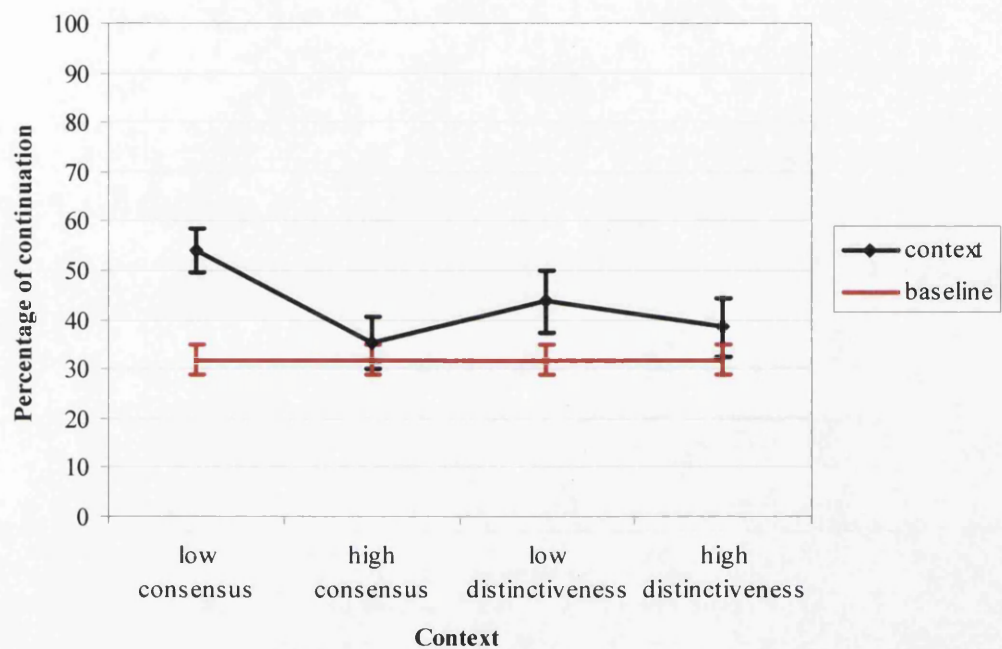


Figure 25: Continuations for NP2 biasing verbs to both the target fragment with no context (baseline) and when presented with covariation information (context). The baseline is shown as a line only for convenience. There is an increase in NP1 continuations when given low consensus and low distinctiveness information, but not an increase when given high consensus and high distinctiveness information.

**4.1.3. Summary**

Covariation theory predicts that low levels of information should increase NP1 continuations, whereas high levels of information should decrease NP1 continuations relative to the baseline. The results are consistent with this, but with some caveats. Unlike previous studies, presenting information consistent with the bias never resulted in a strengthening of the bias. This was not due to ceiling or floor effects as with both NP1 and NP2 biasing verbs the baseline was always at least 30% away from the ceiling/floor. So, movement from the baseline was always possible.

Only information that was inconsistent with the bias affected causal attributions. High levels of information decreased NP1 continuations for NP1 biasing verbs, whereas low levels of information increased NP1 continuations for NP2 biasing verbs. Van

Kleeck *et al.* and Rudolph found that when incongruent information was presented it always reversed the bias. This appeared to be the case in production also. However, there was one exception. When low distinctiveness information was presented with NP2 biasing verbs, there were still more NP2 than NP1 continuations.

It could have been the case that context had the greatest effect on neutral verbs, as there was more room for movement. This would be consistent with the Van Kleeck *et al.* and Rudolph claim that that implicit causality information and explicit covariation information are “pitted against each other”, and that the strongest cue determines the attributions. However, this was not the case. Although there was evidence that the contextual information was affecting attributions, no piece of information produced a reliable shift from baseline. This is quite counterintuitive.

One possible explanation of this could be to do with the properties of neutral verbs. An implicit causality bias can be neutral because neither of the arguments carries a causal weight. If this was the case then context would be the only cue by which to make a causal attribution, and therefore context should completely determine attributions. However, a verb could be neutral because both arguments carry a causal weight, but these cancel each other out. Hence, when context is provided it is always inconsistent with one of the causal weightings of the verb, and therefore always has an attenuated effect on causal attributions.<sup>19</sup>

## 4.2. Experiment 5:

In the previous experiment, covariation information did indeed influence causal attributions when paired with Agent-Patient verbs. Furthermore, it was the dominant cue when making causal attributions. However, this could have been the case because the verbs used in the previous study were not very strongly biased. Therefore, in this experiment an attempt was made to examine the effect of covariation information on verbs that have a strong and robust bias.

As discussed previously (section 2.2.1.1.3), the mapping of covariation roles onto thematic roles is not as simple for state verbs as for action verbs. It was argued that the best mapping is one where consensus information quantifies over NP1 for Experiencer-Stimulus verbs (as is the case for Agent-Patient verbs), but where it quantifies over NP2 for Stimulus-Experiencer verbs, and vice versa for distinctiveness

---

<sup>19</sup> There is some more suggestive evidence in favour of this hypothesis in the next chapter (Chapter 5).

information. Therefore, low consensus information for Experiencer-Stimulus verbs would be *Few people loved Mary*; while for Stimulus-Experiencer verbs it would be *Mary fascinated few people*.

Covariation theory predicts low levels of information for Experiencer-Stimulus verbs should lead to more NP1 attributions. Conversely, high levels of information should lead to fewer NP1 attributions. However, for Stimulus-Experiencer verbs, high levels of information would increase NP1 attributions and low levels would decrease attributions.

A main effect of verb type was also hypothesised, such that Stimulus-Experiencer verbs would lead to more NP1 attributions than Experiencer-Stimulus verbs.

#### **4.2.1. Method**

##### *4.2.1.1. Participants*

Thirty-two undergraduates from the University of Glasgow participated in the experiment. All were native English speakers. None of the participants had taken part in any of the previous experiments.

##### *4.2.1.2. Materials*

The materials were produced in the same manner as the previous experiment with two changes. First, the verbs were state verbs. There were thirty-two state verbs chosen from Chapter 3, Experiment 3. There were 16 Stimulus-Experiencer verbs ( $M = 15.93$ ) and 16 Experiencer-Stimulus verbs ( $M = -19.31$ ) (see Appendix 4). The continuations in Experiment 3 also provided the baseline data for the comparisons conducted.

Second, each item consisted of a context sentence followed by the target fragment (*e.g.*, *Many people liked Mary. Ted liked Mary because...*). The context sentence presented a piece of covariation information. There were four different context sentences presenting high consensus, low consensus, high distinctiveness and low distinctiveness information (see Table 9).



Context	Experiencer-Stimulus	Stimulus-Experiencer
High consensus	Many people liked Mary	Ted fascinated many people
Low consensus	Few people liked Mary	Ted fascinated few people
High distinctiveness	Ted liked few people	Few people fascinated Mary
Low distinctiveness	Ted liked many people	Many people fascinated Mary

Table 9: Examples of covariation information presented for Experiencer-Stimulus verbs and Stimulus-Experiencer verbs.

Note that while consensus information presents information about NP1 for Experiencer-Stimulus verbs; it is distinctiveness information that presents information about NP1 for Stimulus-Experiencer verbs.

4.2.1.3. *Procedure*

Participants were randomly assigned a booklet. They were instructed verbally and given written instructions, which asked them to read each of the two-line passages carefully and then complete the target sentence so that it made sense.

4.2.2. **Results**

Each participant’s response was scored as before. A total of 3.13% responses were discarded in further analyses. Of these, 0.69% referred to NP1 by a repeat name, 1.48% to NP2 and there were 0.95% ‘other’ responses. Responses referring to the pronoun were analysed using a 2 (covariation [consensus vs. distinctiveness]) by 2 (level [high vs. low]) by 2 (verb type [Stimulus-Experiencer vs. Experiencer-Stimulus]) ANOVA.

Because covariation information for the two verb types is different, covariation theory would predict a level by verb type interaction. For Experiencer-Stimulus verbs, low levels of information should lead to an increase in NP1 continuations compared to high levels; while for Stimulus-Experiencer verbs they should lead to a decrease in NP1 continuations compared to high levels. Independently of an effect of covariation, there ought to be a main effect of verb type, such that Stimulus-Experiencer verbs lead to more NP1 continuations and Experiencer-Stimulus verbs lead to fewer NP1 continuations.

4.3.1. Covariation and the implicit causality bias

As predicted, there was a main effect of verb type [ $F_1(1, 31) = 231.84$ ,  $MS_e = 1.12$ ,  $p < .0001$ ,  $F_2(1, 30) = 508.42$ ,  $MS_e = 1.72$ ,  $p < .0001$ ], such that Stimulus-Experiencer verbs led to more NP1 continuations ( $M = 6.84$ ) than did Experiencer-Stimulus verbs ( $M = 1.25$ ).

There was also the hypothesised level by verb type interaction [ $F_1(1, 31) = 10.53$ ,  $MS_e = 1.77$ ,  $p < .003$ ,  $F_2(1, 30) = 16.14$ ,  $MS_e = 1.46$ ,  $p < .0001$ ]. For Stimulus-Experiencer verbs low levels of information decreased NP1 continuations in comparison to high levels of information [ $t_1(1, 31) = 2.49$ ,  $p < .01$ ,  $t_2(1, 15) = 2.74$ ,  $p < .007$ ]. On the other hand, for Experiencer-Stimulus verbs low levels of information increased NP1 continuations in comparison to high levels of information [ $t_1(1, 31) = 3.19$ ,  $p < .002$ ,  $t_2(1, 15) = 2.94$ ,  $p < .005$ ] (see Figure 26).

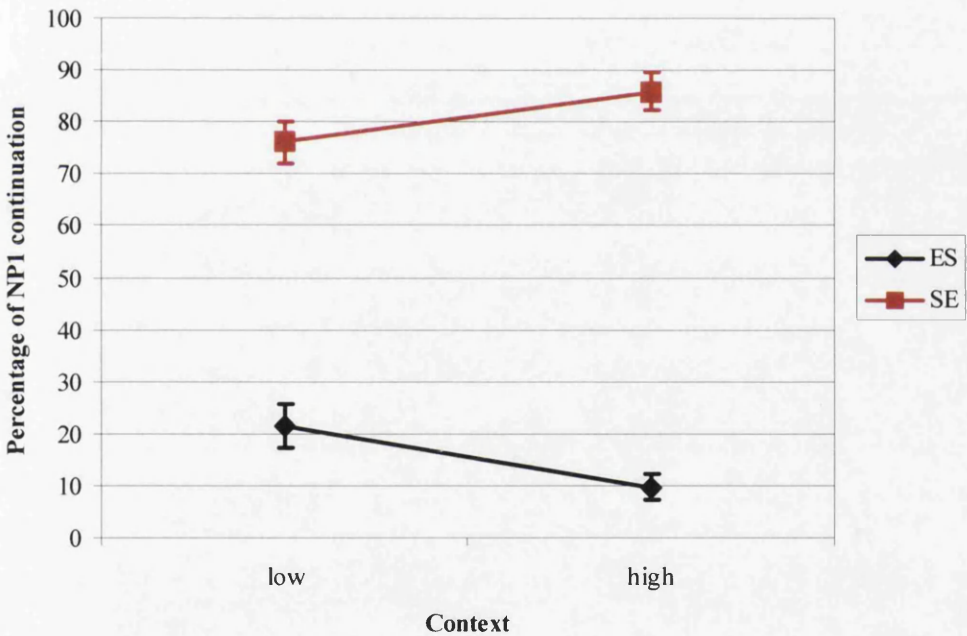


Figure 26: Interaction between verb type and level of information (ES = Experiencer-Stimulus verbs; SE = Stimulus-Experiencer verbs). For Experiencer-Stimulus verbs, low levels of information led to more NP1 continuations than did high levels of information. On the other hand, for Stimulus-Experiencer verbs high levels of information led to more NP1 continuations than did high levels of information.

There was also a verb by covariation interaction [ $F_1(1, 31) = 9.73$ ,  $MS_e = 1.51$ ,  $p < .004$ ,  $F_2(1, 30) = 5.09$ ,  $MS_e = 1.47$ ,  $p < .03$ ]. In post-hoc tests, all differences were reliable at  $p < .05$  (see Figure 27). Crucially, for Stimulus-Experiencer verbs



distinctiveness information led to more NP1 continuations than consensus information; whereas for Experiencer-Stimulus verbs, consensus information led to more NP1 continuations than distinctiveness information. By looking at Table 9, it can be seen that the crucial thing about the increase in NP1 continuations is whether covariation information quantifies over the first or the second noun phrase. For Stimulus-Experiencer verbs distinctiveness information quantifies over the first noun phrase; while for Experiencer-Stimulus verbs, consensus information quantifies over the first noun phrase. Therefore, a context sentence which presents information about the first noun phrase leads to more NP1 continuations. There were no other main effects or interactions.




Figure 27: Interaction between verb type and covariation information (ES = Experiencer-Stimulus verbs; SE = Stimulus-Experiencer verbs). For Experiencer-Stimulus verbs, consensus information led to more NP1 continuations than did distinctiveness information. On the other hand, for Stimulus-Experiencer verbs distinctiveness information led to more NP1 continuations than did consensus information.

4.3.2. Summary

The results of this experiment confirm the hypotheses. There was a strong effect of verb type. Stimulus-Experiencer verbs led to more NP1 continuations than Experiencer-Stimulus verbs. Furthermore, covariation information affected attributions. There were more continuations to NP1 when low covariation information was presented for Experiencer-Stimulus verbs than when high information was presented. On the

other hand, high covariation information for Stimulus-Experiencer verbs led to more NP1 continuations than low levels of information.

As well as the hypothesised effects, there was an unpredicted covariation by verb interaction. This showed that when the context sentence quantified over a particular noun phrase then there were more continuations to that noun phrase. So, if the covariation information quantified over NP1 then there were more continuations to NP1; while if it quantified over NP2 then there were more continuations to NP2. This could be accounted for by focus. By quantifying over a particular noun phrase, more attention may be drawn to it and this could lead to the observed difference in continuations.

To investigate further the relationship between covariation information and the implicit causality bias, a comparison was made between the context and the baseline in order to determine if context over-rides the bias. This would demonstrate whether context still over-rides the implicit causality bias when it is a relatively weak cue to causal attributions.

#### 4.3.3. *Does covariation over-ride the implicit causality bias?*

Each of the four context types, low consensus, low distinctiveness, high consensus and high distinctiveness were compared to the no-context baseline. As before, covariation theory predicts more continuations for NP1 when there is low consensus and low distinctiveness information compared to the baseline condition. However, it predicts fewer NP1 continuations when there is high consensus and high distinctiveness information compared to the baseline.

##### 4.3.3.1. *Stimulus-Experiencer verbs*

There were no reliable differences between context and baseline conditions on any measure, that is, for participant or item analyses (Figure 28).

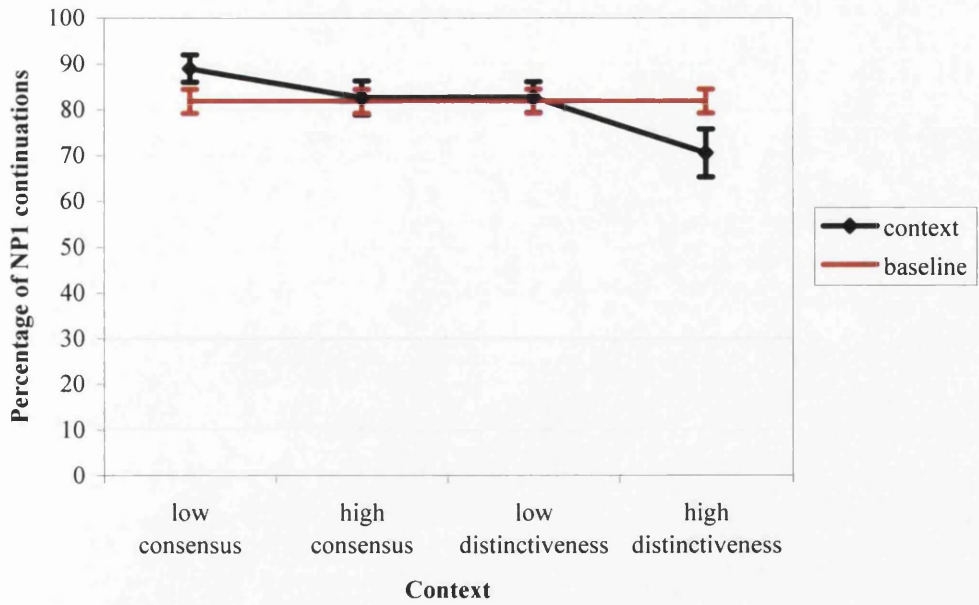


Figure 28: NP1 continuations for Stimulus-Experiencer verbs to both the target fragment with no context (baseline) and when presented with covariation information (context). The baseline is shown as a line only for convenience. There are no reliable differences in continuations between the two conditions.

#### 4.3.3.2. *Experiencer-Stimulus verbs*

Only low consensus information shifted continuations reliably from the baseline. It increased NP1 continuations, as is predicted by covariation theory [ $t_1(1, 54) = 3.62, p < .001$ ;  $t_2(1, 11) = 3.81, p < .002$ ]. However, low consensus information did not reverse the implicit causality bias of Experiencer-Stimulus verbs. This can be seen in Figure29.

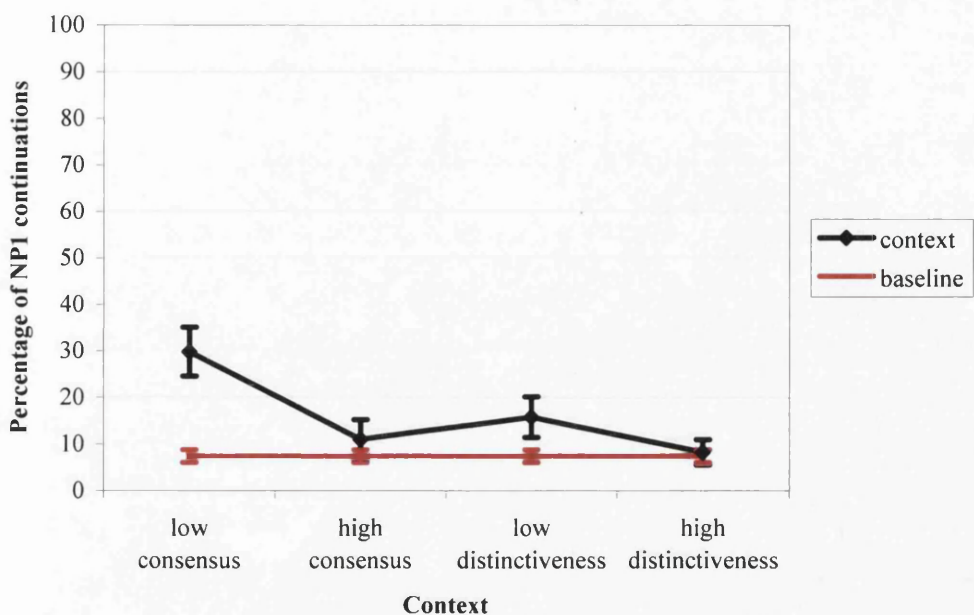


Figure 29: NP1 continuations for Experienter- Stimulus verbs to both the target fragment with no context (baseline) and when presented with covariation information (context). The baseline is shown as a line only for convenience. Low consensus information led to more NP1 continuations than the baseline condition.

#### 4.3.4. Summary

To summarise, when the bias is much stronger, as is the case for these state verbs, context information does not modify the bias to the same extent as it does with weaker biasing verbs. The only evidence for a shift was with Experienter-Stimulus verbs when paired with low consensus information. This is parallel to the shift produced by low consensus information for NP2 biasing verbs in the previous study. However, there were no reliable shifts in any of the other comparisons.

### 4.3. General Discussion

In this chapter, two experiments that compared the relative impact of implicit causality information and explicit covariation information were presented. Thus, two questions were asked: (1) does covariation information influence causal attributions in production and (2) does context information over-ride the implicit causality bias? Both of these questions can now be answered.

First, covariation information did indeed influence causal attributions. In both experiments, the predictions of covariation information were confirmed. In Experiment 4, the effect of covariation information on Agent-Patient verbs was examined. It was demonstrated that low levels of information led to more NP1 continuations while high levels of information led to fewer NP1 continuations, as was predicted.

In Experiment 5, the effect of covariation information on state verbs was examined. Once again, the results were consistent with the predictions of covariation theory. There was a level by verb type interaction, such that for Stimulus-Experiencer verbs there were more NP1 continuations to high than low levels of information; while for Experiencer-Stimulus verbs there were more NP1 continuations to low than high levels of information.

The answer to whether covariation information over-rides the implicit causality bias is more complicated. For Agent-Patient verbs, covariation information incongruent with the implicit causality bias shifts continuations from the baseline. That is, for NP1 biasing verbs, high consensus and high distinctiveness information shifted continuations from the baseline; while for NP2 biasing verbs low consensus and low distinctiveness information shifted continuations. However, there was a slightly different pattern associated with state verbs. For state verbs, only one piece of information shifted continuations. That is, for Experiencer-Stimulus verbs, the incongruent low consensus information shifted continuations from the baseline.

Out of all these cases, there were only three places where the context information completely over-rode the implicit causality bias. In Experiment 5, for NP1 biasing verbs high consensus and high distinctiveness information over-rode the implicit causality bias; while for NP2 biasing verbs low consensus information over-rode the bias. Covariation information over-rode the bias when context was a stronger cue than verb type. That is, while covariation information was the stronger cue to causal attributions for Agent-Patient verbs, context was a much weaker cue for state verbs. However, it is not clear from these two experiments whether the differences are due to the difference in strength of bias, or to type of verb type.

To summarise, it has been demonstrated that covariation information does influence causal attributions in production. This is consistent with the earlier studies (McArthur, 1972; Van Kleeck, Hillger & Brown, 1988; Rudolph, 1997) which demonstrated that when people are explicitly asked to make a causal inference, given covariation information and implicit causality verb information, both sources of information are used. However, it must be reiterated that the fact that both sources of

information are used does not mean that the implicit causality verb bias is reducible to covariation information. This remains an open question. What this chapter has demonstrated is that in a production situation, where people are not explicitly asked a causal question, people still make causal inferences that are influenced by verb information and covariation information. Another question that can be asked is if covariation information, also, influences comprehension. The next chapter attempts to address this question.

## CHAPTER 5:

# COVARIATION THEORY IN COMPREHENSION

### 5.0. Introduction

Chapter 4 demonstrated that during production participants made use of covariation information in order to make causal inferences. To recap on the basic findings, low consensus and low distinctiveness information led to relatively more NP1 attributions, while high consensus and high distinctiveness information led to fewer NP1 attributions. This was the case regardless of the type of verb. There was also an independent effect of verb type, with the most NP1 attributions given to NP1 biasing verbs, followed by neutral verbs, and finally NP2 biasing verbs. Given that both covariation information and verb type influenced causal inferences in production, another question that can be asked is whether these variables affect comprehension, too.

There are a number of studies that have examined the effect of verb type on comprehension (see also sections 1.1 and 3.0). In trying to assess the effect of the implicit causality bias in comprehension, participants are presented with a main clause, containing an implicit causality verb, followed by a subordinate clause which has an explicit cause referring to either NP1 or NP2. Examples are given below (taken from Stewart, Pickering & Sanford, 2000):

- (1) Daniel apologised to Joanne because he had been behaving selfishly.
- (2) Daniel apologised to Joanne because she didn't deserve the criticism.
- (3) Daniel congratulated Joanne because he was very impressed.
- (4) Daniel congratulated Joanne because she had won the championship.

The verb *apologised to* is a NP1 biasing verb. Hence, (1) contains an ending congruent with the verb's implicit causality bias; whereas (2) has an ending incongruent with the verb bias. On the other hand, *congratulated* is a NP2 biasing verb. In this case, it is (4) that presents an ending that is congruent with the verb bias; while (3) presents an ending incongruent with the bias. Overall, the evidence shows that there is a processing advantage for congruent endings. This is evident in naming tasks

(Caramazza, Grober, Garvey & Yates, 1977; Ehrlich, 1980), probe tasks (McDonald & MacWhinney, 1995; Garnham, Traxler, Oakhill & Gernsbacher, 1996; Greene & McKoon, 1995; Long & De Ley, 2000), as well as in self-paced reading (Garnham, Oakhill & Cruttenden, 1992; Stewart, Pickering & Sanford, 2000; Pynte, Kennedy, Murray & Courrieu, 1998).

In a naming task, a participant is presented with a sentence, such as (1) to (4) above. The participant is instructed to name the referent of the pronoun as quickly as possible. Both Caramazza, Grober, Garvey & Yates (1977) and Ehrlich (1980) used this method, in order to test whether readers use implicit causality information from the verb during comprehension. It is assumed that naming will be quicker for a congruent referent than an incongruent referent, due to the higher level of activation of the former type of antecedent. In both studies, this is what was found. Congruent referents were indeed named faster than incongruent referents.

The second method used to examine the causality congruency effect involves using a probe task. In this case, participants are presented with a sentence, either auditorily (as with McDonald & MacWhinney, 1995) or visually (as with Garnham *et al.*, 1996). At a critical point during the presentation of the sentence, a word is presented to the participant. When the target word is presented, he or she has to indicate whether that word has occurred in the sentence so far. Again, based on the assumption that the congruent referent has a higher level of activation, participants should respond to a probe referring to the congruent referent quicker than to a probe referring to an incongruent referent.

McDonald & MacWhinney (1995) and Garnham, Traxler, Oakhill & Gernsbacher (1996) find some evidence to support the causality congruency effect, although they do report differences as to exactly when in the sentence causality information is used by the processor. Overall, McDonald & MacWhinney find evidence in support of early use of causal information; while Garnham *et al.* find that it is not used until much later in the sentence. Further, Greene & McKoon (1995) presented participants with short passages, in which the last sentence was similar to the ones presented above. They found that there was a processing advantage for NP2 probes after NP2 biasing verbs. This is consistent with the causality congruency effect. However, they did not find a similar advantage for NP1 probes with NP1 biasing verbs (see Stewart *et al.*, 2000 for an explanation of this finding).

Recently, Long & De Ley (2000) have demonstrated that while skilled readers show an early effect of implicit causality, unskilled readers do not display a congruency



effect until the end of the sentence. This difference between skilled and unskilled readers could explain some of the discrepancies between the above mentioned studies. However, despite differences in when implicit causality information is used, all four studies show evidence consistent with the hypothesis that implicit causality information is used during comprehension.

Finally, the causality congruency effect can be assessed by investigating the reading process itself, and using reading time as an index of comprehension. If the implicit causality bias makes the processor prefer an ending congruent with the bias, then participants should be faster to read sentences like (1) and (4), which have endings congruent with the verb bias, as opposed to (2) and (3), where the endings are incongruent with the verb bias. Although some early studies failed to find a causality congruency effect in reading times (*e.g.* Garnham & Oakhill, 1985), other studies, which used more rigorous controls, have found a reliable advantage of congruent information (Garnham, Oakhill & Cruttenden, 1992; Stewart, *et al.*, 2000). For example, Garnham, Oakhill & Cruttenden (1992) found that when the explicit reason required a simple inference then there was a stronger congruity effect than when the explicit reason required a complex inference. Furthermore, when Stewart *et al.* (2000) controlled explicit reasons for a number of factors such as plausibility and length, they found strong congruity effects.

Pynte, Kennedy, Murray & Courrieu (1998) also demonstrated the causal congruency effect during reading, in an eye-tracking study. They presented participants with sentences in French, which included an implicit causality verb, and an explicit cause referring to either NP1 or NP2. They found that incongruent endings led to more fixations on the segment containing the implicit causality verb, more re-inspection of the causal referent, as well as more time spent at the sentence ending than the parallel congruent endings.

Therefore, the evidence, using a number of different methodologies, shows that people use the implicit causality in verbs during comprehension: congruent endings have a processing advantage over incongruent endings.

Given that this is the case, one could ask whether causal information is always used during comprehension. Specifically, does covariation information influence comprehension in the same way that the implicit causality bias does? That is, is there a processing advantage for an explicit reason that is congruent with covariation information?

Until relatively recently social psychologists have only been interested in *what kinds of information* people use when making attributions (*e.g.* covariation, causal mechanisms), and the *types of attributions* they make (*e.g.* dispositional, responsibility, instigation). There have been only a limited number of studies that have attempted to address the question of *when* people make such inferences (see Hamilton, 1988, for a brief review). There has not been an attempt to examine whether covariation information is used during the comprehension process. The following studies will attempt to test exactly this issue.

The same logic that is used to examine the effect of the implicit causality bias in production can be used in order to test whether covariation information is used in comprehension, also. If covariation information influences the comprehension process then explicit reasons that are congruent with covariation information should be easier to process than explicit reasons that are incongruent with covariation information. Consider (5) and (6) below:

- (5) Few people confided in Thomas. Debbie confided in Thomas because she was trusting.
- (6) Many people confided in Thomas. Debbie confided in Thomas because she was trusting.
- (7) Few people confided in Thomas. Debbie confided in Thomas because he was trustworthy.
- (8) Many people confided in Thomas. Debbie confided in Thomas because he was trustworthy.

In (5) low consensus information is provided in the sentence, *Few people confided in Thomas*. Low consensus information leads to attributions about NP1. On the other hand, (6) presents high consensus information, which leads to attributions about NP2. In both cases, the explicit reason, which is given in the subordinate clause of the second sentence is about NP1. Because low consensus information is congruent with NP1 causes but high consensus information is not, (5) should be quicker to read than (6). The same reasoning can be applied to (7) and (8). Both (7) and (8) provide explicit reasons about NP2. However, in (7) the covariation information leads to causal attributions about NP1; while in (8) it leads to causal attributions about NP2. Therefore, in this case (8) should be read faster than (7).

So, overall when covariation information and the explicit reason are congruent there should be faster reading times than when covariation information and the explicit reason are incongruent. Covariation information and explicit reasons are congruent with one another when there is low consensus information with a reason referring to NP1, or when there is high consensus information with a reason referring to NP2. On the other hand, the two things are incongruent when there is low consensus information and a reason referring to NP2, or when there is high consensus information and a reason referring to NP1. This can be seen in Table 10.

Covariation information	Referent of explicit reason	
	NP1	NP2
Low consensus	congruent	Incongruent
High consensus	incongruent	Congruent

Table 10: The table shows which information is congruent with which referent.

Of course, in discourse, covariation information cannot be presented without the accompaniment of a verb. And, as was seen from the review at the beginning of this chapter, the type of verb presented affects ease of comprehension. When an explicit cause is congruent with the bias of the verb then it is read faster than when it is incongruent with the verb. Hence, if the effect of covariation information on comprehension is to be examined then the influence of the verb will have to be controlled.

In order to test the effects of covariation information on comprehension weakly biasing verbs were chosen initially. It was hypothesised that the effects of covariation information would be most easily detected when the verb bias was weak, and therefore did not interfere with the contextual information. Hence, the Agent-Patient verbs for Experiment 4 were used.

Before examining the comprehension effects of covariation information, explicit reasons for the Agent-Patient verbs were constructed. Once these materials were constructed, Experiments 6 and 7 tested whether a causality congruency effect was observed. Neither experiment showed such an effect. Because there was no causality congruency effect, these materials were used in order to test whether there was a covariation congruency effect in comprehension. However, Experiment 8 did not show a reliable covariation congruency effect. Therefore, a fourth experiment was conducted,

which presented covariation information with strongly biasing verbs. Experiment 9 was more suggestive of an effect of covariation on comprehension.

## 5.1. Experiment 6:

In this experiment, participants were presented with a single sentence, which contained an implicit causality verb and explicit cause, which referred either to the first or second noun phrase of the main clause. The verbs were all Agent-Patient verbs, taken from the previous chapter (they were the same ones as those used in Experiment 4, section 4.1, see Appendix 3 for a full list). Three types of verbs were used: NP1 biasing (*e.g. defied*), neutral (*e.g. accepted*) and NP2 biasing (*e.g. protected*).

It was hypothesised that there would be a congruency effect such that for NP1 biasing verbs an explicit cause referring to NP1 would be quicker to read than an explicit cause referring to NP2. The opposite is predicted for NP2 biasing verbs. For these verbs, a cause referring to NP2 should be quicker to read than one referring to NP1. For neutral verbs although there is not a causal bias, which would lead to a preference for one argument over another, an advantage for NP1 is still expected. This advantage is hypothesised based on the first mentioned privilege (Gernsbacher, 1988).

### 5.1.1. Method

#### 5.1.1.1. *Participants*

Twenty undergraduates from the University of Glasgow participated in this experiment. All were native speakers of English.

#### 5.1.1.2. *Materials*

There were 36 experimental items (see Appendix 5). Each material consisted of an Agent-Patient verb, which was either NP1 biasing, neutral or NP2 biasing. This was followed by one of two endings, which referred to either the first or second noun phrase in the main clause. Hence, a 3 (verb bias [NP1 vs. neutral vs. NP2]) by 2 (explicit cause [NP1 vs. NP2]) design was used in this study. Referent of the ending was both a within participants and items factor while verb bias was a within participant, but between items factor. Two files were created so that each verb appeared with each ending over the course of the experiment. The items were presented in a different random order for each participant.

There were 108 fillers in the experiment, also. These were created in a similar manner to the experimental items. Each filler consisted of a main clause, a connective<sup>20</sup> and a subordinate clause, of comparable length to the experimental items.

The experimental items consisted of a main clause, which included a male and female name, as well as the implicit causality verb. Half of the items had a male name in the position of the first noun phrase, while the other half had a female name. This was followed by the connective *because* and the subordinate clause, which presented the explicit cause for the event. The explicit cause referred to either NP1 or to NP2. So, for example, for the verb *defied* participants were presented with either *Samantha defied Noel because she was rebellious*, where the cause refers to NP1; or they were presented with *Samantha defied Noel because he was authoritarian*, where the cause refers to NP2.

The explicit cause was always either a state or trait of the character. Further, it was ensured that the cause only ever differed by one word between conditions. This word was controlled for length, so that there was no reliable difference in length between a word referring to the first or second character [ $F(1, 33) = .411$ ,  $MS_e = .89$ ,  $p > .1$ ]. There was, also, no reliable differences of length based on verb bias [ $F(2, 33) = 1.17$ ,  $MS_e = 8.60$ ,  $p > .1$ ]. Nor was there any interaction between referent and verb bias [ $F(2, 33) = 1.6$ ,  $MS_e = .35$ ,  $p > .1$ ].

Similarly, the critical words were controlled for frequency<sup>21</sup>. There was not a reliable difference in frequency of occurrence for words referring to the first character or second character [ $F(1, 33) = .517$ ,  $MS_e = 133386$ ,  $p > .1$ ]. Neither was there a difference of frequency between the causes for the different verb biases [ $F(2, 33) = .625$ ,  $MS_e = 193564$ ,  $p > .1$ ]. Finally, there was no reliable interaction between referent and verb bias [ $F(2, 33) = .535$ ,  $MS_e = 137978$ ,  $p > .1$ ].

In order to ensure that the reason provided for the first character was as good as that provided for the second character a pre-test was conducted. Participants in the pre-test were asked to rate the reasons given for the target event, in order to ensure that the explicit causes to be used in the reading experiment were equivalent across conditions.

Twenty-four undergraduates rated the endings for how good a reason they provided for the target event. Participants were presented with each of the 36 materials in four conditions. The four conditions were: a good reason referring to NP1, a good

---

<sup>20</sup> None of the filler items, in any of the following experiments, used the causal connective *because*.

<sup>21</sup> using written word frequency, per 16.6 million, in Celex (Baayen, Piepenbrock, & Van Rijn, 1993).

reason referring to NP2, a bad reason referring to NP1 and a bad reason referring to NP2. The bad reasons were constructed by switching the referent of an ending from NP1 to NP2 and vice versa. So, for example, if (a) and (b) were constructed to be good reasons then (c) and (d) were the bad reasons.

- (a) Samantha defied Noel because she was rebellious
- (b) Samantha defied Noel because he was authoritarian
- (c) Samantha defied Noel because she was authoritarian
- (d) Samantha defied Noel because he was rebellious

Bad reasons were presented to ensure participants used the full range of the scale. Each item was followed by a scale ranging from 1 (bad reason) to 7 (good reason).

The items were divided into four blocks, with equal numbers of items from each condition in each block. The four blocks were randomised separately. They were then counterbalanced across the experiment.

Participants rated the target reasons as good reasons; that is, they rated them at least above the mid-point of the scale, which was 4. The mean rating for good reasons ( $M = 4.83$ ), which was reliably different from the bad reasons ( $M = 3.33$ ) [ $F(1, 33) = 105.8$ ,  $MS_e = 80.41$ ,  $p < .0001$ ].

The good reasons were further analysed to ensure that there was no difference in goodness of reason between referent or verb bias. Reasons for the first character were rated as highly as those for the second character [ $F(1, 33) = .026$ ,  $MS_e = .01$ ,  $p > .1$ ]. The verb bias comparison did not show a reliable difference [ $F(2, 33) = .52$ ,  $MS_e = 2.37$ ,  $p > .1$ ]. Finally, there was not an interaction of goodness of reason between referent and verb bias [ $F(2, 33) = 1.47$ ,  $MS_e = .67$ ,  $p > .1$ ].

Referent	Verb bias	Length	Frequency	Goodness
NP1	NP1 biasing	7.8 (.66)	439 (174)	4.7 (.37)
	Neutral	8.3 (.75)	150 (69)	4.6 (.24)
	NP2 biasing	6.9 (.63)	354 (118)	5.2 (.35)
NP2	NP1 biasing	7.7 (.51)	356 (112)	4.4 (.40)
	Neutral	7.8 (.62)	359 (174)	4.8 (.37)
	NP2 biasing	6.8 (.59)	487 (222)	5.3 (.36)

Table 11: Means (and  $MS_e$ ) for length, frequency and ratings for goodness of reason.

After the experimental item was presented, half of the items were followed by a forced choice comprehension question (*e.g.*, *Did Samantha defy Noel?*), where the participant had to answer ‘yes’ or ‘no’ by pressing the appropriate button. The correct answer for half of the questions was ‘yes’. Two versions of each question were created. If the correct answer to the question was ‘yes’ in the first list, then the correct answer was ‘no’ in the second. The answers to the questions were not used in any further analyses, but purely as a measure of comprehension.

#### 5.1.1.3. Procedure

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

Participants were randomly assigned to one of the lists. They were then provided with both verbal and written instructions. These were followed by four practice trials before they began the experiment proper.

Each trial began with a fixation spot on the left-hand side of the computer screen. The participant pressed the middle button of the button box to remove the fixation spot. It was then replaced with the main clause plus the causal connective *because*, *e.g.* *Samantha defied Noel because*. A second press of a button presented the subordinate clause *e.g.* *she was rebellious*. A third press of the same button replaced the sentence with a question, if there was one. Participants would answer the question by pressing either the left or right-button on the box. When there was no question the

message on the screen would read NO QUESTION at which point the participant would again press the middle button to bring an end to that trial.

Participants were allowed a break half way through the experiment for as long as they required.

5.1.2. Results

First, responses to the questions were checked for accuracy, in order to ensure that participants comprehended the sentences. All participants answered at least 89% of the questions correctly.

In order to test the experimental hypothesis 2 (referent [NP1 vs. NP2]) by 3 (verb bias [NP1 vs. neutral vs. NP2]) ANOVAs, treating both participants ( $F_1$ ) and items ( $F_2$ ) as random variables. The analyses were conducted on the reading times to the second fragment. Before analysing the data, the data was trimmed according to the definition of extreme values in the SPSS function explore. For each condition, the upper and lower quartiles of the distribution were calculated. Cut-off values were set at three times the interquartile range above the upper quartile, and three times the interquartile range below the lower quartile. Any data points that were above or below these cut-off values were replaced by the relevant cut-off value. The same procedure was used to replace outliers in all of the experiments reported in this chapter.

It was predicted that there would be a causality congruency effect, such that for NP1 biasing verbs, endings referring to NP1 would be read faster than endings referring to NP2; but for NP2 biasing verbs, endings referring to NP2 would be read faster than those referring to NP1. However, there was no evidence of an interaction between explicit cause and verb bias [ $F_1$  (2, 38) = 1.87,  $MS_e$  = 34175,  $p$  > .1;  $F_2$  (2, 33) = .75,  $MS_e$  = 65465,  $p$  > .1] (see Table 12). Indeed, the data shows a trend in the opposite direction.

Referent	Verb bias		
	NP1 biasing	Neutral	NP2 biasing
NP1	1319	1318	1231
NP2	1228	1401	1269
Total	1273	1359	1250

Table 12: Mean reading times for second fragment.



There was a suggestion of a main effect of verb bias, although it was not reliable by items [ $F_1(2, 38) = 5.11$ ,  $MS_e = 21723$ ,  $p < .01$ ;  $F_2(2, 33) = 1.61$ ,  $MS_e = 49679$ ,  $p > .1$ ]. Endings for neutral verbs took longer to read than those for NP1 biasing verbs [ $t_1(1, 19) = 2.07$ ,  $p < .03$ ;  $t_2(1, 22) = 1.33$ ,  $p < .09$ ], or for NP2 biasing verbs [ $t_1(1, 19) = 3.50$ ,  $p < .001$ ;  $t_2(1, 22) = 1.72$ ,  $p < .05$ ].

In the previous chapter, it was hypothesised that neutral verbs were neutral because both arguments of the verb were equally available. If this is the case, then one would expect that causes referring to both NP1 and NP2 would take longer to read than when there is a biasing verb.

As was described above, it is known that there is a processing advantage for a cause referring to a noun phrase that is the implicit cause of the target verb. This is manifest as shorter reading times. On the other hand, there is a processing disadvantage for a noun phrase that is not the implicit cause of the target verb. If neutral verbs have both arguments equally available, then there should be a disadvantage when reading a fragment referring to either NP1 or NP2. This is because in both cases the explicit cause is incongruent with one of the arguments from the implicit cause, which would lead to competition between the two noun phrases. Because there is a disadvantage when referring to both NP1 and NP2 then overall there ought to be longer reading times for sentences with neutral verbs.

However, there remains the possibility that the endings for neutral verbs took longer due to some property of the verb itself. For example, it could be that the neutral verbs in this study are more infrequent than the biasing verbs. It is known that words that are infrequent take longer to read than words which are frequent (*e.g.* Inhoff & Rayner, 1986). However, there was not a reliable difference in frequency between the three categories of verb, *i.e.* between NP1 biasing verbs ( $M = 760$ ), NP2 biasing verbs ( $M = 1442$ ) or neutral verbs ( $M = 833$ ) [ $F(2, 22) = 1.34$ ,  $MS_e = 1263887$ ,  $p > .3$ ].

If neutral verbs differ in some other way from biasing verbs then neutral items would take longer to read than the biasing ones. If some property of the verb, other than the causal weighting of the arguments, is responsible for the increased reading times for neutral verbs, then this should be seen in first fragment reading times also. The reading time for neutral verbs should be longer than the reading times for NP1 and NP2 biasing verbs. However, this is not the case [ $F_1(2, 38) = 1.99$ ,  $MS_e = 41849$ ,  $p > .1$ ;  $F_2(2, 33) = .94$ ,  $MS_e = 48674$ ,  $p > .1$ ], with NP1 biasing verbs taking around 1908 ( $MS_e = 103$ )

milliseconds to read, neutral verbs 1780 ( $MS_e = 108$ ) and NP2 biasing verbs 1826 ( $MS_e = 105$ ) milliseconds.

### 5.1.3. Discussion

In this experiment, there was no reliable causality congruency effect, contrary to the hypothesis. This may be due to a number of reasons. Previous experiments have failed to find a causality congruency effect when the explicit causes were not controlled for a number of factors, such as plausibility, length and ease of inference. However, in this study such factors were taken into consideration. One possible difference between this experiment and others is that the verbs in this study do not have a very strong bias, and a strong bias may be necessary for the causality congruency effect.

There was a suggestion of a main effect of verb type, such that longer reading times for sentences with neutral verbs were observed than for sentences with biasing verbs. The longer reading times were not observed in the first part of the sentence, where the verb is presented. Rather, the longer reading times are observed when participants read the explicit cause in the second part of the sentence. This is not due to differences in the endings because they have been controlled for length, frequency, and goodness of reason. One explanation of the longer reading times could be that both arguments of neutral verbs are equally available. Therefore, whichever ending is read it is incompatible with at least one of the arguments, which leads to the longer times regardless of to whom the ending refers.

To ensure that this main effect was not an artefact of the presentation method, the same experiment was run a second time, but this time using a different presentation method. Instead of presenting the target sentence in two fragments, Experiment 7 presented the complete sentence.

## 5.2. Experiment 7:

In this experiment, participants were presented with the same materials as in Experiment 6, but this time with a whole sentence presentation. As in Experiment 6 a causality congruency effect was predicted, such that for NP1 biasing verbs an explicit cause referring to NP1 would be quicker to read than an explicit cause referring to NP2. The opposite is predicted for NP2 biasing verbs. For these verbs, a cause referring to NP2 should be quicker to read than one referring to NP1. Furthermore, a main effect of

verb type was predicted, such that items with neutral verbs would take longer to read than items with biasing verbs.

### 5.2.1. Method

#### 5.2.1.1. *Participants*

Twenty undergraduates from the University of Glasgow participated in this experiment. All were native speakers of English. None of the participants had taken part in the previous study.

#### 5.2.1.2. *Materials and Procedure*

The materials and procedure were identical to the previous experiment, with one exception. The target sentence was presented as a whole sentence instead of in two fragments. The participant was presented with a fixation spot on the left-hand side of the screen. When the participant pressed a button, the target sentence was displayed. A second press of the same button presented either a question or the participant was told there was no question.

### 5.2.2. Results

All participants answered at least 85% of the questions correctly. Hence, the reading times for the target fragment were subjected to a 2 (referent [NP1 vs. NP2]) by 3 (verb bias [NP1 vs. neutral vs. NP2]) ANOVA.

If there was a causality congruency effect, then there would be an interaction between verb type and explicit cause. However, as in Experiment 6 there was no such reliable interaction [ $F_1(2, 38) = .78$ ,  $MS_e = 239022$ ,  $p > .1$ ;  $F_2(2, 33) = .68$ ,  $MS_e = 247064$ ,  $p > .1$ ] (see Table 13). More specifically, if there was a causality congruency effect then for NP1 biasing verbs there would be an advantage for an explicit cause referring to NP1 over one referring to NP2; while, for NP2 biasing verbs there would be an advantage for an explicit cause referring to NP2 over one referring to NP1. Although there did appear to be some indication that this was the case for NP1 biasing verbs, this was not the case for NP2 biasing verbs. Therefore, there was no causality congruency effect in this experiment.

There was a suggestion of a main effect of verb bias, but this was only reliable by participants [ $F_1(2, 38) = 3.49$ ,  $MS_e = 235965$ ,  $p < .04$ ;  $F_2(2, 33) = 1.86$ ,  $MS_e =$

315067,  $p > .1$ ]. It appeared that sentences with neutral verbs took slightly longer to read than biasing verbs. Items with neutral verbs took longer to read than those with NP2 biasing verbs [ $t_1(1, 19) = 2.00, p < .03$ ;  $t_2(1, 22) = 1.42, p < .09$ ], although, unlike Experiment 6, there was not a reliable difference between neutral verbs and NP1 biasing verbs [ $t_1(1, 19) = .10, p > .1$ ;  $t_2(1, 22) = .59, p > .1$ ].

Explicit Cause	Verb bias		
	NP1 biasing	Neutral	NP2 biasing
NP1	3720	3650	3532
NP2	3836	3926	3537
Total	3778	3788	3535

Table 13: Mean reading times for target sentence.

### 5.2.3. Discussion

Experiment 7 showed a similar pattern of results to Experiment 6. There was no evidence of a causality congruency effect, as outlined in the section 5.1.3. This could be due to a number of factors. There was a main effect of verb bias. However, in this case although the means were in the same direction as in Experiment 6, with the longest reading times for neutral items, the difference between NP1 biasing verbs and neutral verbs was not a reliable one. Hence, although there is some intriguing evidence to suggest that neutral verbs may have both arguments equally available, the evidence is only suggestive.

To summarise, the two studies show that verb information is not being used as a determinant of causal attributions. This made these materials suitable to test whether covariation information is used during comprehension.

## 5.3. Experiment 8:

This experiment investigated whether people use covariation information during comprehension. If covariation information is used during comprehension then there ought to be an interaction between the covariation information and the explicit cause. Low consensus information leads to NP1 attributions, and therefore items with low consensus information should lead to faster reading times for explicit causes referring to

NP1 then those referring to NP2. On the other hand, high consensus information should lead to faster reading times for explicit causes referring to NP2 than those referring to NP1.

In order to test for a covariation congruency effect, participants were presented with one piece of covariation information, either low consensus information (e.g. *Few people defied Noel*), or high consensus information (e.g. *Many people defied Noel*). The covariation information was followed by another sentence, which was of the same form as the previous two experiments (e.g. *Samantha defied Noel because she was rebellious*). As with the above experiments, the second sentence provided an explicit reason for the event that was being described. Therefore, the reading times for this second sentence can be measured in order to test the influence, if any, of the consensus information presented.

### 5.3.1. Method

#### 5.3.1.1. Participants

Forty undergraduates from the University of Glasgow participated in this experiment.<sup>22</sup> None had participated in either of the previous experiments. All were native speakers of English.

#### 5.3.1.2. Materials and Procedure

The materials were the same as in Experiments 6 and 7, with one modification: the addition of a context sentence.

A trial began with the fixation spot, as before. This was followed by a context sentence which presented low consensus information (e.g. *Few people defied Noel*) or high consensus information (e.g. *Many people defied Noel*). Four quantifiers were used to provide variation in the context sentence, with *few* and *very few* being used to denote low consensus information, while *many* and *most* were used for high consensus information.

The context sentence was followed by the target sentence. This was presented in two fragments, as in Experiment 6. A button press after the context sentence presented the first fragment (e.g. *Samantha defied Noel because*). Another button press presented the target fragment, which provided an explicit reason for the event having occurred.

---

<sup>22</sup> Two participants had to be replaced because of technical problems.

The reason referred to either NP1 (*e.g. she was rebellious*); or NP2 (*e.g. he was authoritarian*). On half the trials, the target fragment was followed by a question.

5.3.2. Results

All participants answered at least 80% of the questions correctly, with an overall 89% of the questions being answered correctly. The reading times from the second clause were subjected to a 2 (consensus [low vs. high]) by 3 (verb bias [NP1 vs. neutral vs. NP2]) by 2 (explicit cause [NP1 vs. NP2]) ANOVA.

First, a consensus by explicit cause interaction was predicted. However, there was no such reliable interaction, although the means were in the predicted direction [ $F_1(1, 39) = 1.34, MS_e = 47664, p > .1$ ;  $F_2(1, 33) = .77, MS_e = 16592, p > .1$ ] (see Figure 30). A power analysis based on the participant analysis gave a power of .20 to detect an interaction of consensus and explicit cause, when the effect size was small (Cohen, 1977). This means that a sample size of 1160 would be needed to detect an effect when  $\alpha = .05$ .

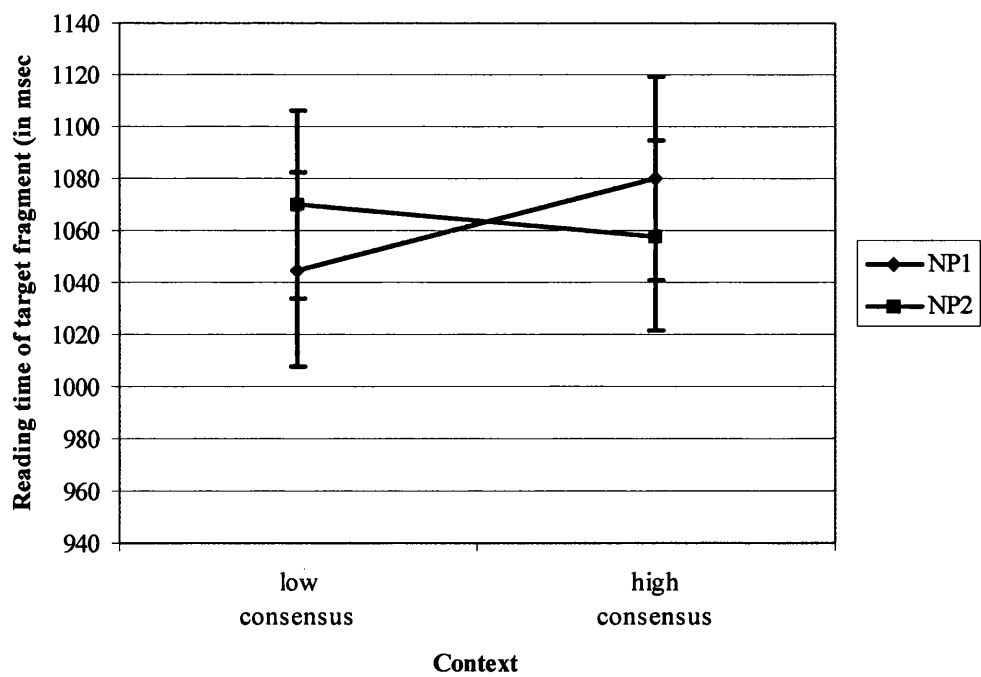


Figure 30: The above graph plots the interaction between context and referent ending (error bars show  $MS_e$ ).

As before there was no reliable verb bias by explicit cause interaction [ $F_1(2, 78) = 2.15$ ,  $MS_e = 48438$ ,  $p > .1$ ;  $F_2(2, 33) = 1.21$ ,  $MS_e = 23656$ ,  $p > .1$ ], although, the means were in the predicted direction. That is, endings referring to NP1 were read faster than endings referring to NP2 for NP1 biasing verbs, and vice versa for NP2 biasing verbs (see Table 14).

Referent	Verb bias		
	NP1 biasing	Neutral	NP2 biasing
NP1	1050	1127	1011
NP2	1110	1106	976
Total	1080	1117	994

Table 14: Mean reading times for second fragment.

Once again there was a main effect of verb bias [ $F_1(2, 78) = 11.96$ ,  $MS_e = 53382$ ,  $p < .0001$ ,  $F_2(2, 33) = 2.87$ ,  $MS_e = 70198$ ,  $p < .07$ ]. Explicit causes for neutral verbs took longer to read than for NP1 biasing verbs [ $t_1(1, 39) = 1.70$ ,  $p < .05$ ;  $t_2(1, 22) = .61$ ,  $p > .1$ ], or for NP2 biasing verbs [ $t_1(1, 39) = 4.70$ ,  $p < .0001$ ;  $t_2(1, 22) = 2.46$ ,  $p < .01$ ] (see Table 14 for means). This is consistent with Experiments 6 and 7. There were no other main effects or interactions.

### 5.3.3. Discussion

In this experiment, an attempt was made to test whether covariation information is used during comprehension. It was hypothesised that there would be a consensus by explicit cause interaction such that an explicit cause referring to NP1 would be read faster when paired with low consensus information; whereas for high consensus information an explicit cause referring to NP2 would be read faster. However, there was no such reliable interaction. Further, a power analysis of the number of participants needed for such an interaction suggested that this is not a fruitful route to take in order to determine the effect of covariation on comprehension.

Up until now, it has been assumed that if an effect of covariation information were to be observed in comprehension, then it would be observed with weakly biasing verbs. With such verbs, there would be less competition between the different sources of causal information. However, this does not appear to be the case. So, in Experiment

9, covariation information was paired with strongly biasing implicit causality verbs. It was assumed that by encouraging participants to make causal attributions, with cues from the context sentence, the verb itself, and the causal connective *because*, an effect of covariation information may be more apparent. On top of this, participants were asked questions after each item, again encouraging a deep reading of the materials.

## **5.4. Experiment 9:**

In this experiment, participants were presented with covariation information with strongly biasing verbs. It was hypothesised that there would be a covariation by explicit cause interaction. Covariation information was in the form of low or high consensus information, as before. Low consensus information should lead to faster reading times for an explicit cause referring to NP1 compared to one referring to NP2; while high consensus information should lead to faster reading times for an explicit cause referring to NP2 than NP1. Furthermore, as strongly biasing verbs are being used, a causal congruency effect was hypothesised, such that explicit causes referring to the implicit cause should be read faster. That is, explicit causes referring to NP1 should be read faster than those referring to NP2 for NP1 biasing verbs, and vice versa for NP2 biasing verbs.

### **5.4.1. Method**

#### *5.4.1.1. Participants*

Forty undergraduates from the University of Glasgow participated in this experiment. All were native speakers of English. None had participated in any of the previous experiments.

#### *5.4.1.2. Materials*

The items were constructed in the same manner as Experiment 8. There was a context sentence followed by a target sentence that introduced two protagonists and an explicit cause. The target sentences were taken from Stewart, Pickering & Sanford (2000). A full set of the materials can be seen in the appendix (Appendix 6). These materials were chosen because they had been rigorously controlled on a number of variables. The verbs were all pre-tested, so only those verbs which were strongly biasing were used. For each target sentence, the explicit cause referred to either NP1 or



NP2, just as in Experiments 6-8. (The bias for the verbs was estimated by taking the difference scores for each item. For NP1 biasing verbs the mean difference scores was 19.1; while for NP2 biasing verbs it was -16.3.)

Each explicit cause was controlled for length and plausibility (for further details see Stewart, Pickering & Sanford, 2000). The overall design was a 2 (consensus [low vs. high]) by 2 (verb bias [NP1 vs. NP2]) by 2 (explicit cause [NP1 vs. NP2]).

Each item consisted of a context sentence, which presented either high or low consensus information. This was followed by the target sentence, which consisted of a main clause presenting the implicit causality verb and a subordinate clause that presented an explicit cause referring to either NP1 or NP2. As before, the two clauses were connected by *because*.

There were 24 implicit causality verbs. Each verb was presented twice in order to increase the overall power of the experiment. If on first presentation, a verb appeared with low consensus information and an ending referring to NP1; then on the second presentation it was paired with high consensus information and an ending referring to NP2. To further minimise the similarities of the items different names were paired with each presentation of the verb. Further, the two occurrences of a verb were presented in two separate blocks, which were randomised independently of one another and were separated by a break.

There were 144 filler items, so as before, the experimental items only constituted a quarter of the materials. The fillers consisted of a context sentence, and a target sentence consisting of a main clause, connective and subordinate clause. A third of the context sentences for the filler items presented a quantified statement (*e.g. Many people were at Ascot. Tina was wearing a large hat and was spotted by Yvonne.*).

To encourage integration, a question followed each item. The question was a forced choice comprehension question. So, for example, following the item *Few people defied Noel. Samantha defied Noel because she was rebellious*, the question would be *Who was rebellious?* where the participant had to answer *Samantha* or *Noel* by pressing the appropriate button. The two possible answers were presented simultaneously with the question. One answer appeared on the left-hand side of the screen; the other on the right-hand side. Participants were told to press the left-hand button of the button box if they thought the correct answer was on the left-hand side of the screen, and the right button if they thought the correct answer was on the right-hand side of the screen. The position of the correct answer was counterbalanced across files.

The questions for the experimental items made reference to the subordinate clause and required resolution of the pronoun. For the filler items that presented a quantified context sentence a question about the quantifier was asked (*e.g. How many people were at Ascot? few/many*). This encouraged participants to attend to the context information. The remaining fillers had questions referring to a mixture of the context sentence, main clause and subordinate clause.

#### 5.4.1.3. Procedure

The experiment began with a fixation spot on the left-hand side of the screen. On pressing of the appropriate button, the spot was replaced by the context sentence. For the experimental items, the context sentence presented low or high consensus information. A second press of the button replaced the context sentence with the target sentence. The target sentence consisted of a main clause, which contained the implicit causality verb, plus the subordinate clause, which presented the explicit cause. A third press of the button presented the question along with two answers – one of which was the correct one. If the participant thought the correct answer was on the left-hand side of the screen then they had to press the left most button in front of them; if they thought it was on the right-hand side of the screen they had to press the right most button. When the participant had answered the question, they were asked to press the yellow button, which resulted in the fixation spot being presented again.

#### 5.4.2. Results

Because the questions for the experimental items in this study were comprehension questions which required the resolution of the pronoun in the subordinate clause, the question answering data could also be analysed. Hence, there were three dependent variables: firstly, the time taken to read the target sentence; secondly, the question answering accuracy; and finally, the time taken to answer the questions. Each of these will be discussed in turn.

##### 5.4.2.1. Reading time for target sentence

There was a main effect of explicit cause [ $F_1(1, 39) = 25.86$ ,  $MS_e = 236769$ ,  $p < .0001$ ;  $F_2(1, 22) = 15.54$ ,  $MS_e = 118193$ ,  $p < .001$ ]. Sentences containing a subordinate clause referring to NP1 ( $M = 3014$ ) were read faster than those referring to NP2 ( $M =$

3317) by around 300 milliseconds. This can be interpreted as the first mentioned privilege (Gernsbacher, 1988).

There was also a reliable effect of verb bias [ $F_1(1, 39) = 35.48, MS_e = 203374, p < .0001; F_2(1, 22) = 6.57, MS_e = 329428, p < .02$ ]. When the target sentence contained a NP2 biasing verb ( $M = 3029$ ) it was read around 300 milliseconds faster than when it contained an NP1 biasing verb ( $M = 3329$ ). This was also found by Stewart, *et al.* (2000).

Thirdly, there was a verb bias by explicit cause interaction [ $F_1(1, 39) = 26.11, MS_e = 274777, p < .0001; F_2(1, 22) = 18.21, MS_e = 118193, p < .0001$ ] (see Figure 31).




Figure 31: The above graph shows an interaction between the implicit causality bias and the explicit cause. The explicit cause can refer to either NP1 (black) or NP2 (red). There is a penalty seen in reading time when the explicit cause refers to NP2 when the implicit cause is NP1.

This can be interpreted as the causality congruency effect, where there is an advantage for reading an explicit cause that is congruent with the implicit cause of the verb. However, this interaction is actually due to a penalty in reading time following an ending referring to NP2 when paired with an NP1 biasing verb. Stewart *et al.* do not report such an interaction, as they were always interested in a main effect of congruency. Hence, they reported a main effect of congruency, such that congruent endings were read faster than incongruent endings.

Finally, a consensus by explicit cause interaction was hypothesised. However, there was no such reliable interaction [ $F_1(1, 39) = 1.98, MS_e = 238809, p > .1$ ;  $F_2(1, 22) = 1.21, MS_e = 117125, p > .1$ ] (see Figure 32). Covariation theory predicts that for low consensus information explicit causes referring to NP1 should be read faster than those referring to NP2 should. A simple main effect analysis showed that this was indeed the case [ $t_1(1, 39) = 5.08, p < .0001$ ;  $t_2(1, 23) = 3.35, p < .001$ ]. Covariation theory also predicts that the opposite should be true for high consensus information. However, NP1 endings were read faster than NP2 endings for high consensus information also [ $t_1(1, 39) = 3.62, p < .001$ ;  $t_2(1, 22) = 1.59, p > .1$ ]. This means that there is a general advantage for explicit causes referring to NP1. This will be discussed after the question answering data is reported.




Figure 32: The above graph shows how context information affects the reading of the explicit cause. This shows that there is an advantage for explicit causes referring to NP1, regardless of the type of contextual information.

#### 5.4.2.2. Question response accuracy

Correct responses to the questions were analysed. Once again, a 2 (consensus [low vs. high]) by 2 (verb bias [NP1 vs. NP2]) by 2 (explicit cause [NP1 vs. NP2]) ANOVA was conducted. This analysis only showed a main effect of referent, and even then, it was only reliable by participants [ $F_1(1, 39) = 13.23, MS_e = .051, p < .001$ ;  $F_2(1,$

22) = 1.69,  $MS_e = .012$ ,  $p > .1$ ]. This showed that participants were more accurate in answering questions referring to NP1 (see Table 15).

Context	Verb bias	Explicit cause	
		NP1	NP2
Low consensus	NP1	96.3%	92.9%
	NP2	95.4%	92.5%
High consensus	NP1	96.3%	92.9%
	NP2	95.0%	92.9%

Table 15: Percentage of correct answers in each condition

#### 5.4.2.3. Question response time

Finally, the time taken to answer the questions was analysed, in order to determine whether there was a late effect of covariation information. Overall, the question response data showed the same pattern as the reading data. There was a main effect of explicit cause [ $F_1(1, 39) = 28.88$ ,  $MS_e = 90776$ ,  $p < .0001$ ;  $F_2(1, 22) = 14.79$ ,  $MS_e = 52427$ ,  $p < .001$ ], as there was of verb bias [ $F_1(1, 39) = 13.73$ ,  $MS_e = 81391$ ,  $p < .001$ ;  $F_2(1, 22) = 5.43$ ,  $MS_e = 84729$ ,  $p < .03$ ]. In both cases, the means were in the same direction as in the reading time data. Explicit causes referring to NP1 ( $M = 1618$ ) were read faster than those referring to NP2 ( $M = 1743$ ) were. On the other hand, sentences with NP2 biasing verbs ( $M = 1649$ ) were read faster than those with NP1 biasing verbs ( $M = 1712$ ) were.

Additionally, there was a main effect of consensus [ $F_1(1, 39) = 11.69$ ,  $MS_e = 49150$ ,  $p < .001$ ;  $F_2(1, 22) = 9.08$ ,  $MS_e = 246581$ ,  $p < .006$ ]. The questions were answered quicker when the context sentence presented low consensus information ( $M = 1680$ ) as opposed to high consensus information ( $M = 1765$ ).

There was also evidence of a causality congruency effect, as seen by the verb bias by explicit cause interaction [ $F_1(1, 39) = 21.87$ ,  $MS_e = 82237$ ,  $p < .0001$ ;  $F_2(1, 22) = 11.37$ ,  $MS_e = 52428$ ,  $p < .003$ ]. However, this was modified by a three-way interaction between consensus, referent and verb bias [ $F_1(1, 39) = 7.06$ ,  $MS_e = 81815$ ,  $p < .01$ ;  $F_2(1, 22) = 2.99$ ,  $MS_e = 63771$ ,  $p < .09$ ]. Because the question of interest was whether there was an effect of covariation, this three-way interaction was broken down into two, two-way interactions. The first two-way interaction examined NP1 biasing



verbs, and the second NP2 biasing verbs. This way the covariation effects could be examined for each verb type.

For NP1 biasing verbs, there was a main effect of referent [ $F_1(1, 39) = 38.47$ ,  $MS_e = 113893$ ,  $p < .0001$ ;  $F_2(1, 22) = 23.36$ ,  $MS_e = 58404$ ,  $p < .001$ ] and a main effect of consensus [ $F_1(1, 39) = 17.55$ ,  $MS_e = 75655$ ,  $p < .0001$ ;  $F_2(1, 22) = 20.46$ ,  $MS_e = 20127$ ,  $p < .001$ ].

Covariation theory predicts a consensus by explicit cause interaction, such that low consensus information should lead to quicker response times with a NP1 explicit cause than a NP2 explicit cause; whereas high consensus information should lead to quicker response times to a NP2 explicit cause than a NP1 explicit cause. There was, indeed, a consensus by referent interaction, but this was reliable only by participants [ $F_1(1, 39) = 8.37$ ,  $MS_e = 99315$ ,  $p < .006$ ;  $F_2(1, 22) = 1.79$ ,  $MS_e = 87739$ ,  $p > .1$ ].




Figure 33: Interaction of consensus and explicit cause for NP1 biasing verbs

Low consensus information did indeed lead to quicker response times to NP1 than NP2 endings [ $t_1(1, 39) = 6.17$ ,  $p < .0001$ ;  $t_2(1, 11) = 4.30$ ,  $p < .0005$ ]. However, for high consensus information NP2 explicit causes did not lead to quicker responses than NP1 explicit causes. In fact, NP1 explicit causes were quicker again [ $t_1(1, 39) = 2.72$ ,  $p < .01$ ;  $t_2(1, 11) = 1.80$ ,  $p < .09$ ] (see Figure 33).

There were no reliable effects for NP2 biasing verbs (see Figure 34). This was true for the main effect of explicit cause [ $F_1(1, 39) = .65$ ,  $MS_e = 59120$ ,  $p > .1$ ;  $F_2(1, 22) = .12$ ,  $MS_e = 49118$ ,  $p > .1$ ], as well as for consensus information [ $F_1(1, 39) = .12$ ,

$MS_e = 55128, p > .1; F_2(1, 22) = .17, MS_e = 19055, p > .1$ ]. Finally, there was not a reliable consensus by explicit cause interaction [ $F_1(1, 39) = .71, MS_e = 37668, p > .1; F_2(1, 22) = .33, MS_e = 39666, p > .1$ ].




Figure 34: Interaction of consensus and explicit cause for NP2 biasing verbs

5.4.3. Discussion

In this experiment, participants were encouraged to make causal attributions through presentation of numerous causal cues (strongly biased implicit causality verbs, explicit covariation information, and the presence of a causal connective). Furthermore, participants had to attend to the information because a question was asked with every trial. However, despite this, there is not a straightforward effect of covariation information.

In both the reading time data from the target sentence and the time taken to answer questions, there was evidence of a consensus by explicit cause interaction. Covariation theory predicts a crossover interaction between consensus information and the explicit cause: the covariation congruency effect. There should be an advantage for reading an explicit cause referring to NP1 with low consensus information; whereas, there should be an advantage for NP2 information with high consensus information. However, in this experiment a covariation congruency effect was not found. Although there was an advantage for explicit causes referring to NP1 for low consensus

information as predicted, the same advantage was seen for high consensus information. This was the case for both the reading data and the question answering data.

So, despite the fact that these findings are very suggestive for the hypothesis that covariation information is used during comprehension, a second rival theory has to be ruled out. It could be the case that the above interaction is not due to the covariation information per se, but the result of more general discourse processes. The context sentences all present consensus information, which quantify over the first noun phrase, *i.e. Quantifier people apologised to Joanne*. The fact that this is always the case may focus attention more on the subject noun phrase. If this were true, then one would expect an overall advantage for NP1. This is, indeed, what was found. Regardless of whether there was low or high consensus information, explicit causes referring to NP1 were at an advantage. The fact that there was an interaction between the two types of consensus information could be due to the fact that the quantifier *few* focused attention even more than *many* because the former quantifier is marked. Therefore, there could be an effect of context that is not sensitive to the covariation information, but other linguistic factors.

In order to disentangle these accounts, a second experiment would need to be undertaken. Instead of presenting consensus information, participants would be presented with distinctiveness information. For example, participants would be presented with either low distinctiveness information such as *Daniel apologised to many people*, or high distinctiveness information such as *Daniel apologised to few people*. Distinctiveness information quantifies over the second noun phrase. Already one can see that if it is linguistic factors which gave rise to the effects in the current experiment then the same effects would not be seen with distinctiveness information. Furthermore, covariation theory predicts that there would be faster reading times for explicit causes referring to NP1 than NP2, with low distinctiveness information. However, if markedness of the quantifier is affecting reading time then one would expect faster reading times for NP1 with high distinctiveness information.

## 5.5. General discussion

In this chapter, an attempt was made to test whether covariation information is used during comprehension. In an experiment with weakly biasing verbs (Experiment 7), there was a small and non-significant effect of covariation on reading. On closer examination of the data, it was seen that a rather large number of participants would be required to make this effect a reliable one, should there be one. In Experiment 8, a



number of cues were used to maximise the possibility of participants using covariation information during comprehension. Covariation theory predicts that there should be a covariation congruency effect. This was not the case. Instead, there was an overall advantage for explicit causes referring to NP1 regardless of the context information. There was some evidence to indicate that the effect was stronger for low consensus information. This finding is consistent with covariation theory, which predicts that low consensus information should lead to NP1 attributions. However, it is not clear whether it is the covariation information that is giving rise to a preference for NP1, or whether it is a combination of linguistic factors, such as the noun phrase that is being quantified over and the markedness of the quantifier.

Covariation theory assumes that natural language quantifiers are providing set size information, and it is this set size information that is being used to make causal attributions. However, from the above discussion it is apparent that the way that covariation information is presented is confounded with a number of linguistic factors. Chapters 6 and 7 take a closer examination at the assumptions underlying covariation theory and examine the effect of presentation of covariation information on attributions.

# CHAPTER 6:

## COVARIATION THEORY AND NATURAL LANGUAGE

### QUANTIFIERS

#### 6.0. Introduction

In this chapter, the assumptions underlying covariation theory will be examined more closely. Chapter 4 clearly demonstrated that covariation information does influence causal attributions. But, exactly what information is being used to make these attributions? Covariation theory is based on the assumption that people use frequency information to make causal attributions (Kelley, 1967; McArthur, 1972, *etc.*). It is assumed that people sample from multiple instances of the same, or similar, events in order to make a causal attribution about a target event. The target event is compared to other instances of the event in order to establish the relative frequency of various attributes. As covariation theory is based on three variables, the occurrence of three things is checked: the relative frequency of People behaving as in the target event (consensus information); the relative frequency of Entities behaving as in the target event (distinctiveness information); and the relative frequency of the Person reacting to the Entity as in the target event (consistency information).

In each of these cases, the thing that belongs to the smaller class is taken to be the cause. So, for example, for consensus information the relative frequency of People acting in the target manner is crucial. Hence, when trying to establish what the cause of *Ted loves Mary* is, it is important to know if Ted is one of many people who loves Mary, or one of few people who loves Mary.

However, in order to test covariation theory multiple instances of an event are not presented to participants. Instead, participants are presented with a short vignette that gives information about the incidence of the target event in a population. This approach was established by McArthur (1972) when she tested covariation theory and has remained the dominant paradigm since (*e.g.* Hewstone & Jaspers, 1983; Hilton & Slugoski, 1986; Pruitt & Insko, 1980; Hilton, Smith & Kim, 1995).

Implicit covariation theory (see section 2.2.1.) also emphasises the fact that it is frequency information that is crucial when making causal attributions (Brown & Fish,

1983a; 1983b; Brown, 1986; Gordon & Randall, 1998). Hence, when describing the implicit covariation information associated with verbs Brown & Fish (1983a) refer to the underlying frequency information.

“It follows from the fact that the class  $p_1$  is smaller than the class  $p_2$  that for the purpose of predicting the occurrence or not of a predicate ‘x’ (e.g. *flattery*) between two unknown individuals, there would be more information (in the mathematical sense) in the knowledge that one participant belonged to the class  $p_1$  (the smaller class, not all humans) than that one participant belonged to the class  $p_2$  (the larger class, almost all humans).” (Brown & Fish, 1983a, 265)

Further, when speaking of NP1 biasing Agent-Patient verbs, Brown (1986) says,

“the Patient classes are very large, approximately coincident with the number of humans in the world. Agent classes would be smaller”. (Brown, 1986, 270)

As stated above, in order to test whether frequency information is indeed used in attributions, vignettes using natural language quantifiers are constructed. Typically people are presented with a passage such as the following (taken from McArthur, 1972):

(9) John laughs at the comedian.

Almost everyone who hears the comedian laughs at him.

John does not laugh at almost any other comedian.

In the past John has almost always laughed at the same comedian.

Thus, the participant is presented with the relevant base-rate information for consensus, distinctiveness and consistency. In order to provide maximally contrasting information, researchers use the quantifiers *almost everyone* and either *not...almost any other* or *hardly any* (McArthur, 1972, Orvis *et al.*, 1975, Hewstone & Jaspers, 1983, Hilton & Slugoski, 1986; Hilton & Jaspers, 1987, *etc.*). Only these quantifiers are used to denote high and low levels of covariation information. Similarly, when researchers test implicit covariation theory they use a limited range of quantifiers, commonly *few* and *many* (Brown & Fish, 1983a, 1983b; van Kleeck *et al.*, 1988; Rudolph, 1997).

To recap, participants are presented with a quantified statement giving base-rate information about the target event. The quantified statement presents one of a limited range of natural language quantifiers. It has been shown that people are sensitive to such base-rate information when making causal attributions. On this basis, researchers have concluded that frequency information influences causal attributions.

However, it has been shown that the relationship between natural language quantifiers and frequency information is not a straightforward one (see Moxey & Sanford, 1993b; 2000 for reviews). If natural language quantifiers are indeed associated with a number, then people ought to be able to translate between the two types of quantifiers. There is evidence to show that people can translate natural language quantifiers into vague numeric information. For example, Moxey (1986) showed that participants assigned the quantifiers *few* and *a few* amounts of between 10% and 25%; whereas they assigned the quantifier *almost everybody* amounts between 75% and 100%. This shows that people can translate natural language quantifiers into vague amounts: *few* denotes a smaller amount than *almost everybody*, but that there is considerable variability in the exact number people estimate a quantifier to denote.

However, there have been a number of demonstrations showing that people find it very difficult to convert natural language quantifiers into precise amounts (Moxey & Sanford, 1993a; Mullet & Rivett, 1991; Newstead & Coventry, 2000; Newstead, Pollard & Reizbos, 1987). For example, it has been demonstrated that the number assigned to a natural language quantifier depends on context. Moxey & Sanford (1993a) showed that as the base-rate expectation increases the value assigned to it also increases. If a person were attempting to translate the quantifier *many* into an amount then their estimate would differ depending on the context. Hence, *many* would receive higher numerical estimates when presented in the context of *Many people enjoyed the party*, which has a high base-rate expectation than when presented with *Many of the doctors in the hospital were female*, which has a low base-rate expectation.

Based on such evidence Moxey & Sanford (1993b; 2000) argue that although natural language quantifiers can be used to vaguely represent information about amount, they are primarily used because of the inference inducing properties each quantifier has. They argue that quantifiers are used for the perspective they provide about the information being conveyed. For example, if Mary were to say to Ted *Few of my presents were wrapped* then it can be inferred that Mary expected a certain amount of her presents to be wrapped. That is, information about Mary's beliefs are included in the naturally quantified statement, as opposed to a numerically quantified statement such as *30% of my presents were wrapped*. In the later case, it is not clear what inference can be made about Mary's beliefs (Moxey & Sanford, 1997). Did she expect fewer or more presents to be wrapped?

Natural language quantifiers function by quantifying over sets of entities. For example, *Some boys like Mary* presupposes at least three sets of entities:

- (i) The set of boys who like Mary
- (ii) The set of boys who do not like Mary
- (iii) The set of people who are not boys and who like Mary

Moxey & Sanford have shown that different quantifiers focus on different sets. Overall, quantifiers can be divided into two groups, positive quantifiers and negative quantifiers. A quantifier can be classified as positive or negative by using a battery of linguistic tests (Klima, 1964). One such test is whether the quantifier is acceptable in the environment of a negative polarity item, such as *anymore*. Negative quantifiers are acceptable in combination with such expressions, whereas positive quantifiers are not. Therefore, it is possible to say:

(10) Not many people like jazz anymore.

Or even:

(11) Few people like jazz anymore.

Whereas it is not acceptable to say:

(12) Many people like jazz anymore.

Examples of negative quantifiers are *few*, *hardly any* and *not quite all*; while examples of positive quantifiers are *a few*, *many* and *almost all*. Both negative and positive quantifiers can be used to refer to different amounts. Hence, one can denote a small amount by using either *few* or *a few*, the former of which is negative, the latter positive.

These positive and negative quantifiers show different focussing properties. One way of demonstrating this is by showing which set is referred to by means of an anaphor. Sanford & Garrod (1981) have shown that the most prominent and focussed elements in a discourse can be referred to by using a pronoun. When a set is referred to by use of a pronoun, this can be taken as evidence of that set being in focus in the mental representation. For example, after the sentence:

(13) Some of the boys played basketball.

It is possible to continue with:

(a) They were exhausted after half an hour.

But not with:

(b) They played football instead.

In (a) *they* refers to the boys that played basketball. This is equivalent to set (i) above, which Moxey & Sanford call the Standard Reference Set (or the Refset). In (b) *they* refers to the boys who did not play basketball. This is equivalent to set (ii), which Moxey & Sanford call Complement Set Reference (or the Compset). After *Some of the boys played basketball* a Refset continuation, such as in (a), is acceptable but a Compset continuation, as in (b), is not. Hence, the quantifier *some* focuses attention on the Refset.

Furthermore, it is the case that positive quantifiers focus attention on the Refset; while negative quantifiers license Compset continuations. Hence, (c) is acceptable after (14) but not after (15); however, (d) is more acceptable after (15) but not as acceptable after (14).

(14) Few people did the recommended amount of exercise this week.

(c) They watched television instead.

(15) A few people did the recommended amount of exercise this week.

(d) They had an orange juice together to celebrate.

Another difference between positive and negative quantifiers is in the content of the continuations that follow the quantified sentence. Negative quantifiers predominantly produce reasons why the predicate did not hold. On the other hand, positive quantifiers predominantly produce continuations that refer to what happened next (Moxey & Sanford, 1993b). This can be seen in the above examples also. (14) is followed by a reason why people do not do the recommended amount of exercise; whereas (15) is followed by what happened next.

To summarise, according to Moxey & Sanford a natural language quantifier is not used to denote amount alone. They argue that although quantifiers can be used to vaguely represent numerical information, specific quantifiers are used because of the inference-inviting properties that they also possess, such as emphasising different sets (Refset or Compset in the above examples). They also comment on the amount being denoted. Therefore, they lead to continuations that give a reason why something is not the case, or what happened next, *etc.*. Finally, they can also carry cues as to the speaker's or reader's beliefs. Hence, Moxey & Sanford emphasise many properties of quantifiers, which are neglected in traditional attribution studies.

Barton & Sanford (1990) attempted to examine the effects of some neglected properties of quantifiers on causal attributions. They argue that frequency information is not sufficient when making causal attributions. According to the frequency signalling account if the target individual is one of a small number of people who behave in the manner described by the verb, then the target individual is the cause. In contrast to this, Barton & Sanford propose that the focussing properties of quantifiers are the crucial cue to making attributions. According to the focussing account, quantifiers can be of two types: those which draw attention to the amount, and invite the participant to think of reasons why this is the case, *e.g. few*; and those that do not bring attention to the amount, and invite other inferences, *e.g. a few*.

In order to distinguish these two accounts, they used the traditional attribution paradigm, where participants were presented with a vignette similar to those used by McArthur (1972, example 9). When they did this, they found that the focussing property of the quantifier was the major determinant of causal attributions. Whereas *few* produced attributions consistent with the frequency signalling account, *a few* did not even though it also denotes a small amount. Therefore, Barton & Sanford argue that causal attributions are made on the principles of the focussing account.

In Chapter 4, it was demonstrated that although covariation information did influence causal attributions, other factors also affected when covariation information was used: when covariation information congruent with the implicit causality verb bias was presented there was no change in attributions. On the other hand, when the covariation information presented was incongruent with the verb bias then causal attributions were shifted in accordance with the covariation information. So, if low consensus information (a cue to NP1 attributions) was presented with a NP1 biasing verb, then although attributions were made to NP1, there was no increase in NP1 attributions from the baseline condition, when there was no covariation information. In

contrast, when high consensus information (a cue to NP2 attributions) was presented with a NP1 biasing verb, then there were fewer attributions made to NP1 than in the baseline condition.

One interpretation of these results is that covariation information is used when it is *informative*. In the case where the covariation information is congruent with the verb bias, the covariation information is not informative in making a causal bias; however, when it is incongruent with the verb bias then it is providing extra information which may help in making a correct causal attribution. In the later case, the covariation information is more informative, as it is providing extra cues to the cause of the event.

Hence, although covariation information is used to make causal attributions, it is used under particular circumstances – which demonstrates the need for a revised version of covariation theory. In order to establish which cues from covariation theory are used and when they are used, a series of experiments were conducted to isolate the relevant information that is used when making attributions.

In Experiments 10 and 11 frequency and focus information are contrasted, in order to identify which source of information determines causal attributions. According to Barton & Sanford's focus account, covariation information is only used when it is focussed; however according the frequency signalling account causal attributions are made towards the entity in the smaller class. Participants are presented with both low and high consensus and distinctiveness information in order to discover what information is used.

## 6.1. Experiment 10:

In this experiment, the effect of consensus information on attributions was examined. Participants were presented with a context sentence (*e.g. Few people flattered Mick*) and a target fragment (*e.g. Carol flattered Mick because...*). They then had to produce a continuation to the target fragment. The continuations are then analysed for whether causal attribution is made to the first or the second noun phrase.

Consensus information provides base-rate information that quantifies over the first noun-phrase (NP). The frequency signalling account predicts that when the frequency presented by the first quantified NP increases the number of completions to NP1 should decrease. Four quantifiers were used: *few*, *a few*, *many*, and *almost everybody*. *Few* and *a few* both denote small amounts and hence provide low consensus information, while *many* and *almost everybody* denote large amounts, and consequently



provide high consensus information. This is based on evidence from Moxey (1986). Moxey presented participants with pin-figure drawings, which showed different proportions of male to female surgeons. The participants were asked not to count the figures but to just describe the sketches using natural language expressions. From this data, she was able to test which quantifiers were used to denote which proportions. Participants used both *few* and *a few* to denote 10-25%, *many* to denote 25-90% and *almost all* to denote 75-100% (see also Chapter 7).

Based on this evidence, the frequency signalling account would predict that *few* and *a few* should both lead to more NP1 completions than *many* and *almost everybody*. Further, there should not be a significant difference between *few* and *a few* as they both provide the same range of frequency information. But there should be a difference between *a few* and *many* as the latter denotes a larger amount than the former.

However, according to the focussing account *few* and *a few* should give rise to different pattern of completions. Because *few* has a comment signalling function, which leads to participants focussing on reasons why an event happened, the focussing account predicts that *few* will lead to more NP1 completions than *a few*. Further, because *few* and *many* and *few* and *almost everybody* also differ in focussing properties, *few* should lead to more NP1 completions than both of these quantifiers. On the other hand, because *a few* and *many* and *almost everybody* share the same focussing properties they should not lead to different attributions.

### 6.1.1. Method

#### 6.1.1.1. *Participants*

Forty undergraduates from the University of Glasgow participated in the experiment. All were native speakers of English. None has participated in the previous experiments.

#### 6.1.1.2. *Materials*

A booklet consisting of thirty-six materials was given to the participants. The materials were based on the Agent-Patient verbs used in Experiment 4 (see section 4.1.1.2. and Appendix 3). These included 12 NP1 biasing verbs, 12 neutral verbs and 12 NP2 biasing verbs. The completions made for these verbs when there was no context sentence – just the target fragment – comprised the baseline, as in Experiment 4.

Each of the experimental items consisted of a context sentence, followed by a target sentence that the participant had to complete. There were four different context sentences, each of which provided consensus information. For example, for the verb *flattered* subjects were given one of the following context sentences:

*Few people flattered Mick* (LC)

*A few people flattered Mick* (LC)

*Many people flattered Mick* (HC)

*Almost everybody flattered Mick* (HC)

Followed by the target sentence:

*Carol flattered Mick because...*

Each verb was randomly assigned a male and a female name, the order of which was counterbalanced across booklets. A Latin-square design was employed. Altogether, a total of eight booklets were prepared. Four different random orders were used, with the remaining four being presented in the reverse random order.

#### 6.1.1.3. Procedure

Participants were asked to read each of the passages and then to complete the final sentence given the information preceding it. Participants took between 20 and 40 minutes to complete the booklet.

#### 6.1.2. Results

Each of the participant's responses were scored as either a NP1 or NP2 continuation based on whether the first pronoun in the continuation referred to the subject (NP1) or object (NP2) of the main clause in the target sentence. Continuations referring to the main clause by a repeat name were coded separately. There were 0.97% responses referring to the sentence subject with a repeat name, while 3.13% of responses referred to the sentence object with a repeat name. Any remaining continuations were coded as 'other', and these comprised 5.9% of the total responses.

In order to test whether the base-rate information was treated in the same manner as the previous experiments (section 4.1) a comparison was made between the context statements and the baseline condition. Secondly, to test the experimental hypotheses, the crucial comparisons between *few*, *a few*, *many* and *almost everybody* were made. To reiterate, if the frequency signalling account is correct then there should not be a difference in attributions between *few* and *a few* because they both signal a

small amount. However, there should be a difference between *a few* and *many*, as well as between *a few* and *almost everybody* because *a few* signals a smaller amount than both *many* and *almost everybody*. On the other hand, if the focussing account is correct then there ought to be a difference between *few* and *a few* because they differ in their focussing properties. On the other hand, there should be no difference between *a few* and *many*, and *a few* and *almost everybody*.<sup>23</sup>

Hence, the baseline comparisons are presented first so that the overall effect of consensus information can be examined. This is followed by the *few/a few* comparisons, which test the frequency and the focussing accounts.

#### 6.1.2.1. *Baseline comparisons*

Each quantifier was compared to the baseline in order to determine which, if any, of the context sentences shifted attributions. Both participant and item analyses were conducted, treating subjects as a between variable, and items as a within variable. Once again, because analyses based on NP1 continuations and NP2 continuations did not differ, only the results for NP1 continuations are presented.

The baseline comparison was conducted in order to verify that the base-rate information being presented in the context sentence was being utilised by participants in their response. Based on Experiment 4 (section 4.1.), it is hypothesised that information congruent with the bias would not shift attributions from the baseline; whereas information that was incongruent with the bias would produce a shift. To recap, *few* and *a few* provide low consensus information and this is congruent with the NP1 biasing verbs, whereas *many* and *almost everybody* provide high consensus information, which is congruent with NP2 biasing verbs.

Hence, for NP1 biasing verbs high consensus (*many* and *almost everybody*) information should decrease NP1 attributions from the baseline, whereas low consensus (*few* and *a few*) information should not shift attributions. On the other hand, for NP2 biasing verbs low consensus (*few* and *a few*) information should increase NP1 attributions from the baseline, but high consensus (*many* and *almost everybody*) should not shift attributions. Neutral verbs should be somewhere in the middle, showing some increase in NP1 completions for low consensus information and some decrease for high consensus information.

---

<sup>23</sup> . As was the case for Experiments 4 and 5, because multiple comparisons were made alpha was set to less than .01.

6.1.2.1.1. NP1 biasing verbs

Only quantifiers depicting high consensus information shifted completions from the baseline (see Figure 35). This is in line with the findings from Experiment 4. *Many* produced reliably fewer NP1 completions than the baseline [ $t_1(1, 62) = 2.07, p < .02$ ;  $t_2(1, 11) = 5.63, p < .0001$ ], as did *almost everybody* [ $t_1(1, 62) = 4.28, p < .0001$ ;  $t_2(1, 11) = 6.57, p < .0001$ ]. In fact, there was a reversal of the bias with *almost everybody*, as there were fewer than 40% completions to NP1. This reversal is again consistent with the findings of Experiment 4.




Figure 35: Percentage of NP1 completions for context and baseline conditions for NP1 biasing verbs (error bars show  $MS_e$ ). The baseline is shown as a line only for convenience. There is a decrease in NP1 completions for high consensus information (*i.e.* *many* and *almost everybody*).

6.1.2.1.2. Neutral verbs

For neutral verbs, it was predicted that there would be a trend for a shift from the baseline with all four context sentences. There was evidence of *many* causing fewer NP1 attributions than the baseline condition [ $t_1(1, 62) = 2.05, p < .02$ ;  $t_2(1, 11) = 3.67, p < .002$ ], as there was for *almost everybody* (although this is only reliable by participants) [ $t_1(1, 62) = 2.39, p < .01$ ;  $t_2(1, 11) = 1.39, p < .08$ ]. (See Figure 36.)




Figure 36: Percentage of NP1 completions for context and baseline conditions for Neutral verbs (error bars show  $MS_e$ ). The baseline is shown as a line only for convenience. There is a slight increase in NP1 completions for low consensus information (*i.e. few* and *a few*) and a decrease in NP1 completions for high consensus information (*i.e. many* and *almost everybody*).

### 6.1.2.1.3. NP2 biasing verbs

As predicted, low consensus information shifted attributions from the baseline. Both *few* [ $t_1(1, 62) = 2.05, p < .02$ ;  $t_2(1, 11) = 3.67, p < .002$ ] and *a few* [ $t_1(1, 62) = 2.05, p < .02$ ;  $t_2(1, 11) = 3.67, p < .002$ ] led to reliably more NP1 completions than the baseline condition (see Figure 37.). However, the context information did not lead to a reversal of the bias, as was seen in Experiment 4. There was no shift for high consensus information.




Figure 37: Percentage of NP1 completions for context and baseline conditions for NP2 biasing verbs (error bars show  $MS_e$ ). The baseline is shown as a line only for convenience. There was an increase in NP1 completions for low consensus information (*i.e.* *few* and *a few*).

#### 6.1.2.1.4. Summary

To summarise, the baseline comparisons were conducted to test whether consensus information has the same effect that it had in Experiment 4. Overall, this was indeed the case. As with Experiment 4 information that was incongruent with the bias of the verb led to a shift in attributions. High consensus information decreased NP1 completions when presented with NP1 biasing verbs; whereas low consensus information increased NP1 completions when presented with NP2 biasing verbs. Neutral verbs showed some evidence of both low and high consensus information shifting attributions. Previously, although there had been a trend for an increase in NP1 completions with low consensus and a decrease with high consensus these differences had not been reliable.

#### 6.1.2.2. Frequency signalling or focussing?

Now, the frequency signalling account and the focussing account are tested. The frequency signalling account predicts that there should be no difference in attributions between *few* and *a few*, but that there should be a difference between *a few* and *many*, as well as between *a few* and *almost everybody*. On the other hand, the focussing account predicts that there ought to be a difference between *few* and *a few*, but no difference

between *a few* and *many*, and *a few* and *almost everybody*. These hypotheses were tested separately for each verb type: NP1 biasing, neutral and NP2 biasing verbs.

#### 6.1.2.2.1. NP1 biasing verbs

The results for NP1 biasing verbs were consistent with the focussing account. As predicted by the focus account there was a reliable difference between *few* and *a few* [ $t_1(1, 39) = 3.55, p < .0005$ ;  $t_2(1, 11) = 3.95, p < .001$ ]. Also, as predicted by the focus account there was not a reliable difference between *a few* and *many* [ $t_1(1, 39) = .73, p > .1$ ;  $t_2(1, 11) = .22, p > .1$ ]. However, there was a difference between *a few* and *almost everybody*, but this was reliable by participants only [ $t_1(1, 39) = 3.02, p < .002$ ;  $t_2(1, 11) = .66, p > .1$ ] (see Figure 35).

#### 6.1.2.2.2. Neutral verbs

There was not a reliable difference between *few* and *a few* [ $t_1(1, 39) = .92, p > .1$ ;  $t_2(1, 11) = 1.34, p > .1$ ], contrary to the focussing account, but consistent with the frequency signalling account. There was a difference between *a few* and *many* [ $t_1(1, 39) = 2.79, p < .004$ ;  $t_2(1, 11) = 2.08, p < .03$ ], as there was between *a few* and *almost everybody* (although this was not reliable by items) [ $t_1(1, 39) = 2.91, p < .003$ ;  $t_2(1, 11) = .23, p > .1$ ]. These differences are consistent with the frequency signalling account (see Figure 36).

#### 6.1.2.2.3. NP2 biasing verbs

There was not a reliable difference between *few* and *a few* [ $t_1(1, 39) = .14, p > .1$ ;  $t_2(1, 11) = 1.08, p > .1$ ] (see Figure 37). This was consistent with the frequency signalling account. However, the frequency signalling account also predicts a difference in attributions for *a few* and *many*. This was not the case [ $t_1(1, 39) = 1.16, p > .1$ ;  $t_2(1, 11) = 1.21, p > .1$ ]. However, the difference between *a few* and *almost everybody* predicted by the frequency signalling account was found (although this was not reliable by items) [ $t_1(1, 39) = 2.31, p < .01$ ;  $t_2(1, 11) = 1.20, p > .1$ ].

#### 6.1.2.2.4. Summary

According to the frequency signalling account, because *few* and *a few* both denote the same amounts they should give rise to the same pattern of attributions. And this should differ from the attributions found for quantifiers denoting large amounts such as *many* and *almost everybody*. However, according to the focussing account *few*



and *a few* ought to give rise to different patterns of attributions, as they differ in their focussing properties. *Few* should lead to more NP1 completions than *a few* when providing base-rate information about the first NP, because it serves a comment signalling function. It invites the participant to think of reasons why, whereas *a few* does not. *A few* shares its focussing properties with *many* and *almost everybody*. Hence, *a few* should not differ in attributions from *many* and *almost everybody*.

The results above provide mixed results. For neutral and NP2 biasing verbs, it seems that the majority of the evidence was in favour of the frequency signalling account. There was not a reliable difference between *few* and *a few*, but there was evidence of a difference between *a few* and the quantifiers denoting large amounts, *many* and *almost everybody*.

However, the results for NP1 biasing verbs overwhelmingly supported the focussing account. There was a reliable difference between *few* and *a few*: *few* led to more NP1 completions than *a few*. However, *a few* and *many* produced the same pattern of attributions, which was consistent with the focussing account but not the frequency signalling account.

### 6.1.3. Discussion

Covariation theory is based on the assumption that providing base-rate information about the first NP, that is consensus information, influences participants' causal attributions. It is assumed that it is frequency information which influences attributions, such that when information is given that a small number of people behave in a particular manner then an increase in NP1 attributions occurs; however, when a large number of people are said to act in a particular manner then a decrease in NP1 attributions occurs.

The results of this experiment support this theory but with some notable modifications. Firstly, covariation information modified attributions only when it was incongruent with the implicit causality associated with the target verb. That is, high consensus information decreased the number of NP1 completions for NP1 biasing verbs; whereas low consensus information increased NP1 completions for NP2 biasing verbs.

Secondly, it seems that covariation information does not affect attributions by providing information about frequency alone. Previous experiments have used natural language quantifiers to test whether frequency information influenced causal



attributions. However, such quantifiers do not just provide frequency information, but also carry various inferential cues too. In this experiment, it was found that although frequency does change attributions for neutral and NP2 biasing verbs, a different cue is used when making attributions for NP1 biasing verbs. According to a frequency signalling account, quantifiers denoting small amounts should lead to the same pattern of attributions and this should differ from quantifiers denoting large amounts. However, according to a focussing account *few* and *a few* ought to lead to different attributions although they both denote small amounts. This was indeed found to be the case for NP1 biasing verbs. *Few* led to more NP1 completions than did *a few*.

However, the above results introduce the question of why there is a *few/a few* difference for NP1 biasing verbs but not for neutral or NP2 biasing verbs. The next experiment attempted to answer exactly this question.

## 6.2. Experiment 11:

In the previous experiment *few* and *a few* led to a different pattern of attributions, as was predicted by the focussing theory. However, this difference was restricted to NP1 biasing verbs. There are at least two explanations for this finding.

Firstly, it could be the case that *few* and *a few* are distinguished when they present information that is congruent with the underlying cause. Hence, for NP1 biasing verbs low consensus information (*i.e. Few people flattered Mick* and *A few people flattered Mick*) which points towards NP1 as being the cause. However, for NP2 biasing verbs it is high consensus information that is congruent with the implicit cause (*i.e. Many people protected Mick* and *Almost everybody protected Mick*). Hence, according to this explanation, when information is presented that is congruent with the underlying cause of the verb, focussing differences become more important. So, the difference between *few* and *a few* is seen for NP1 biasing verbs, but not NP2 biasing verbs. This will be called the *congruency hypothesis*.

A second possible explanation could be that in order for the difference in focus between *few* and *a few* to be processed, it has to be highly salient in the reader's attention. This is the case for consensus information presented with NP1 biasing verbs because all the cues are drawing the reader's attention to the first character: the information presented quantifies over the first character. There is evidence that the first character is more salient than the second character (Sanford & Garrod, 1981; Gernsbacher, 1989). The first character is also salient because it is taken as the cause

for NP1 biasing verbs. However, for NP2 biasing verbs, the information and the bias are directing attention to different participants. The information to the first character, but the verb bias to the second. Because of this, the difference between *few* and *a few* is used in the first case but not in the second. This will be called the *saliency hypothesis*.

Presenting distinctiveness information can separate these two explanations. If the congruency hypothesis is correct then a difference between *few* and *a few* would be predicted when distinctiveness information is presented. This explanation states that the *few* and *a few* difference is seen when information congruent with the underlying bias is presented. This was the case when low consensus information was presented with NP1 biasing verbs, but not when low consensus information was presented with NP2 biasing verbs. As with consensus information there are two types of distinctiveness information: high distinctiveness and low distinctiveness. Examples of high distinctiveness statements are *Carol defied few people* and *Carol defied a few people*; whereas examples low distinctiveness statements are *Carol defied many people* and *Carol defied almost everybody*. For distinctiveness information, it is low distinctiveness information that is congruent with NP1 biasing verbs, but high distinctiveness information that is congruent with NP2 biasing verbs. High distinctiveness information uses the quantifiers denoting small quantities, that is *few* and *a few*. Hence, according to this explanation a difference between *few* and *a few* should be seen for NP2 biasing verbs.

However, if the saliency hypothesis is correct then there should be no difference between *few* and *a few* with distinctiveness information. This is because, according to this explanation, the difference between *few* and *a few* will only be seen when it is in maximal focal attention. This is the case when consensus information is presented as it is quantifying over the first character. Distinctiveness information, on the other hand, quantifies over the second noun phrase. Therefore, if information being maximally salient determines the difference between *few* and *a few* then no difference should be observed when distinctiveness information is presented. (See Table 16.)

The present experiment attempts to test these two explanations by presenting participants with a context that provides distinctiveness information. To recap, for distinctiveness information, both the congruency hypothesis and the saliency hypothesis do not predict a difference between *few* and *a few* for NP1 biasing verbs. However, the congruency hypothesis predicts that there will be a difference between *few* and *a few* for NP2 biasing verbs; while the saliency hypothesis does not predict such a difference.

Type of information	Congruency hypothesis		Saliency hypothesis	
	NP1 basing	NP2 biasing	NP1 basing	NP2 biasing
LC	difference	no difference	difference	no difference
HC	–	–	–	–
LD	–	–	–	–
HD	no difference	difference	no difference	no difference

Table 16: The predictions of the congruency hypothesis and the saliency hypothesis for differences between *few* and *a few*. There are no predictions made for high consensus (HC) or low distinctiveness (LD) information because they present information about large quantities. The crucial difference between the two hypotheses is in the case of high distinctiveness information for NP2 biasing verbs. The congruency hypothesis predicts that there will be a difference between *few* and *a few*; while the saliency hypothesis predicts no difference.

6.2.1. Method

6.2.1.1. Participants

Forty undergraduates from the University of Glasgow participated in this experiment. One participant did not complete the booklet and so was excluded from further analyses. All were native speakers of English.

6.2.1.2. Materials and Procedure

The materials and procedure were exactly the same as the previous experiment, with one exception. Instead of participants being presented with consensus information they were presented with one of four pieces of information about distinctiveness, that is:

- Carol flattered many people* (LD)
- Carol flattered almost everybody* (LD)
- Carol flattered few people* (HD)
- Carol flattered a few people* (HD)

Followed by the target sentence:

*Carol flattered Mick because...*

## 6.2.2. Results

Responses were coded as before. A total of 1.91% of responses referred to the sentence subject with a repeat name, while 4.72% of responses referred to the sentence object with a repeat name. Any remaining continuations were coded as ‘other’, and these comprised 5.31% of the total responses.

First, a comparison was made between the context statements and the baseline condition. Secondly, the crucial comparisons between *few* and *a few* were made.

### 6.2.2.1. Baseline comparisons

Each quantifier was compared to the baseline in order to determine which, if any, of the context sentences shifted attributions. It was hypothesised that information incongruent with the implicit bias of the verbs would shift attributions, as was found in Experiment 10, and in Chapter 4. For NP1 biasing verbs high distinctiveness information (*few* and *a few*) should shift attributions from the baseline, whereas low distinctiveness information (*many* and *almost everybody*) should not. On the other hand, for NP2 biasing low distinctiveness information (*many* and *almost everybody*) should shift attributions from the baseline, high distinctiveness information (*few* and *a few*) should not. Neutral verbs should be somewhere in the middle, showing some decrease in NP1 completions for low distinctiveness information and some increase for high distinctiveness information.

#### 6.2.2.1.1. NP1 biasing verbs

High distinctiveness information caused a shift from the baseline as was predicted. Both *few* [ $t_1(1, 61) = 1.67, p < .05$ ;  $t_2(1, 11) = 3.66, p < .002$ ] and *a few* [ $t_1(1, 61) = 1.63, p < .05$ ;  $t_2(1, 11) = 3.44, p < .003$ ] decreased NP1 completions from the baseline. Also as hypothesised, there were no reliable shifts with low distinctiveness information (see Figure 38).




Figure 38: Percentage of NP1 completions for context and baseline conditions for NP1 biasing verbs (error bars show  $MS_e$ ). The baseline is shown as a line only for convenience. There was a decrease in NP1 completions for high distinctiveness information (*i.e. few and a few*).

#### 6.2.2.1.2. Neutral verbs

There was a reliable difference between *many* and the baseline [ $t_1(1, 61) = 1.93$ ,  $p < .03$ ;  $t_2(1, 11) = 2.38$ ,  $p < .02$ ]. There was also evidence of a shift with *almost everybody*, but this was only reliable by items [ $t_1(1, 61) = 1.36$ ,  $p < .09$ ;  $t_2(1, 11) = 2.72$ ,  $p < .01$ ]. (See Figure 39.)




Figure 39: Percentage of NP1 completions for context and baseline conditions for Neutral verbs (error bars show  $MS_e$ ). The baseline is shown as a line only for convenience. There was a slight increase of NP1 completions for low distinctiveness information (i.e. *many* and *almost everybody*) and a slight decrease for high distinctiveness information (i.e. *few* and *a few*).

#### 6.2.2.1.3. NP2 biasing verbs

It was predicted that low distinctiveness information would increase attributions from the baseline. This was indeed the case, for both *many* [ $t_1(1, 61) = 3.19, p < .001$ ;  $t_2(1, 11) = 3.33, p < .004$ ] and *almost everybody* [ $t_1(1, 61) = 3.16, p < .001$ ;  $t_2(1, 11) = 3.62, p < .002$ ]. However, there was also a shift with *a few* [ $t_1(1, 61) = 1.87, p < .07$ ;  $t_2(1, 11) = 2.68, p < .02$ ]. *A few* presented high distinctiveness information, and so according to covariation theory ought to decrease attributions from the baseline. However, this was not the case. *A few* reliably increased attributions, in a manner comparable to *many* and *almost everybody*. However, there was no reliable difference in attributions between *few* and the baseline. (See Figure 40.)




Figure 40: Percentage of NP1 completions for context and baseline conditions for NP2 biasing verbs (error bars show  $MS_e$ ). The baseline is shown as a line only for convenience. There was an increase in NP1 completions for low distinctiveness information (i.e. *many* and *almost everybody*).

#### 6.2.2.1.4. Summary

The previous experiments in this thesis showed that information incongruent with the implicit causality bias caused a shift in attributions. Overall, this was found in the present study also. High distinctiveness information caused a shift in attributions for NP1 biasing verbs; whereas low distinctiveness information caused a shift in attributions for NP2 biasing verbs. However, there was also a shift in attributions for NP2 biasing verbs when high distinctiveness information contained the quantifier *a few*. *A few* behaved more like *many* and *almost everybody* than it did like *few*.

#### 6.2.2.2. Is there a focus difference?

In this section, a direct test of the congruency hypothesis and the saliency hypothesis was made. To recap, according to the congruency hypothesis the *few* ~ *a few* difference appears when the two quantifiers are presenting information congruent with the underlying cause. High distinctiveness information presents information about small quantities. Further, it is congruent with NP2 biasing verbs, but incongruent with NP1 biasing verbs. Hence, the congruency hypothesis predicts a difference in attributions between *few* and *a few* for NP2 biasing verbs, but not for NP1 biasing verbs.

On the other hand, according to the saliency hypothesis information has to be maximally in focus in order for a difference between *few* and *a few* to be apparent. This was the case in the previous experiment when low consensus information was paired with NP1 biasing verbs. However, it is not the case for any combination of information in this experiment. Therefore, the saliency hypothesis does not predict any differences between *few* and *a few*.

So, both accounts predict no difference between *few* and *a few* for NP1 biasing verbs or neutral verbs. The critical comparison is between *few* and *a few* for NP2 biasing verbs. According to the congruency hypothesis there should be a difference between *few* and *a few*, but not between *a few* and *many*; while according to the saliency hypothesis there should not be a difference between *few* and *a few*, but there should be a difference between *a few* and *many*.

#### 6.2.2.2.1. NP1 biasing verbs

There was no difference between *few* and *a few* [ $t_1(1, 38) = .00, p > .1$ ;  $t_2(1, 11) = .00, p > .1$ ]. However, there was a reliable difference between *a few* and *many* [ $t_1(1, 38) = 4.54, p < .0001$ ;  $t_2(1, 11) = 4.34, p < .0005$ ]. There was also a difference between *a few* and *almost everybody* [ $t_1(1, 38) = 2.14, p < .04$ ;  $t_2(1, 11) = 2.05, p < .07$ ]. These results are consistent with both the congruency hypothesis and the saliency hypothesis.

#### 6.2.2.2.2. Neutral verbs

For completeness sake, the results for neutral verbs are presented, also. Neutral verbs behaved in the same manner as NP1 biasing verbs. There was not a reliable difference between *few* and *a few* [ $t_1(1, 38) = .77, p > .1$ ;  $t_2(1, 11) = .65, p > .1$ ]. But once again *a few* and *many* were reliably different [ $t_1(1, 38) = 2.08, p < .02$ ;  $t_2(1, 11) = 2.03, p < .04$ ]. There was not a reliable difference between *a few* and *almost everybody* [ $t_1(1, 38) = 1.75, p < .08$ ;  $t_2(1, 11) = 2.15, p < .05$ ]. These results are consistent with both the congruency hypothesis and the saliency hypothesis.

#### 6.2.2.2.3. NP2 biasing verbs

NP1 completions: The crucial comparison in order to distinguish the above two explanations is between *few* and *a few* for NP2 biasing verbs. There was not a reliable difference between *few* and *a few* [ $t_1(1, 38) = 1.02, p > .1$ ;  $t_2(1, 11) = 1.00, p > .1$ ]. The responses to *a few* and *many* did differ, but this was reliable only by participants [ $t_1(1,$



38) = 1.84,  $p < .04$ ;  $t_2(1, 11) = 1.38, p < .09$ ]. The same held true for the difference between *a few* and *almost everybody* [ $t_1(1, 38) = 2.23, p < .03$ ;  $t_2(1, 11) = 1.48, p > .1$ ].

6.2.2.2.4.     *Summary*

Overall, the results supported the saliency hypothesis, but congruency appeared to playing a role, also. For all verb types, the difference between *few* and *a few* was not reliable, but the difference between *a few* and *many* was. This was consistent with the saliency hypothesis. However, the latter difference was only marginally significant for NP2 biasing verbs. And indeed, if on comparison of NP1 biasing verbs, through to neutral verbs to NP2 biasing, the numerical difference between *few* and *a few* becomes larger and the difference between *a few* and *many* and between *a few* and *almost everybody* becomes smaller. This is consistent with the congruency hypothesis.

Another implication of the congruency hypothesis is that information congruent with the bias is differentiated more than information incongruent with the bias. Hence, a further test of the congruency hypothesis is to make a direct comparison between information congruent and incongruent with the bias. Low consensus information is congruent with NP1 biasing verbs, but high distinctiveness information is incongruent with NP1 biasing verbs. On the other hand low consensus information is incongruent with NP2 biasing verbs, and high distinctiveness information is congruent with NP2 biasing verbs. Both low consensus and high distinctiveness information present quantifiers denoting small amounts, namely *few* and *a few* (see Table 17).

Covariation information	NP1 biasing verbs	NP2 biasing verbs
Low consensus (Experiment 10)	congruent	incongruent
High distinctiveness (Experiment 11)	incongruent	congruent

Table 17: Both low consensus information and high distinctiveness information use small quantifiers, *i.e. few* and *a few*. For NP1 biasing verbs, low consensus information is congruent with the bias and for NP2 biasing verbs, high distinctiveness is congruent with the bias. If congruency underlies the *few/a few* difference then there should be a bigger difference between *few* and *a few* when covariation information is congruent with the verbs bias.

If information is attended to more when it is congruent with the bias than when it is incongruent, an interaction would be expected when consensus and distinctiveness information is compared directly. For NP1 biasing verbs when *a few* is congruent with the bias it should produce less NP1 responses than when it is incongruent with the bias, whereas *few* should produce more NP1 responses when it is congruent with the bias than when it is incongruent with the bias. The opposite pattern of results would be predicted for NP2 biasing verbs. *A few* should produce more NP1 continuations when it is congruent with the bias than when it is incongruent with the bias; whereas *few* should produce less NP1 completions when is it congruent than when it is incongruent with the bias.

#### 6.2.2.3. *Congruency or Saliency?*

For each verb type a 2 (quantifier [*few* vs. *a few*]) by 2 (congruency [congruent vs. incongruent]) ANOVA was performed. It was hypothesised that there would be a quantifier by congruency interaction. That is, for NP1 biasing verb congruent *few* would lead to more NP1 completions than incongruent *few*; whereas, congruent *a few* would lead to less NP1 completions than incongruent *a few*. On the other hand, congruent *few* would lead to less NP1 completions than incongruent *few*, and vice versa for congruent and incongruent *a few*.

##### 6.2.2.3.1. *NP1 biasing verbs*

A quantifier by congruency interaction was predicted. This was confirmed [ $F_1(1, 77) = 5.20, MS_e = .05, p < .03$ ;  $F_2(1, 11) = 7.05, MS_e = .01, p < .02$ ] (see Figure 41). Furthermore, it was predicted that when *few* was congruent with the bias then it would lead to more NP1 completions than when it was incongruent with the bias. This was confirmed [ $t_1(1, 77) = 3.04, p < .002$ ;  $t_2(1, 11) = 3.25, p < .004$ ]. Finally, it was hypothesised that when *a few* was congruent with the bias then it would lead to less NP1 completions than when it was incongruent. This was not substantiated by the data, as there was not a reliable difference between congruent *a few* and incongruent *a few* [ $t_1(1, 77) = .41, p > .1$ ;  $t_2(1, 11) = 1.00, p > .1$ ].




Figure 41: Percentage of NP1 completions for congruent and incongruent information for NP1 biasing verbs (error bars show  $MSE$ ).

#### 6.2.2.3.2. NP2 biasing verbs

There was not a reliable quantifier by congruency interaction [ $F_1(1, 77) = .67$ ,  $MS_e = .07$ ,  $p > .1$ ;  $F_2(1, 11) = .26$ ,  $MS_e = .05$ ,  $p > .1$ ] (see Figure 42). Further, there were no reliable main effect of quantifier [ $F_1(1, 77) = .39$ ,  $MS_e = .07$ ,  $p > .1$ ;  $F_2(1, 11) = .48$ ,  $MS_e = .02$ ,  $p > .1$ ]. There was some evidence of an effect of congruency, but this was reliable by items only [ $F_1(1, 77) = 2.61$ ,  $MS_e = .12$ ,  $p > .1$ ;  $F_2(1, 11) = 8.25$ ,  $MS_e = .01$ ,  $p < .01$ ]. This showed that incongruent information ( $M = 48.8\%$ ) led to more NP1 completions than did congruent information ( $M = 38.8\%$ ). Again, this can be seen as a further demonstration of the effect of incongruent information.




Figure 42: Percentage of NP1 completions for congruent and incongruent information for NP2 biasing verbs (error bars show  $MS_e$ ).

#### 6.2.2.3.3. Summary

To summarise, it was hypothesised that participants would be more sensitive to differences in congruent information than they would be to incongruent information. This would be evident in a quantifier by congruency interaction. There is some evidence to support this. There was an interaction but this appeared on only one measure: for NP1 biasing verbs in NP1 completions. Under these circumstances there were more completions to *few* when it was congruent than there were to it when it was incongruent.

In all other cases, however, only a main effect of congruency was observed. This showed that incongruent information caused a weakening of the verb bias. For both NP1 and NP2 biasing verbs presentation of incongruent information, regardless of whether it was *few* or *a few*, shifted completions to around 50%.

### 6.2.3. Discussion

Experiments 10 and 11 together demonstrated that the frequency signalling account was not sufficient to account for the pattern of attributions produced when participants were given covariation information. In fact, there were modifying factors

such as whether the covariation information was congruent with the implicit causality bias and whether it presented information that made the participant think of reasons why this was the case. That is, there was a difference between *few* and *a few* which was not predicted by the frequency signalling account. Evidence from Experiment 11 suggested that the focussing account was mediated by a mixture of saliency and congruency. The saliency hypothesis predicts that the *few/a few* difference only occurs when consensus information is presented with NP1 biasing verbs. This is because all of the cues: from the verb, context and general cognitive preference draws attention to NP1. Therefore, information about NP1 is distinguished more. However, the congruency hypothesis predicts that the difference between *few* and *a few* is also apparent when distinctiveness information is presented. When high distinctiveness information is presented with NP2 biasing verbs then the *few/a few* distinction should be apparent.

Consistent with the saliency hypothesis and contrary to the congruency hypothesis, there was not a reliable difference between *few* and *a few* when high distinctiveness information was presented with NP2 biasing verbs. However, there was a numerical trend in the predicted direction. Furthermore, when the quantifiers were congruent with the bias then there did appear to be a bigger difference between *few* and *a few* than when the information was incongruent with the bias.

In the next chapter, the effects of covariation information were examined further. Experiments 12 and 13 examined the effects of presenting participants with numerical information, while Experiments 14 and 15 presented percentage information and measured how this information affected causal attributions.

# CHAPTER 7:

## COVARIATION THEORY AND

### NUMERICAL AND PERCENTAGE INFORMATION

#### 7.0. Introduction

In Chapter 6, it was shown that covariation theory assumes people use frequency information in order to make causal attributions: attribution is made to the person in the relatively smaller group. This was contrasted with the focussing account of Barton & Sanford, which claims that information about a small amount is not enough to make attributions; rather, attention also has to be drawn to the fact that the proportion is indeed small and then make the individual think of reasons why this might be the case. Based on the results of Experiments 10 and 11, it is clear that the pattern of attributions people produce are best accounted for using a mixed model, where frequency and focus are both important. Frequency information does influence attributions, but this is modified in a number of ways, including the type of quantifier used and whether the covariation information is congruent with other cues to causality, such as the implicit causality bias.

In this chapter, a closer examination of the effects of frequency information on attributions is made. A number of studies have shown that it is possible to translate natural language quantifiers into vague numeric information (see section 6.0.). People can assign natural language quantifiers numerical values on a vague scale, such that quantifiers like *few* are assigned a much smaller amount than quantifiers like *almost everybody* (e.g. Moxey, 1986). Based on such evidence, the previous chapter assumed that *few* and *a few* denote the same amount because they are known to be used for the same range of numerical values.

However, it is also known that the interpretation of a quantifier can change depending on the context in which it is presented (Hormann, 1983; Moxey & Sanford, 1993a; Newstead & Coventry, 2000). Therefore, before accepting a mixed model of attributions it is important to rule out the possibility that in the present context *few* and *a few* actually represent different amounts. If it is the case that *a few* is translated into a

larger amount than *few* is then the difference between *few* and *a few* found in the previous chapter could be the result of frequency information alone.

Therefore, in Experiment 12, participants were asked to translate the covariation information presented in Experiment 10 into a numerical value, in order to establish people take *few* and *a few* to denote different amounts. This was followed by an attribution experiment (Experiment 13), where covariation information was presented in a numerical form. Experiment 14 further investigated the amount *few* and *a few* are taken to denote, by asking participants to translate covariation information into proportional values. This is followed by a final attribution experiment (Experiment 15), where the affect of proportional information on attributions was examined.

## 7.1. Experiment 12:

It could be argued that the difference in attributions found between *few* and *a few* for NP1 biasing verbs is due to the two quantifiers denoting different frequencies. If *few* denotes a smaller quantity than *a few* then covariation theory would be able to account for the above results purely in terms of frequency. Covariation theory predicts that as the frequency of the quantifier increases the number of attributions to NP1 will decrease, and those to NP2 increase. Consequently, in this experiment an attempt was made to discover whether the difference between *few* and *a few* could be accounted for in terms of a difference in absolute number.

Participants were given each of the context sentences in Experiment 10 (consensus information) and were asked to indicate how many people behaved in the manner described in the sentence.

### 7.1.1. Method

#### 7.1.1.1. *Participants*

Sixteen undergraduates from the University of Glasgow participated in this experiment. None of the participants had participated in the previous study.

#### 7.1.1.2. *Materials*

The context sentences from the previous experiment were presented to participants in a booklet, such that each participant saw each of the 36 verbs once, with one of the four quantifiers: *few*, *a few*, *many*, *almost everybody*. Following every

context sentence, participants were asked to indicate how many people acted in the manner described in the context sentence. For example:

*Few people flattered Mick.*  
*How many people flattered Mick?* .....

A total of eight booklets were created, using four random orders, with the remaining four in the reverse order.

7.1.1.3. Procedure

Participants were presented with a booklet. They were asked to read each of the sentences carefully and then indicate how many people had acted in the manner described. It took participants around 10 minutes to complete the booklet.

7.1.2. Results

The estimates given for each class of verb is presented below, in Table 18. Because the elicited numerical scores were skewed, the median scores were used in further analyses.

The critical comparison was between the numbers provided for *few* and *a few* for NP1 biasing verbs. The frequency signalling account would predict that *a few* denotes a larger amount than *few* (because *a few* led to fewer NP1 continuations than *few*). There was a reliable difference between *few* and *a few* in the predicted direction [ $t_1(1, 15) = 3.16, p < .003$ ;  $t_2(1, 11) = 1.72, p < .06$ ]. Hence, it could be argued that the difference in attributions found in Chapter 6 for these quantifiers was due to this underlying numerical difference.

However, there is a significant numerical difference between *few* and *a few* for neutral verbs [ $t_1(1, 15) = 2.39, p < .02$ ;  $t_2(1, 11) = 2.30, p < .02$ ], and NP2 biasing verbs, also [ $t_1(1, 15) = 2.64, p < .009$ ;  $t_2(1, 11) = 2.27, p < .02$ ]. Neither of these numerical differences gave rise to a difference in completions in Experiment 10, in the previous chapter. If the numerical difference between *few* and *a few* for NP1 biasing verbs was the cause of the difference in attributions witnessed in Experiment 10, then a difference in attributions for neutral verbs and NP2 biasing verbs would, also, be expected. However, this was not the cause. This suggests that numerical differences are not sufficient to give rise to the *few/a few* difference.



Quantifier	NP1 biasing verbs		Neutral verbs		NP2 biasing verbs	
	mean	median	mean	median	mean	median
Few	2.77	2.83	3.26	2.85	3.19	3.59
A few	4.19	3.60	3.86	3.64	8.81	3.88
Many	13.80	12.04	16.05	8.51	13.86	12.31
Almost	16.54	14.39	19.01	13.48	56.96	13.48

Table 18: Mean and median numerical estimates for quantifiers by verb type.

### 7.1.3. Summary

It could be argued that the difference in attribution found between *few* and *a few* for NP1 biasing verbs in the previous experiment was not due to the difference in focus, but actually due to a difference in the numbers these quantifiers denote. When participants were asked to generate numbers for the context sentences in the previous experiment then there was indeed a reliable difference in the numbers that *few* and *a few* denote. *A few* was taken to denote a higher number than *few*. This is consistent with the frequency signalling account.

Further, there was a difference between *few* and *a few* for neutral verbs and NP2 biasing verbs, as was the case for NP1 biasing verbs. However, neither of the former two classes of verbs showed a difference between *few* and *a few*. This suggests that the number information is not sufficient to produce the attribution difference in the previous study. In order to test whether this numerical difference is a sufficient cue to influence causal attributions numerical covariation information was presented in Experiment 13.

## 7.2. Experiment 13:

This experiment tested whether information presented in numerical form influences causal attributions in discourse. There were three main questions that were addressed.

First, and most importantly, does the numerical difference found in Experiment 12 give rise to a difference in attributions? According to a frequency signalling account the difference in attributions found in Experiment 10 is due to the quantifiers denoting different frequencies. There was evidence supporting this in the previous study. If this is the correct explanation, then when the numerical information is presented as context,

there should be the same pattern of completions as found when natural language quantifiers were presented. However, according to the focus theory, presentation of the equivalent numerical information should not lead to a difference in attributions, as there is no focus difference.

Second, some general predictions can be made about the use of numerical information on attributions. According to the covariation theory of attributions there should be an overall decrease in the number of NP1 completions as the number presented in the context sentence increases. And finally, if numbers are equivalent to natural language quantifiers, in the relevant respects, then certain predictions can be made when comparing the numerical context to the baseline.

For NP1 biasing verbs, numbers that give high consensus information, namely *fifteen* and *thirty* should decrease NP1 completions. For neutral verbs, there should be no reliable difference between context and baseline conditions. For NP2 biasing verbs, numbers that provide low consensus information, namely *three* and *six*, should increase NP1 completions.

### **7.2.1. Method**

#### **7.2.1.1. *Participants***

Forty undergraduates from the University of Glasgow participated in this experiment. All were native speakers of English. None had participated in any previous experiment.

#### **7.2.1.2. *Materials***

Agent-Patient verbs from Experiment 4 and the previous chapter (see section 4.1.1.2. and Appendix 3) were used to construct the materials. These included 12 NP1 biasing verbs, 12 neutral verbs and 12 NP2 biasing verbs. The completions made for these verbs when there was no context sentence – just the target fragment – comprised the baseline, as before.

In order to generate the context sentences, the numerical estimates provided for Experiment 12 were collapsed over verb type (see Table 19), so that four numerical quantifiers would be used, equivalent to the four natural language quantifiers used in Experiment 10. From this data, the mean estimates were used as a basis for the context information as they maximised the numerical differences between quantifiers. This allowed the fairest test of the frequency signalling account.

Quantifier	Mean	Median
Few	3.07	2.59
A few	5.62	3.88
Many	14.42	12.31
Almost every	30.19	13.48

Table 19: Overall mean and median numerical estimates for quantifiers by verb type.

The materials were constructed in exactly the same manner as in Experiment 10. The context sentences were:

*Three people flattered Mary*

*Six people flattered Mary*

*Fifteen people flattered Mary*

*Thirty people flattered Mary*

The target sentence fragment followed each of these:

*John flattered Mary because...*

### 7.2.1.3. Procedure

Participants were asked to read each of the passages and then to complete the final sentence given the information preceding it. Participants took between 20 and 40 minutes to complete the booklet.

## 7.2.2. Results

Continuations were coded and analysed as in Experiment 10. Of the responses, 0.35% referred to the sentence subject using a repeat name, while 3.06% of the responses referred to the sentence object using a repeat name. The remaining other responses constituted 11.25% of the total responses. As before, only the results for NP1 continuations are presented here.

The results of the critical comparison between *three* and *six* are presented first, followed by the overall effect of numbers on attributions and then by the baseline comparisons.

### 7.2.2.1. *Is the focus difference due to a numerical difference?*

In order to answer the question of whether the difference between *few* and *a few* found with NP1 biasing verbs in Experiment 10 was driven by the numerical difference found in Experiment 12, a crucial comparison was made. If *few* led to more NP1 completions because it referred to a smaller number than *a few*, then *three* should lead to reliably more NP1 completions than *six* for NP1 biasing verbs. Contrary to the frequency signalling account, there was no difference between *three* and *six* [ $t_1(1, 39) = .15, p > .1$ ;  $t_2(1, 11) = .11, p > .1$ ]. There was no evidence that *six* produced fewer NP1 completions than *three* did.

### 7.2.2.2. *Is there an overall effect of numerical information?*

Four (quantifier [three vs. six vs. fifteen vs. thirty]) by 3 (verb type [NP1 vs. neutral vs. NP2]) linear trends analyses were conducted by participants ( $F_1$ ) and items ( $F_2$ ). It was predicted that as the number provided in the context sentence increased, the number of completions to NP1 would decrease. A main effect of verb type was predicted, such that NP1 biasing verbs would produce the most NP1 completions followed by neutral verbs, and then NP2 biasing verbs. No other effects were predicted.

There was, indeed, a linear effect of quantifier on completions [ $F_1(1, 39) = 9.39, MS_e = 4.68, p < .004$ ;  $F_2(1, 33) = 6.87, MS_e = 13.07, p < .01$ ]. As the number increased, there were fewer continuations to NP1. (See Figure 43.) There was also a main effect of verb type [ $F_1(2, 78) = 17.07, MS_e = 8.65, p < .0001$ ;  $F_2(2, 33) = 8.01, MS_e = 25.77, p < .001$ ]. There were most completions to NP1 biasing verbs ( $M = 5.58$ ), followed by neutral verbs ( $M = 4.73$ ) and finally NP2 biasing verbs ( $M = 4.13$ ). There were no other effects.




Figure 43: The overall effect of quantifier on NP1 completions.

To summarise, an overall effect of quantifier was found on completions. There was a main effect of verb type, which was seen in all measures. This was the same effect that was found in the previous experiments using Agent-Patient verbs.

#### 7.2.2.3. Baseline comparisons

Each context sentence was compared to the baseline in order to determine if numbers behaved in the same way as natural language quantifiers. According to the frequency signalling account: for NP1 biasing verbs *fifteen* and *thirty* (high consensus information) should decrease the number of NP1 completions in comparison to the baseline; whereas for NP2 biasing verbs *three* and *six* (low consensus information) should increase the number of NP1 completions made.

##### 7.2.2.3.1. NP1 biasing verbs

Overall, the numerical context information produced a decrease in NP1 attributions – effectively causing a weakening of the effect of the implicit causality bias. (These differences can be seen in Figure 44.) As for the predicted shifts for *fifteen* and *thirty*, only *thirty* produced reliably fewer NP1 completions [*fifteen*:  $t_1(1, 62) = 1.36, p < .09$ ;  $t_2(1, 11) = 1.16, p > .1$ ; *thirty*:  $t_1(1, 62) = 2.33, p < .01$ ;  $t_2(1, 11) = 2.65, p < .01$ ].




Figure 44: Percentage of NP1 completions for context and baseline conditions for NP1 biasing verbs (Error bars show MSe).

#### 7.2.2.3.2. Neutral verbs

There were no reliable differences on any measure for neutral verbs. Figure 45 shows the means in the two conditions.




Figure 45: Percentage of NP1 completions for context and baseline conditions for Neutral verbs (Error bars show MSe).

#### 7.2.2.3.3. NP2 biasing verbs

There was a relative increase in the number of NP1 completions for all quantifiers. As with the NP1 biasing verbs, the context information appeared to have weakened the effect of the implicit causality bias (see Figure 46).

It was predicted that *three* and *six* would increase the number of NP1 completions. *Three* did indeed increase NP1 completions, although this effect was marginal by participants [ $t_1(1, 62) = 1.76, p < .04$ ;  $t_2(1, 11) = 2.57, p < .01$ ]. There was, also, evidence for a shift when *six* was presented [ $t_1(1, 62) = 1.69, p < .05$ ;  $t_2(1, 11) = 2.11, p < .01$ ].




Figure 46: Percentage of NP1 completions for context and baseline for NP2 biasing verbs (error bars show MSe)

#### 7.2.2.3.4. Summary

According to the frequency signalling account (1) numbers *fifteen* and *thirty* should decrease NP1 completions for NP1 biasing verbs, (2) there should be no reliable difference between context and baseline conditions for neutral verbs and (3) *three* and *six* should increase NP1 completions for NP2 biasing verbs. The hypotheses were confirmed.

Further, it appeared that there was a relative weakening of the bias with all four numerical quantifiers. There was a numerical decrease in the number of NP1



completions for NP1 biasing verbs and a numerical decrease in NP2 completions for NP2 biasing verbs.

### 7.2.3. Discussion

According to the frequency signalling account, causal attributions are made to a person if (s)he is one of a small proportion of people which acts in the target manner. When participants were presented with natural language quantifiers, then this did indeed seem to be the case: the smaller the group a person was in the more likely they were to be seen as causal. When the same natural language quantifiers were translated into numerical information, and that numerical information was given as covariation information, then once again, it could be seen that as the group size grew the individual was seen as less causal.

However, it is crucial to note that numerical information does not account for the whole story. In Experiment 10, *a few* produced less NP1 attributions for NP1 biasing verbs than did *few*. When the equivalent numerical quantifiers were contrasted there was no such difference. Therefore, the results from Experiments 12 and 13 are consistent with the mixed model of attributions, where frequency information is mediated by focus.

Before accepting this conclusion, however, there is a limitation of the above studies. Asking participants to translate a natural language quantifier into an absolute value is a very difficult task. It is unclear what absolute numbers mean when there is no norm information given. For example, if a participant were to say that *Few people flattered Mary* is equivalent to *Three people flattered Mary*, does that mean three people in the room, three people in the world...? Pure numerical information does not provide any norm information. One way of providing some sort of norm information is by presenting proportional information. Proportional information gives population parameters, and allows information to be transmitted about how common or uncommon an event is relative to the population. Hence, the effect of percentage information on attributions was tested in Experiment 15. Before that, it can be asked whether relative information about the population is sufficient to give rise to the causal attributions observed in Experiment 10. That is, is there a difference in the proportion that *few* and *a few* denote?



**7.3. Experiment 14:**

Participants were given each of the context sentences in Experiment 10 and asked to indicate what percentage of people behaved in the manner described in the sentence. According to the frequency signalling account *a few* should give higher percentage estimates than *few*.

**7.3.1. Method**

*7.3.1.1. Participants*

Sixteen undergraduates from the University of Glasgow participated in this experiment. None of the participants had participated in any of the previous studies.

*7.3.1.2. Materials and Procedure*

The materials and procedure were the same as those used for Experiment 13 in this chapter, except instead of asking participants to indicate what number of people a sentence referred to, they were asked to indicate what percentage it referred to.

**7.3.2. Results**

Below is a table (Table 20) giving both means and medians of the percentage estimations for the four different quantifiers. As with the numerical estimates, the percentage estimates were skewed, so median scores were used in further analyses.

If the frequency signalling account is correct in assuming a translation process between natural language quantifiers and percentages, then *a few* should denote a higher percentage than *few* for NP1 biasing verbs. However, there was no reliable difference between the percentages [ $t_1(1, 15) = 1.07, p > .1$ ;  $t_2(1, 11) = .17, p > .1$ ].

Quantifier	NP1 biasing verbs		Neutral verbs		NP2 biasing verbs	
	mean	median	mean	median	mean	median
Few	14.9%	11.3%	13.4%	10.5%	13.9%	14.3%
A few	12%	11.6%	14%	11.5%	16.5%	14.9%
Many	74.9%	73.1%	73.9%	74.1%	72.9%	73.1%
Almost	87.6%	91.4%	88.8%	90.5%	88.2%	90.1%

Table 20: Mean and median percentage estimates for quantifiers by verb type.

### 7.3.3. Summary

The frequency signalling account could argue that the difference in attributions for NP1 verbs between *few* and *a few* is due to a translation process, whereby the natural language quantifiers are translated into percentage information. If this was the case, then *a few* should denote a higher percentage than *few* should, as the former led to *fewer* NP1 completions. However, when participants were asked to indicate what percentage the quantifiers denoted there was no reliable difference between the percentages.

The next experiment investigated the effect of percentage base-rate information on causal attributions.

## 7.4. Experiment 15:

It has been demonstrated that proportional information does not underlie the difference between *few* and *a few*. Hence, a translation between the natural language quantifier into a proportional quantifier is not sufficient to give rise to the attribution pattern seen in Experiment 10. In this study, the effect of pure proportional information on attributions was examined.

According to the covariation theory of attributions there should be an overall decrease in the number of NP1 completions as the percentage information increases.

And as with the previous base-rate manipulations predictions can be made about the effect of the context when compared to the baseline. For NP1 biasing verbs, percentage information that provides high consensus information, namely 73% and 88%, should decrease NP1 completions. For neutral verbs, there should be no reliable difference between context and baseline conditions. For NP2 biasing verbs, percentage

information that provides low consensus information, namely 14% and 20%, should increase NP1 completions.

7.4.1. Method

7.4.1.1. Participants

Forty undergraduates from the University of Glasgow participated in this experiment. All were native speakers of English. None had participated in any of the previous studies.

7.4.1.2. Materials

The Agent-Patient verbs from Experiments 4 were used again. These included 12 NP1 biasing verbs, 12 neutral verbs and 12 NP2 biasing verbs. The completions made for these verbs when there was no context sentence – just the target fragment – comprised the baseline, as before.

In order to generate the context sentences, percentage estimates provided for Experiment 14 were collapsed over verb type (see Table 21), so that four proportional quantifiers would be used, equivalent to the four natural language quantifiers used in Experiment 10. From this data, the mean estimates were chosen to provide context information as they showed the biggest absolute differences between quantifiers. This allowed the fairest test of the frequency signalling account.

However, unlike the case when constructing materials in the numerical experiment there were not four different proportional quantifiers, as the estimates given for *few* and *a few* were the same. Hence, a small difference was introduced in the quantifiers. This allowed testing the effect of proportional information in the same manner as numerical information was tested.

Quantifier	Mean	Median
Few	14.06%	12.91%
A few	14.18%	12.51%
Many	73.42%	74.49%
Almost every	88.23%	89.91%

Table 21: Overall mean and median percentage estimates for quantifiers by verb type.

The materials were constructed in exactly the same manner as the previous production studies. The context sentences were:

*14% of people flattered Mary*

*20% of people flattered Mary*

*73% of people flattered Mary*

*88% of people flattered Mary*

The target sentence fragment followed the context sentence:

*John flattered Mary because...*

#### 7.4.1.3. Procedure

Participants were asked to read each of the passages and then to complete the final sentence given the information preceding it. Participants took between 20 and 40 minutes to complete the booklet.

### 7.4.2. Results

Continuations were coded as in the previous experiments. Out of the continuations, 0.56% referred to the sentence subject using a repeat name, while 3.68% referred to the sentence object in this manner. ‘Other’ completions constituted 7.15% of the total responses. All analyses are based on NP1 continuations. First, the effect of proportional information on causal completions is examined. This is followed by a comparison of the proportional information to the baseline.

#### 7.4.2.1. Is there an overall effect of proportional information?

Four (quantifier [14% vs. 20% vs. 73% vs. 88%]) by 3(verb type [NP1 vs. neutral vs. NP2]) linear trends analyses were conducted by participants ( $F_1$ ) and items ( $F_2$ ). According to the frequency signalling account as the percentage increases the number of completions to NP1 should decrease. A main effect of verb type was predicted, such that NP1 would produce the most NP1 completions followed by neutral verbs, and then NP2 biasing verbs. No other effects were predicted. The converse was predicted for NP2 completions.

There was indeed a linear effect of quantifier on completions, consistent with the frequency signalling account [ $F_1$  (1, 39) = 17.26,  $MS_e$  = 13.35,  $p$  < .0001;  $F_2$  (1, 33) = 33.17,  $MS_e$  = 44.50,  $p$  < .0001]. (See Figure 47.) As the proportion increased, there were fewer continuations to NP1. Although there was an overall decrease in

attributions to NP1, there appears to be a slight increase for the largest proportion, *i.e.* 88%.

There was also a main effect of verb type [ $F_1(2, 78) = 16.31$ ,  $MS_e = 10.75$ ,  $p < .0001$ ;  $F_2(2, 33) = 9.43$ ,  $MS_e = 33.58$ ,  $p < .001$ ]. Again, this effect confirmed the hypothesis: there were most completions to NP1 biasing verbs ( $M = 5.60$ ), followed by neutral verbs ( $M = 4.64$ ) and finally NP2 biasing verbs ( $M = 3.94$ ). There were no other effects.




Figure 47: The overall effect of percentage information on NP1 completions.

Overall, as the proportional quantifier increased the number of NP1 completions decreased. For 88%, there appeared to be an increase in NP1 completions, which was unexpected. This could be the result of random fluctuations in the data. There was also a main effect of verb type, which replicated the earlier experiments.

#### 7.4.2.2. Baseline comparisons

Each context sentence was compared to the baseline. To recap on the predictions: for NP1 biasing verbs, 73% and 88% (high consensus information) should decrease the number of NP1 completions in comparison to the baseline; whereas for NP2 biasing verbs, 14% and 20% (low consensus information) should increase the number of NP1 completions made.

7.4.2.2.1. NP1 biasing verbs

It was hypothesised that 73% and 88% would shift attributions from the baseline. 73% did produce a shift from the baseline [ $t_1(1, 62) = 3.11, p < .002$ ;  $t_2(1, 11) = 3.81, p < .002$ ]. There was not a reliable effect of 88% for participants [ $t_1(1, 62) = 1.48, p < .07$ ;  $t_2(1, 11) = 1.91, p < .04$ ]. (See Figure 48.)




Figure 48: Percentage of NP1 completions for context and baseline conditions for NP1 biasing verbs (Error bars show MSE).

7.4.2.2.2. Neutral verbs

Although no differences were predicted 14% shifted attributions from the baseline [ $t_1(1, 62) = 2.36, p < .01$ ;  $t_2(1, 11) = 2.83, p < .008$ ]. (See Figure 49.)




Figure 49: Percentage of NP1 completions for context and baseline conditions for Neutral verbs (Error bars show MSe).

#### 7.4.2.2.3. NP2 biasing verbs

It was hypothesised that 14% and 20% would increase NP1 completions. This was confirmed for 14% [ $t_1(1, 62) = 2.88, p < .003$ ;  $t_2(1, 11) = 3.09, p < .005$ ] and there was evidence that this was the case for 20%, also [ $t_1(1, 62) = 1.44, p < .08$ ;  $t_2(1, 11) = 2.35, p < .02$ ]. (See Figure 50.)






Figure 50: Percentage of NP1 completions for context and baseline conditions for NP2 biasing verbs (error bars show MSe)

To recap, according to the frequency signalling account: (1) 73% and 88% should decrease NP1 completions for NP1 biasing verbs, (2) there should be no reliable difference between context and baseline conditions for neutral verbs and (3) 14% and 20% should increase NP1 completions for NP2 biasing verbs. Overall, the hypotheses were confirmed. Although, for NP1 biasing verbs the difference between 88% and the baseline was not reliable.

### 7.4.3. Discussion

Consistent with the frequency signalling account, as percentage information increased the number of attributions to NP1 decreased. This was the same pattern of results found when natural language quantifiers were presented and when numerical information was presented. However, this overall pattern did not account for the difference that was found between *few* and *a few* in the previous chapter. The next section examines these differences more closely.



## 7.5. Comparing base-rate information

It has been argued that natural language quantifiers provide information other than frequency information, and that this information independently influences causal attributions. However, all the evidence presented has been indirect. There was a difference between *few* and *a few* when they were presented in context sentences for NP1 biasing verbs. Although these two quantifiers led to different numerical estimates, this difference did not lead to a difference in completions as was observed with the natural language quantifiers. Further, there was no difference in percentage estimates for the two quantifiers.

The focus account states that *few* should focus the reader's attention onto reasons why the event happened, and so lead to more causal attributions. If this is the case then not only would there be the difference between *few* and *a few* demonstrated in Experiment 10, there would also be a difference between *few* and both numerical and percentage information. However, because *a few* does not have this focussing property, there should be no difference between *a few* and numerical or percentage information. Further, although percentage information appears to behave more like natural language quantifiers than numerical information, percentage information is not marked, so there should not be a difference between numerical information and percentage information.

Hence, three comparisons were made: (1) between natural language quantifiers and numerical information; (2) between natural language quantifiers and percentage information; and (3) between numerical information and percentage information.

### 7.5.1. Natural language quantifiers and numerical information

The focus theory predicts that there should be a difference between *few* and numerical information, but not between *a few* and such information. Hence, *few* was compared to *three* and *a few* to *six*. It was hypothesised that *few* would lead to more NP1 completions than *three*. Analyses treated participants as a between factor, and items as a within factor.

There was a reliable difference between *few* and *three* [ $t_1(1, 78) = 2.88, p < .003$ ;  $t_2(1, 11) = 2.07, p < .03$ ] as predicted. However, there was not a reliable difference between *a few* and *six* [ $t_1(1, 78) = .12, p > .1$ ;  $t_2(1, 11) = .16, p > .1$ ]. A

power analysis based on the participant analysis gave a power of .88 to detect an effect as large as that obtained for *few* and *three*, when  $\alpha = .05$ . (See Figure 51.)




Figure 51: The percentage of NP1 completions for natural language quantifiers compared to numerical quantifiers (error bars show MSe).

7.5.2. Natural language quantifiers and percentage information

The focus theory predicts that there should be a difference between *few* and percentage information, but not between *a few* and such information. Because the percentage estimate given for both *few* and *a few* was 14%, *few* was compared to 14% and so was *a few*. It was hypothesised that *few* would lead to more NP1 completions than 14%.

There was a reliable difference between *few* and 14% as predicted by the focus theory [ $t_1(1, 78) = 2.01, p < .02$ ;  $t_2(1, 11) = 2.21, p < .03$ ]. But there was no difference between *a few* and 14% [ $t_1(1, 78) = .76, p > .1$ ;  $t_2(1, 11) = .60, p > .1$ ]. Once again, a power analysis based on the subject analysis was conducted. This gave a power of .74 to detect an effect as large as that obtained for *few* and 14%, when  $\alpha = .05$ . (See Figure 52.)




Figure 52: Percentage of NP1 completions for natural language quantifiers compared to percentage information (error bars show MSe).

### 7.5.3. Numerical information and percentage information

Because numerical information and percentage information do not differ in focussing properties, no differences were predicted: neither for *three* and *14%* nor for *six* and *14%*.

There was not a reliable difference between *three* and *14%* [ $t_1(1, 78) = .53, p > .1$ ;  $t_2(1, 11) = .44, p > .1$ ]. Similarly, there was no difference between *six* and *14%* [ $t_1(1, 78) = .37, p > .1$ ;  $t_2(1, 11) = .24, p > .1$ ]. (See Figure 21.)




Figure 53: Percentage of NP1 completions for numerical information compared to percentage information (error bars show MSE).

## 7.6. General Discussion

In Chapters 6 and 7, the frequency signalling account and the focussing account of covariation theory were examined in detail. Covariation theory states that frequency information governs the causal attributions that people make. When a person has performed an action, and they are one of a small number of people who perform that action then they are more likely to be perceived as the cause. However, as the number of people who perform that action increases, the less likely they are to be perceived as the cause. Hence, low consensus and low distinctiveness information, which state that a small number of people did X, produce attributions to NP1 (the person); whereas, high consensus and high distinctiveness information, which state that a large number of people did X, produce attributions to NP2 (the entity).

However, this frequency signalling account does not account for all of the data. In a sentence completion task, participants used covariation information as a cue to causality when it was incongruent with implicit causality verb bias. However, when the information was congruent with the causal bias of the verb, then covariation information did not affect completions.

Frequency information did affect attributions. In both chapters, it was clear that as the base-rate increased, the number of attributions to NP1 decreased: this was the case when participants were given natural language quantifiers, numerical information and percentage information. However, there was an additional effect of focus. When information that was congruent with the verb bias was put into focus by making the person ask why this was the case, then that information was given more weight than when it was not in focus. This was demonstrated by the difference in attributions between *few* and *a few*, which was observed for NP1 biasing verbs when consensus information was presented; and similarly, for NP2 biasing verbs when distinctiveness information was presented. Therefore, frequency signalling is not a sufficient account of how covariation information affects causal attributions.

## General Discussion

Causality plays a very important role in many different aspects of cognition, including learning, perception, development, social cognition, categorisation and problem solving. This thesis has examined the causal inferences people make when presented with a social scenario.

It is clear that the implicit causality verb bias plays a significant role in making causal attributions. Not only is the implicit causality bias apparent when an individual is asked a causal question about a situation between two people, it is also evident when no causal question is asked. Chapter 3 demonstrated that the predominant inference people make when given a sentence such as *Mary fascinated Ted* is about why this is the case. Furthermore, people maintain a bias to the argument that they refer to when they are asked a causal question. In the case of *Mary fascinated Ted*, people will make the inference that there is something about Mary which caused the situation. One interpretation of this finding is that there is information in the lexical entry for interpersonal verbs, which specifies the cause. However, another interpretation – consistent with the fact that causality is fundamental to so many processes in cognition – is that the cognitive system is predisposed to seek causal relations. Hence, given minimal information about the cause of an event, a causal explanation is sought. Causal information is most informative in this context. On the other hand, one can imagine other situations where causal information is not informative (*e.g.* when the individual already knows what the cause is) and other inferences become more important.

It remains an open question whether causes are sought in all situations. If they are not then when are they sought? One plausible explanation is that they are sought when there is something abnormal about the situation (Hilton & Slugoski, 1986). However, for the implicit causality bias verbs, the situations that they describe are not always abnormal. For example, *flatter*, *like*, *protect* *etc.* do not denote abnormal situations. Nevertheless, people make causal inferences when they are presented with these verbs. The biggest question about the implicit causality verb bias that remains unanswered is what underlies the causal bias. In order to understand the implicit causality bias better more research is needed on how the implicit causality bias interacts with other sources of causal information.

Chapter 4 examined one other source of causal information and investigated the relationship between the implicit causality bias and covariation theory. According to covariation theory, cause is attributed to that thing which covaries with the event. In production, both the implicit causality bias and explicit covariation information influence causal attributions. However, contrary to classic covariation theory, where low consensus and low distinctiveness information increases NP1 attributions and high consensus and high distinctiveness information increases NP2 attributions, a more complex pattern of attributions was evident. When covariation information was congruent with the implicit causality bias then it did not effect causal attributions. This was the case when low consensus information was presented with NP1 biasing verbs and when high consensus information was presented with NP2 biasing verbs. Covariation information only appeared to have an effect on attributions when it was incongruent with the verb bias. This was the case when high consensus information was presented with NP1 biasing verbs and low consensus information was presented with NP2 biasing verbs. Once again, it appeared that causal attributions were made on the most informative cues. When there were redundant cues (*e.g.* when implicit causality and covariation information were congruent) only one of them was used to make the attribution. However, when the cues were in competition (*e.g.* when implicit causality and covariation information were incongruent) both cues were given some weight. Hence, when the causal cues were incongruent with one another the “strongest” cue was given the most weight, and it largely determined causal attributions.

A slightly different picture emerged in comprehension. Chapter 5 examined whether covariation information affected ease of comprehension. If covariation information is used automatically, then when presented with low consensus information explicit causes referring to NP1 should be easier to read than those referring to NP2. On the other hand, when presented with high consensus information explicit causes referring to NP2 should be easier to read than those referring to NP1. The hypothesised interaction was not found. Although there was some suggestive evidence that covariation information was not being overlooked completely, the overwhelming factor in influencing ease of comprehension was the verb bias.

Therefore, there is an asymmetry between production and comprehension when making causal attributions. Whereas in production covariation information had a noticeable impact on attributions, especially when the verb bias was weak, in comprehension covariation information played a negligible role. This could be because covariation information requires the participant to engage in quite a complex set of

inferences before a correct causal attribution can be made. Although the time to make these inferences is available during production, there is not sufficient time to engage in such elaborative inferences during comprehension. One could test this hypothesis in at least two different ways: first, by limiting the time that people had to produce completions and second by increasing the time that people have to make a causal inference in comprehension.

If participants had to do a speeded production task then there would be less time for them to be able to engage in elaborative inferences. Therefore, based on the above explanation, it would be hypothesised that there would be less of an effect of covariation information in the speeded production task than in the normal production task. In the comprehension task, one could increase the time that people have to make causal inferences in a number of ways. For example, instead of presenting participants with one sentence at a time, one word at a time could be presented, using a fixed-pace RSVP procedure. By doing this, one could manipulate the rate of presentation. The standard presentation rate is 300 milliseconds, plus 25 milliseconds per character. However, to encourage elaborative inferences during reading one could slow down the rate of presentation. If the neglect of covariation information was because of insufficient time then, it would be hypothesised that there would be larger influence of covariation using the above method.

However, another factor that may have affected comprehension results was the fact that the readers did not have to integrate the covariation information with the target sentence in order to understand the passage. Participants were presented with sentences such as *Few people confided in Thomas. Debbie confided in Thomas because she was trusting.* The first and second sentences are not very well connected, and there is nothing in the text to encourage participants to integrate the two. However, by adding a connective, e.g., *Few people confided in Thomas but Debbie confided in Thomas because she was trusting,* the covariation information suddenly requires more integration. Therefore, future experiments could whether covariation information affects comprehension when it is actually integrated with the rest of the text.

Finally, Chapters 6 and 7 examined exactly what cues people were using to make causal attributions when they did use covariation information. Two theories were contrasted: the frequency signalling account and the focussing account. According to the frequency signalling account if the target individual is one of a small number of people who behave in the manner described by the verb, then they are the cause; whereas if they are one of a large group of people then they are less likely to be the



cause. Therefore, the important information that covariation information provides is about the set sizes which the Person and Entity<sup>24</sup> belong to. However, according to the focussing account causal attributions are not made on the basis of set size information, but are made when the covariation information draws attention to the fact that there is a particular amount, and invites a comment on that amount. This is illustrated by the difference between *few* and *a few*: *few* indicates there is a small amount and invites the participant to think of reasons why this is the case; *a few* indicates that there is a small amount, also, but does not invite people to think why.

In Chapter 7, it was demonstrated that people did in fact use information about set size, be it absolute number or proportion, to make causal attributions. This is consistent with the frequency signalling account. However, numerical and proportional information did not affect causal attributions to the same extent as the natural language quantifiers did. Furthermore, from Chapter 6 it was quite clear the frequency signalling account was not sufficient to account for the covariation cues made in typical attribution studies. It was found that focus affected causal attributions, also. When the covariation information was congruent with the implicit causality bias then participants were affected by whether the information was inviting a reason-why inference.

This demonstrated that a modified version of covariation theory is necessary. Covariation information is used when: (1) it is informative – it is only used when it provides information that is incongruent with the verb. Moreover, covariation information influences attributions more when it is presented in natural language quantifier than in numerical or proportional information because the former includes more information about the beliefs of the writer/reader. (2) There is enough incentive/time – it appears that covariation information is not used when there is insufficient time, or need, to make elaborative inferences. This was the case in the comprehension experiments, and may also be the case with the numerical and proportional information in Chapter 7. If having exact set size information was important to the task then no doubt a larger effect of numerical and proportional information would be seen.

Overall, the evidence suggests that a more sophisticated model of causal attributions is required. This would detail when people make causal attributions, and which details the different types of causal cues (covariation information, verb bias, *etc.*) are used by people to make causal inferences.

---

<sup>24</sup> And also modality

## References

- Abelson, R. P., & Kanouse, E. E. (1966). Subjective acceptance of verbal generalizations. In S. Feldman (Ed.), *Cognitive Consistency: Motivational Antecedents and Behavioural Consequents* (pp. 171-197). New York: Academic Press.
- Ahn, W. K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, 54(3), 299-352.
- 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.
- Baayen, H. R., Piepenbrock, R., & Van Rijn, H. (1993). The CELEX lexical database (CDROM). Philadelphia: Linguistic Data Consortium, University of Pennsylvania.
- Baillargeon, R., Kotovsky, L., & Needham, A. (1995). The acquisition of physical knowledge in infancy. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal Cognition : A Multidisciplinary Debate* (pp. 79-116). Oxford: Clarendon Press.
- Barton, S. B., & Sanford, A. J. (1990). The control of attributional patterns by the focusing properties of quantifying expressions. *Journal of Semantics*, 7, 81-92.
- Breedin, S. D., Saffran, E. M., & Schwartz, M. F. (1998). Semantic factors in verb retrieval: An effect of complexity. *Brain and Language*, 63, 1-31.
- Brown, R. (1986). Linguistic relativity. In S. H. Hulse & B. F. Green Jr (Eds.), *One hundred years of psychological research in America: G. Stanley Hall and the John Hopkins tradition* (pp. 241-276). Baltimore: John Hopkins University Press.
- Brown, R., & Fish, D. (1983a). Are there universal schemas of psychological causality? *Archives de Psychologie*, 51, 145-153.
- Brown, R., & Fish, D. (1983b). The psychological causality implicit in language. *Cognition*, 14, 237-273.
- Brown, R., & Van Kleeck, M. H. (1989). Enough said: Three principles of explanation. *Journal of Personality and Social Psychology*, 57(4), 590-604.

- Burgess, C., Livesay, K., & Lund, K. (1998). Explorations in context space: Words, sentences, discourse. *Discourse Processes*, 25(2-3), 211-257.
- Cacciari, C., Carreiras, M., & Cionini, C. B. (1997). When words have two genders: Anaphor resolution for Italian functionally ambiguous words. *Journal of Memory and Language*, 37(4), 517-532.
- Caramazza, A., Grober, E., Garvey, C., & Yates, J. (1977). Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behaviour*, 16, 601-609.
- Carlson, G. N., & Tanenhaus, M. K. (1988). Thematic roles and language comprehension. *Syntax and Semantics*, 21, 263-299.
- Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Cohen, J. (1977). *Statistical Power Analysis for the Behavioral Sciences*. New York London : Academic Press
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). Psyscope: A new graphic interactive environment for designing psychology experiments. *Behavioural Research Methods, Instruments and Computers*, 25, 257-271.
- Corrigan, R. (1988). Who dun it? The influence of actor-patient animacy and type of verb in the making of causal attributions. *Journal of Memory and Language*, 27, 447-465.
- Corrigan, R. (1992). The relationship between causal attributions and judgements of the typicality of events described by sentences. *British Journal of Social Psychology*, 31(4), 351-368.
- Corrigan, R. (1993). Causal attributions to states and events described by different classes of verbs. *British Journal of Social Psychology*, 32(Pt4), 335-348.
- Corrigan, R., & Denton, P. (1996). Causal understanding as a developmental primitive. *Developmental review*, 16, 162-202.
- Corrigan, R., & Stevenson, C. (1994). Childrens causal attributions to states and events described by different classes of verbs. *Cognitive Development*, 9(2), 235-256.
- Cunningham, J., Starr, P., & Kanouse, D. (1979). Self as actor, active observer, and passive observer: Implications for causal attributions. *Journal of Personality and Social Psychology*, 37, 1146-1152.
- Dickinson, A., & Shanks, D. (1995). Instrumental action and causal representation. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal Cognition : A Multidisciplinary Debate* (pp. 5-25). Oxford: Clarendon Press.

- Dowty, D. (1991). Thematic proto-roles and argument selection. *Language*, 67(3), 547-619.
- Edwards, D., & Potter, J. (1993). Language and causation - A discursive model of description and attribution. *Psychological Review*, 100(1), 23-41.
- Ehrlich, K. (1980). Comprehension of pronouns. *Quarterly Journal of Experimental Psychology*, 32, 247-255.
- Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, 99(1), 3-19.
- Ekman, P. (1973). Cross-cultural studies of facial expression. In P. Ekman (Ed.), *Darwin and Facial Expression: A Century of Research in Review* (pp. 169-222). New York: Academic Press.
- Fiedler, K., & Semin, G. R. (1988). On the causal information conveyed by different interpersonal verbs: The role of implicit sentence context. *Social Cognition*, 6(1), 21-39.
- Fillmore, C. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory*. New York: Holt.
- Fillmore, C. J., & Langendoen, D. T. (Eds.). (1971). *Studies in linguistic semantics*. New York: Holt, Rinehart & Winston.
- Fiske, S. T., & Taylor, S. E. (1991). *Social Cognition* (2nd ed.). London: McGraw-Hill.
- Franco, F. (1987). Implicit causality in verbal descriptions of adult-infant interaction. *Bulletin of the British Psychological Society*, 40, A91-A92.
- Franco, F., & Arcuri, L. (1990). Effect of semantic valence on implicit causality of verbs. *British Journal of Social Psychology*, 29(Pt2), 161-170.
- Garnham, A., & 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., & Cruttenden, H. (1992). The role of implicit causality and gender cue in the interpretation of pronouns. *Language and Cognitive Processes*, 7(3/4), 231-255.
- Garnham, A., Traxler, M., Oakhill, J., & Gernsbacher, M. A. (1996). The locus of implicit causality effects in comprehension. *Journal of Memory and Language*, 35, 517-543.
- Garrod, S., & Sanford, A. (1977). Interpreting anaphoric relations: The integration of semantic information while reading. *Journal of Verbal Learning and Behaviour*, 16, 77-90.

- Garvey, C., & Caramazza, A. (1974). Implicit causality in verbs. *Linguistic Inquiry*, 5(459-464).
- Garvey, C., Caramazza, A., & Yates, J. (1976). Factors influencing assignment of pronoun antecedents. *Cognition*, 3(3), 227-243.
- Gernsbacher, M. A., & Hargreaves, D. J. (1988). Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language*, 27, 699-717.
- Gilovich, T., & Regan, D. T. (1986). The actor and the experiencer: Divergent patterns of causal attribution. *Social Cognition*, 4(3), 342-352.
- Gordon, P. C., Grosz, B. J., & Gilliom, L. A. (1993). Pronouns, names and the centering of attention in discourse. *Cognitive Science*, 17, 311-347.
- Gordon, P. C., & Hendrick, R. (1998). The representation and processing of coreference in discourse. *Cognitive Science*, 22, 389-424.
- Greene, S. B., & McKoon, G. (1995). Telling something we can't know - experimental approaches to verbs exhibiting implicit causality. *Psychological Science*, 6(5), 262-270.
- Grier, S. A., & McGill, A. L. (2000). How we explain depends on whom we explain: The impact of social category on the selection of causal comparisons and causal explanations. *Journal of Experimental Social Psychology*, 36(6), 545-566.
- Grober, E. H., Beardsley, W., & Caramazza, A. (1978). Parallel function strategy in pronoun assignment. *Cognition*, 6, 117-133.
- Hamilton, D. L. (1988). Causal attribution viewed from an information-processing perspective. In Bar-Tal, D., & Kruglanski, A. W. (Eds.). *The Social Psychology of Knowledge*.
- Heider, F. (1958). *The psychology of interpersonal relations*. London New York: Lawrence Erlbaum Associates.
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behaviour. *American Journal of Psychology*, 57, 243-249.
- Hewstone, M., & Jaspers, J. (1983). A re-examination of the roles of consensus, consistency and distinctiveness: Kelley's cube revisited. *British Journal of Social Psychology*, 22, 41-50.
- Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgement. *Psychological Bulletin*, 118(2), 248-271.
- Hilton, D. J., & Jaspers, J. M. F. (1987). The explanation of occurrences and non-occurrences: A test of the inductive logic model of causal attribution. *British Journal of Social Psychology*, 26, 189-201.

- Hilton, D. J., & Slugoski, B. R. (1986). Knowledge-based causal attribution: The abnormal conditions focus model. *Psychological Review*, 93(1), 75-88.
- Hilton, D. J., Smith, R. H., & Kim, S. H. (1995). Processes of causal explanation and dispositional attribution. *Journal of Personality and Social Psychology*, 68(3), 377-387.
- Hoffman, C., & Tchir, M. A. (1990). Interpersonal verbs and dispositional adjectives: The psychology of causality embodied in language. *Journal of Personality and Social Psychology*, 58(5), 765-778.
- Holtgraves, T., & Raymond, S. (1995). Implicit causality and memory: Evidence for a priming model. *Personality and social psychology bulletin*, 21(1), 5-12.
- Howell, D. C. (1992). *Statistical Methods for Psychology* (3rd ed.). Belmont, California: Duxbury Press.
- Inhoff, & Rayner. (1986). Parafoveal word-processing during eye fixations in reading - effects of word frequency. *Perception & Psychophysics*, 40(6), 431-439.
- Jackendoff, R. (1972). *Semantics in generative grammar*. Cambridge, Mass.: MIT Press.
- Jackendoff, R. (1983). *Semantics and Cognition*. Cambridge, Mass.: MIT Press.
- Jackendoff, R. (1990). *Semantic Structures*. Cambridge, Mass.: MIT Press.
- Jaeger, J. (1992). "Not by the chair of my hinny hin hin": Some general properties of slips of the tongue in young children. *Journal of Child Language*, 19, 335-366.
- Johnson-Laird, P. N., & Oatley, K. (1989). The language of emotions: An analysis of a semantic field. *Cognition and emotion*, 3(2), 81-123.
- Jones, & Nisbett . (1972). The actor and the observer: Divergent perceptions of the causes of behaviour. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins , & B. Weiner (Eds.), *Attribution: Perceiving the Causes of Behaviour* (pp. 79-94). Morristown, NJ: General Learning Press.
- Kanouse, D. E. (1972). Verbs as implicit quantifiers. *Journal of Verbal Learning and Verbal Behaviour*, 11, 141-147.
- Kasof, J., & Lee, J. Y. (1993). Implicit causality as implicit salience. *Journal of Personality and Social Psychology*, 65(5), 877-891.
- Katz, J. J., & Fodor, J. A. (1963). The structure of semantic theory. *Language*, 39, 170-210.
- Kelley, H. H. (1967). Attribution theory in social psychology. *Nebraska Symposium on Motivation*, 15, 192-238.

- Klima, E. S. (1964). Negation in English. In J. A. Fodor & J. J. Katz (Eds.), *The Structure of Language*. Englewood Cliffs, NJ: Prentice-Hall.
- Kummer, H. (1995). Causal knowledge in animals. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 26-36). Oxford: Oxford University Press.
- Lafrance, M., Brownell, H., & Hahn, E. (1997). Interpersonal verbs, gender, and implicit causality. *Social Psychology Quarterly*, 60(2), 138-152.
- Lee, J. Y., & Kasof, J. (1992). Interpersonal verbs and interpersonal experiences. *Journal of Social Psychology*, 132(6), 731-740.
- Levin, B. (1993). *English verb classes and alternations: A preliminary investigation*. Chicago: The University of Chicago Press.
- Long, D. L., & De Ley, L. (2000). Implicit causality and discourse focus: The interaction of text and reader characteristics in pronoun resolution. *Journal of Memory and Language*, 42(4), 545-570.
- MacDonald, M. C., & MacWhinney, B. (1990). Measuring inhibition and facilitation from pronouns. *Journal of Memory and Language*, 29, 469-492.
- Mackie, J. L. (1980). *The Cement of the Universe: A Study of Causation*. Oxford: Clarendon Press.
- Mannetti, L., & Degrada, E. (1991). Interpersonal verbs - implicit causality of action verbs and contextual factors. *European Journal of Social Psychology*, 21(5), 429-443.
- McArthur, L. Z. (1972). The how and what of why: Some determinants and consequences of causal attributions. *Journal of Personality and Social Psychology*, 22, 171-193.
- McArthur, L. Z., & Post, D. L. (1977). Figural emphasis and person perception. *Journal of Experimental Social Psychology*, 13, 520-535.
- McDonald, J. L., & Macwhinney, B. (1990). Implicit causality and the time course of referent activation. *Bulletin of the Psychonomic Society*, 28(6), 522.
- McDonald, J. L., & MacWhinney, B. (1995). The time course of pronoun resolution: Effects of implicit verb causality and gender. *Journal of Memory and Language*, 34, 543-566.
- McGill, A. L. (1993). Selection of a causal background: Role of expectation versus feature mutability. *Journal of personality and social psychology*, 64(5), 701-707.
- Michotte, A. (1946/1963). *The perception of causality*. London: Methuen.

- Miller, D. T., & Ross, M. (1975). Self-serving biases in the attribution of causality: Fact or fiction? *Psychological Bulletin*, 82, 213-225.
- Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and Perception*. Cambridge: Cambridge University Press.
- Moens, M., & Steedman, M. J. (1988). Temporal ontology and temporal reference. *Computational linguistics*, 14, 15-28.
- Mollaret, P. (1998). Redefining traits. *Annee Psychologique*, 98(3), 545-567.
- Moxey, L. M., & Sanford, A. J. (1993a). Prior expectation and the interpretation of natural language quantifiers. *European Journal of Cognitive Psychology*, 5, 73-91.
- Moxey, L. M., & Sanford, A. J. (1993b). *Communicating Quantities: A Psychological Perspective*. Hove: Erlbaum.
- Moxey, L. M., & Sanford, A. J. (1997). Choosing the right quantifier. In T. Givon (Ed.), *Conversation* (Vol. 34, pp. 207-231). Amsterdam: John Benjamins Publishing Company.
- Moxey, L. M., & Sanford, A. J. (2000). Communicating quantities: A review of the psycholinguistic evidence of how expressions determine perspectives. *Applied Cognitive Psychology*, 14, 237-255.
- Mullet, E., & Rivett, I. (1991). Comprehension of verbal probability expressions in children and adolescents. *Language and Comprehension*, 11, 217-225.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 289-316.
- Myers, J. L., Shinjo, M., & Duffy, S. A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language*, 26(4), 453-465.
- Natsopoulos, D., Grouios, G., Bostantzopoulou, S., Mentenopoulos, G., Katsarou, Z., & Logothetis, J. (1993). Algorithmic and heuristic strategies in comprehension of complement clauses by patients with Parkinson's-disease. *Neuropsychologia*, 31(9), 951-964.
- Newstead, S. E., & Coventry, K. R. (2000). The role of context and functionality in the interpretation of quantifiers. *European Journal of Cognitive Psychology*, 12(2), 243-259.
- Newstead, S. E., Pollard, P., & Reizbos, D. (1987). The effect of set size on the interpretation of quantifiers used in rating scales. *Applied Ergonomics*, 18, 178-182.



- Norman, D. A., Rumelhart, D. E., & Group, L. N. R. R. (1975). *Explorations in Cognition*: Freeman.
- Orvis, B. R., Cunningham, J. D., & Kelley, H. H. (1975). A closer examination of causal inference: The roles of consensus, distinctiveness and consistency information. *Journal of personality and social psychology*, 32(4), 605-616.
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning*. Urbana: University of Illinois.
- Pavlov, I. (1927). *Conditioned reflexes*. Oxford: Oxford University Press.
- Perfetti, C. A. (1998). The limits of co-occurrence: Tools and theories in language research. *Discourse processes*, 25(2&3), 363-377.
- Power, M., & Dagleish, T. (1997). *Cognition and Emotion: From Order to Disorder*. Sussex: Psychology Press.
- Pruitt, D. J., & Insko, C. A. (1980). Extension of the Kelley attribution model: The role of comparison-object consensus, target-object consensus, distinctiveness and consistency. *Journal of Personality and Social Psychology*, 39(1), 39-58.
- Pynte, J., Kennedy, A., Murray, W. S., & Courrieu, P. (1988). *The effects of Spatialisation on the processing of ambiguous pronominal reference*. Paper presented at the Proceedings of Fourth European Conference on Eye Movements.
- Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 10, pp. 174-221). New York: Academic Press.
- Rudolph, U. (1997). Implicit verb causality: Verbal schemas and covariation information. *Journal of Language and Social Psychology*, 16(2), 132-158.
- Rudolph, U., & Forsterling, F. (1997). The psychological causality implicit in verbs. *Psychological Bulletin*, 121(2), 192-218.
- Sanford, A. J., Moar, K., & Garrod, S. C. (1988). Proper names as controllers of discourse focus. *Language and Speech*, 31, 43-56.
- Schlesinger, I. M. (1992). The experiencer as agent. *Journal of Memory and Language*, 31(315-332).
- Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Science*, 4(8), 299-309.
- Semin, G. R., & Marsman, J. G. (1994). Multiple inference-inviting properties of interpersonal verbs - event instigation, dispositional inference, and implicit causality. *Journal of Personality and Social Psychology*, 67(5), 826-849.

- Semin, G. R., Rubini, M., & Fiedler, K. (1995). The answer is in the question - the effect of verb causality on locus of explanation. *Personality and Social Psychology Bulletin*, 21(8), 834-841.
- Skinner, B. F. (1953). *Science and human behaviour*. New York: Free Press London.
- Spelke, E. S., Philips, A., & Woodward, A. L. (1995). Infants' knowledge of object motion and human action. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal Cognition : A Multidisciplinary Debate* (pp. 44-78). Oxford: Clarendon Press.
- Sperber, D., Premack, D., & Premack, A. J. (1995). *Causal Cognition : A Multidisciplinary Debate*. Oxford: Clarendon Press.
- Stevenson, R. J., Crawley, R. A., & Kleinman, D. (1994). Thematic roles , focus and the representation of events. *Language and cognitive processes*, 9(4), 519-548.
- Stevenson, R. J., Knott, A., Oberlander, J., & McDonald, S. (2000). Interpreting pronouns and connectives: Interactions among focusing, thematic roles and coherence relations. *Language and Cognitive Processes*, 15(3), 225-262.
- Stewart, A. J., Pickering, M. J., & Sanford, A. J. (1998). The relationship between implicit causality and implicit consequentiality. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 1266). University of Wisconsin-Madison: Lawrence Erlbaum Associates.
- Stewart, A. J., Pickering, M. J., & Sanford, A. J. (2000). The time course of the influence of implicit causality information: Focusing versus integration accounts. *Journal of Memory and Language*, 42(3), 423-443.
- Taylor, S. E., & Fiske, S. T. (1978). Salience, attention, and attribution: Top of the head phenomena. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 11, pp. 249-288). New York: Academic Press.
- Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and importance of story events. *Journal of Memory and Language*, 24, 595-611.
- Trabasso, T., & van den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language*, 24, 612-630.
- Turnbull, W., & Slugoski, B. (1988). Conversational and linguistic processes in causal attribution. In D. J. Hilton (Ed.), *Contemporary Science and Natural Explanation: Commonsense Conceptions of Causality* (pp. 66-93). Sussex: Harvester Press Limited.

- Tversky, A., & Kahneman, D. (1986). Evidential impact of base rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgement Under Uncertainty : Heuristics and Biases* (pp. 153-162). Cambridge: Cambridge University Press.
- Van Kleeck, M. H., & Brown, R. (1988). Pitting verbal schemas against information variables in attribution. *Social Cognition*, 6(2), 89-106.
- Vorster, J. (1985). Implicit causality in language: Evidence from Afrikaans. *South African Journal of Psychology*, 15(2), 62-67.
- Wierzbicka, A. (1992). Semantics, Culture and Cognition: Universal Human Concepts in Culture-Specific Configurations. Oxford: Oxford University Press.
- Wierzbicka, A. (1996). Semantics: Primes and Universals. Oxford: Oxford University Press.
- Wilson, G., & Stevenson, R. J. (1998). *Do situational roles or thematic roles provide the best cues to pronoun assignment?* Paper presented at the New approaches to discourse anaphora: Proceedings of the second colloquium on discourse anaphora and anaphora resolution (DAARC2).

## Appendix 1

A full list of the verbs (and their thematic roles) used in Experiments 1, 2 and 3. The first column shows the verb and the second the associated thematic role (TR). For Experiments 1 and 2 the proportion of causal (C) and consequential (Q) completions are shown, as well as the proportion of subject (NP1) and object (NP2) references within each category. Experiment 3 shows the proportion of subject (NP1) and object (NP2) causal completions only, as no consequential completions were possible.

VERB	TR	EXPERIMENT 1				EXPERIMENT 2				EXPERIMENT 3	
		C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2
Answer	AE	0	0.17	0	0.04	0.25	0.13	0.04	0.17	0.17	0.58
Applaud	AE	0.04	0.71	0	0	0	0.75	0.08	0	0.42	0.5
Calm down	AE	0	0.38	0.04	0.04	0	0.54	0.13	0	0.13	0.79
Censure	AE	0.04	0.42	0	0.04	0.21	0.25	0.08	0.13	0.29	0.54
Criticise	AE	0.04	0.54	0	0.08	0	0.79	0	0	0.38	0.58
Heal	AE	0.04	0.08	0	0	0.25	0.21	0	0.13	0.58	0.38
Help	AE	0	0.33	0	0	0.08	0.54	0.04	0.04	0.29	0.58
Judge	AE	0.04	0.54	0	0.04	0.21	0.54	0.17	0	0.58	0.29
Obey	AE	0.25	0.29	0.08	0	0.38	0.25	0.13	0.13	0.42	0.58
Placate	AE	0.17	0.21	0	0	0.13	0.25	0	0.13	0.33	0.5
Praise	AE	0.04	0.79	0	0.04	0.13	0.71	0	0.04	0.21	0.75
Reassure	AE	0.17	0.33	0	0	0.13	0.33	0.04	0.04	0.25	0.58
Reprimand	AE	0.04	0.58	0	0.04	0.17	0.63	0	0.08	0.17	0.75
Accept	AP	0	0.38	0	0	0.04	0.54	0	0.08	0.46	0.46
Accuse	AP	0	1	0	0	0.13	0.67	0	0.17	0.46	0.54
Agree with	AP	0.04	0.08	0	0.04	0.21	0.17	0.04	0.08	0.33	0.5
Approach	AP	0.21	0	0.08	0.08	0.42	0.08	0.08	0	0.58	0.33
Banter with	AP	0.08	0.04	0.08	0.04	0.21	0.08	0	0	0.25	0.13
Betray	AP	0.04	0	0.04	0	0.21	0.21	0	0.29	0.67	0.29
Chase	AP	0	0.08	0.04	0	0.13	0.25	0.04	0.17	0.33	0.58
Cheat	AP	0	0.04	0	0.04	0.17	0.08	0.13	0.13	0.79	0.17
Cheer	AP	0.08	0.38	0	0	0.25	0.25	0	0	0.38	0.63

VERB	TR	EXPERIMENT 1				EXPERIMENT 2				EXPERIMENT 3	
		C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2
Cheer up	AP	0.04	0.04	0	0	0.17	0.5	0	0.04	0.21	0.67
Command	AP	0	0	0	0	0.33	0.13	0	0	0.58	0.38
Compete with	AP	0.04	0	0.08	0	0.25	0.08	0.13	0	0.5	0.17
Confide in	AP	0.38	0.17	0	0.13	0.83	0.08	0	0.04	0.46	0.42
Defy	AP	0.08	0.04	0	0	0.21	0.25	0.08	0.13	0.79	0.13
Denounce	AP	0.04	0.46	0.04	0.04	0.04	0.46	0	0.17	0.25	0.63
Deride	AP	0.08	0.38	0	0	0.13	0.58	0	0.04	0.42	0.46
Discourage	AP	0	0.83	0	0	0.13	0.67	0	0.04	0.71	0.29
Encourage	AP	0	0.58	0	0.04	0	0.67	0	0.13	0.5	0.46
Flatter	AP	0.08	0	0.04	0.13	0.29	0.17	0.04	0.25	0.79	0.21
Greet	AP	0	0.04	0	0.04	0.08	0	0.04	0.17	0.54	0.33
Harass	AP	0.21	0.08	0.08	0.13	0.17	0.21	0.04	0.25	0.75	0.25
Harm	AP	0.08	0.08	0	0.08	0.25	0.08	0.13	0.25	0.5	0.46
Hit	AP	0	0.25	0.04	0.08	0.13	0.29	0.04	0.21	0.38	0.5
Hug	AP	0.13	0.25	0.04	0.13	0.13	0.29	0.04	0.25	0.29	0.67
Hurt	AP	0.17	0.04	0	0.13	0.58	0.08	0	0.08	0.54	0.25
Insult	AP	0.04	0.17	0	0.13	0.04	0.33	0	0.17	0.54	0.38
Interrupt	AP	0.04	0.17	0	0.13	0.29	0.17	0.04	0.08	0.46	0.54
Joke with	AP	0	0.21	0.04	0.04	0.17	0.17	0	0.33	0.33	0.46
Kick	AP	0.04	0.08	0.04	0.13	0.17	0.33	0.04	0.21	0.38	0.58
Kill	AP	0.08	0.13	0.04	0.04	0.42	0.17	0.08	0	0.63	0.38
Manipulate	AP	0.33	0.17	0	0.13	0.21	0.29	0	0.08	0.58	0.38
Murder	AP	0.08	0.25	0.08	0.04	0.25	0.25	0.21	0.08	0.5	0.5
Order around	AP	0.29	0.08	0	0.25	0.21	0.38	0	0.25	0.58	0.29
Protect	AP	0.04	0.33	0.04	0	0.08	0.38	0	0	0.25	0.71
Push	AP	0	0.08	0	0.04	0.08	0.17	0.04	0.29	0.21	0.67
Restrict	AP	0.04	0.04	0	0.04	0.33	0.25	0	0.13	0.33	0.58
Rush to	AP	0.29	0.21	0.13	0.08	0.25	0.33	0.17	0.17	0.5	0.5
Slander	AP	0.04	0.21	0	0.13	0.08	0.21	0.13	0.33	0.71	0.29
Snub	AP	0.17	0.38	0.08	0	0.5	0.33	0	0.13	0.46	0.46
Stare at	AP	0.29	0.21	0	0.04	0.29	0.5	0	0.13	0.25	0.75
Stop by	AP	0.38	0.08	0.13	0	0.17	0.21	0.46	0	0.67	0.29
Support	AP	0	0.58	0	0.04	0.04	0.71	0.04	0	0.25	0.71
Tie up	AP	0.08	0.08	0.25	0	0.08	0.17	0.13	0.17	0.5	0.38
Warn	AP	0.13	0.33	0	0	0.04	0.67	0	0.04	0.42	0.46
Admire	ES	0	0.75	0	0	0	0.75	0.04	0.17	0.08	0.92
Adore	ES	0	0.54	0	0	0.08	0.46	0.13	0.13	0.13	0.83

VERB	TR	EXPERIMENT 1				EXPERIMENT 2				EXPERIMENT 3	
		C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2
Desire	ES	0.04	0.13	0	0.04	0.21	0.38	0.08	0.08	0.21	0.75
Despise	ES	0.04	0.54	0.04	0	0.13	0.54	0	0.13	0.08	0.83
Detest	ES	0.04	0.67	0	0	0.08	0.88	0	0	0	1
Dislike	ES	0	0.63	0.04	0.04	0.17	0.67	0.08	0.04	0.04	0.88
Distrust	ES	0.17	0.5	0	0	0.08	0.79	0.04	0	0.08	0.83
Dote on	ES	0.08	0.25	0.04	0	0.04	0.38	0.17	0.17	0.38	0.58
Dread	ES	0.04	0.29	0	0	0.13	0.67	0.04	0	0.13	0.79
Fear	ES	0.04	0.5	0	0	0.25	0.58	0.04	0.04	0.08	0.83
Gape at	ES	0.08	0.38	0.08	0	0.13	0.71	0	0	0.21	0.79
Hate	ES	0.04	0.46	0.08	0.04	0.04	0.54	0.13	0	0.08	0.92
Identify	ES	0	0.42	0	0	0.08	0.38	0	0	0.33	0.63
Like	ES	0	0.38	0.08	0.08	0.04	0.17	0.25	0.13	0	0.92
Loath	ES	0.04	0.5	0.04	0.04	0.04	0.67	0.08	0.04	0.04	0.83
Love	ES	0	0.29	0	0.17	0	0.29	0.08	0.21	0.17	0.79
Notice	ES	0	0.29	0	0	0	0.71	0.08	0.04	0	1
Pity	ES	0.08	0.71	0.04	0.04	0.08	0.92	0	0	0.08	0.88
Recognise	ES	0	0.04	0	0.04	0.13	0.54	0	0.04	0.38	0.58
Resent	ES	0.08	0.79	0	0	0	0.83	0	0.13	0.17	0.79
Respect	ES	0.04	0.58	0.04	0	0.04	0.63	0.17	0	0	1
See	ES	0.08	0	0.08	0	0	0.17	0.33	0.04	0.5	0.46
Trust	ES	0.08	0.08	0	0	0.21	0.21	0.17	0.29	0.33	0.67
Appal	SE	0.17	0.08	0	0	0.67	0.13	0	0.13	0.58	0.29
Appease	SE	0.08	0.04	0.04	0	0.21	0.29	0.04	0	0.58	0.38
Attract	SE	0.13	0.04	0	0	0.25	0.04	0	0.08	0.88	0.04
Daunt	SE	0.29	0.08	0	0	0.46	0.21	0	0.04	0.75	0.17
Distract	SE	0	0.08	0.08	0.13	0.17	0.21	0	0.04	0.88	0.08
Embarrass	SE	0.13	0	0	0	0.08	0.13	0	0.13	0.67	0.25
Fascinate	SE	0.21	0.04	0	0.04	0.54	0.17	0	0.08	0.83	0.17
Frighten	SE	0.25	0.08	0	0.08	0.58	0	0	0.08	0.92	0.04
Impress	SE	0.04	0	0	0	0.38	0	0	0.04	0.92	0.08
Intrigue	SE	0.29	0.08	0	0.08	0.33	0.21	0	0.08	0.71	0.29
Irritate	SE	0.21	0	0.04	0	0.5	0.13	0	0.08	0.88	0.13
Obsess	SE	0.21	0.04	0	0.17	0.29	0.17	0	0.25	0.54	0.38
Offend	SE	0.08	0.04	0	0.04	0.25	0.13	0.08	0.08	0.71	0.25
Overexcite	SE	0.08	0	0.04	0.17	0.38	0.17	0.08	0.13	0.75	0.21
Shame	SE	0	0.17	0	0.13	0.17	0.13	0.08	0.21	0.58	0.38
Shock	SE	0.17	0	0	0	0.29	0.38	0	0	0.88	0.13

VERB	TR	EXPERIMENT 1				EXPERIMENT 2				EXPERIMENT 3	
		C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2	Q NP1	Q NP2	C NP1	C NP2
Terrify	SE	0.29	0	0	0.04	0.58	0.17	0	0	0.92	0.08
Trouble	SE	0.38	0	0	0.13	0.71	0.13	0	0.04	0.71	0.25
Uplift	SE	0.17	0	0.04	0	0.21	0.38	0.08	0.04	0.67	0.33
Worry	SE	0.46	0.04	0	0.04	0.67	0.13	0	0.04	0.88	0.08

## Appendix 2

The instructions given to coders for categorising sentence completions in Experiments 2 and 3.

- If the continuation is a reason why the event in the main clause happened then code as CAUSE (C), that is, if under the circumstances had the continuation not happened then the event in the main clause would not have happened.
  - If the cause is NP1 then code as C1
  - If the cause is NP2 then code as C2
  - If the cause is neither of the above, or ambiguous then code as Ca*e.g.* Ted spoke to Mary **because he was bored.** (C1)
- If under the circumstances had the event in the main clause not happened then the continuation would not have happened, then code as CONSEQUENCE (Q)
  - If the consequence is NP1 then code as Q1
  - If the consequence is NP2 then code as Q2
  - If the consequence is neither of the above, or ambiguous then code as Qa*e.g.* Ted spoke to Mary **and so she answered him back.** (Q2)
- If the continuation describes the method by which the event in the main clause was carried out then code as METHOD (M)
  - If the method was carried out by NP1 then code as M1
  - If the method was carried out by NP2 then code as M2*e.g.* Ted spoke to Mary **by using piece of chalk and a blackboard.** (M1)
- If the continuation mentions a property of one of the NPs then code as PROPERTY (P). That is, if the continuation modifies one of the NPs then the continuation is a property of that NP.
  - If the property belongs to NP1 then code as P1
  - If the property belongs to NP2 then code as P2*e.g.* Ted spoke to Mary's **friend.** (P2)
- If the continuation mentions the location the event happened then code as LOCATION (L)  
*e.g.* Ted spoke to Mary **in the garden.** (L)
- If the continuation mentions the time the event happened then code as TIME (T)  
*e.g.* Ted spoke to Mary **at midnight.** (T)
- If the continuation mentions the manner in which the event was carried out in then code as MANNER (er). That is if the continuation modifies the verb, then the continuation is of the manner type.  
*e.g.* Ted spoke to Mary **loudly.** (er)



- If the continuation does not fall into any of these categories then code as OTHER (X)  
*e.g.* Ted spoke to Mary **and Tom spoke to Sue.** (X)

## Appendix 3

Verbs used in Experiments 4, 10, 11, 12, 13, 14 and 15

VERB	NP1 continuations	NP2 continuations	Difference scores	Bias
defied	19	3	16	NP1
cheated	19	4	15	NP1
flattered	19	5	14	NP1
harassed	18	6	12	NP1
slandered	17	7	10	NP1
betrayed	16	7	9	NP1
competed with	12	4	8	NP1
hurt	13	6	7	NP1
approached	14	8	6	NP1
killed	15	9	6	NP1
commanded	14	9	5	NP1
greeted	13	8	5	NP1
bantered with	6	3	3	neutral
tied up	12	9	3	neutral
insulted	12	10	2	neutral
confided in	11	10	1	neutral
harmed	12	11	1	neutral
encouraged	12	11	1	neutral
murdered	12	12	0	neutral
accepted	11	11	0	neutral
rushed to	12	12	0	neutral
snubbed	11	11	0	neutral
warned	10	11	-1	neutral
interrupted	11	13	-2	neutral
accused	11	14	-3	NP2
joked with	8	11	-3	NP2
hit	9	12	-3	NP2
kicked	9	14	-5	NP2
chased	8	14	-6	NP2
cheered	9	15	-6	NP2
restricted	8	14	-6	NP2
helped	7	14	-7	NP2
hugged	7	16	-9	NP2
supported	6	17	-11	NP2
pushed	5	16	-11	NP2
protected	6	17	-11	NP2

## Appendix 4

Verbs used in Experiment 5. (ES = Experiencer-Stimulus verbs; SE = Stimulus-Experiencer verbs)

VERB	Thematic roles	NP1 continuations	NP2 continuations	Difference scores	Bias
frightened	SE	22	1	21	NP1
attracted	SE	21	1	20	NP1
terrified	SE	22	2	20	NP1
impressed	SE	22	2	20	NP1
worried	SE	21	2	19	NP1
distracted	SE	21	2	19	NP1
shocked	SE	21	3	18	NP1
irritated	SE	21	3	18	NP1
fascinated	SE	20	4	16	NP1
appalled	SE	19	4	15	NP1
daunted	SE	18	4	14	NP1
overexcited	SE	18	5	13	NP1
troubled	SE	17	6	11	NP1
offended	SE	17	6	11	NP1
embarrassed	SE	16	6	10	NP1
intrigued	SE	17	7	10	NP1
loved	ES	4	19	-15	NP2
resented	ES	4	19	-15	NP2
dreaded	ES	3	19	-16	NP2
adored	ES	3	20	-17	NP2
despised	ES	2	20	-18	NP2
distrusted	ES	2	20	-18	NP2
feared	ES	2	20	-18	NP2
loathed	ES	1	20	-19	NP2
pitied	ES	2	21	-19	NP2
disliked	ES	1	21	-20	NP2
admired	ES	2	22	-20	NP2
hated	ES	2	22	-20	NP2
liked	ES	0	22	-22	NP2
detested	ES	0	24	-24	NP2
noticed	ES	0	24	-24	NP2
respected	ES	0	24	-24	NP2

## Appendix 5

Materials used in Experiments 6, 7 and 8. Experiment 8 included a context sentence presenting consensus information.

Items 1-12 contain NP1 biasing verbs; items 13-24 neutral verbs and items 25-36 NP2 biasing verbs. For each item (a) contains an explicit cause which refers to NP1; while (b) contains an ending referring to NP1. This means that for items 1-12, (a) presents a causality congruent ending while (b) presents a causality incongruent ending. On the other hand, for items 25-36, (b) presents a causality congruent ending while (a) presents an incongruent ending.

- 1a Samantha defied Noel because she was rebellious.
- 1b Samantha defied Noel because he was authoritarian.
- 2a Ray flattered Rachel because he was charming.
- 2b Ray flattered Rachel because she was beautiful.
- 3a Mick cheated Carol because he was ruthless.
- 3b Mick cheated Carol because she was gullible.
- 4a Rose betrayed Peter because she was very angry with him.
- 4b Rose betrayed Peter because he was very horrid to her.
- 5a Edward competed with Tracy because he was ambitious.
- 5b Edward competed with Tracy because she was successful.
- 6a Kevin harassed Eileen because he was a stalker.
- 6b Kevin harassed Eileen because she was an actress.
- 7a Claire slandered Gordon because she was jealous.
- 7b Claire slandered Gordon because he was arrogant.
- 8a Charles hurt Mary he was very thoughtless.
- 8b Charles hurt Mary because she was very sensitive.
- 9a Susan approached Jack because she was friendly.
- 9b Susan approached Jack because he was lonely.
- 10a Ted killed Jane because he was annoyed.
- 10b Ted killed Jane because she was annoying.
- 11a David commanded Theresa because he was the chief.
- 11b David commanded Theresa because she was the slave.
- 12a Nicola greeted Jonathan because she was sociable.
- 12b Nicola greeted Jonathan because he was alone.
- 13a Elizabeth bantered with Tony because she was extroverted.
- 13b Elizabeth bantered with Tony because he was opinionated.
- 14a Sarah insulted Andrew because she was rather rude.
- 14b Sarah insulted Andrew because he was rather stupid.
- 15a Simon tied up Anna because he was a sadist.
- 15b Simon tied up Anna because she was a masochist.
- 16a Debbie confided in Thomas because she was trusting.
- 16b Debbie confided in Thomas because he was trustworthy.
- 17a Ian harmed Mary because he was very cruel.

- 17b Ian harmed Amy because she was very irritating.
- 18a Sue murdered Luke because she was a psychopath.
- 18b Sue murdered Luke because he was a millionaire.
- 19a Pamela accepted Philip because she was welcoming.
- 19b Pamela accepted Philip because he was repentant.
- 20a Rita encouraged Paul because she was concerned.
- 20b Rita encouraged Paul because he was insecure.
- 21a Grant rushed to Sandra because he felt anxious.
- 21b Grant rushed to Sandra because she felt hurt.
- 22a Caroline snubbed Trevor because she was very snooty.
- 22b Caroline snubbed Trevor because he was very dirty.
- 23a Nancy interrupted Roy because she was rather impatient.
- 23b Nancy interrupted Roy because he was rather long-winded.
- 24a Richard accused Suzanne because he was a policeman.
- 24b Richard accused Suzanne because she was a thief.
- 25a Bruce joked with Meg because he was a comic.
- 25b Bruce joked with Meg because she was a giggler.
- 26a Jean warned Fraser because she was very worried.
- 26b Jean warned Fraser because he was very careless.
- 27a Douglas hit Kathy because he was aggressive.
- 27b Douglas hit Kathy because she was naughty.
- 28a Rob chased Beth because he was very athletic.
- 28b Rob chased Beth because she was very attractive.
- 29a Liz cheered Stephen because she was excited.
- 29b Liz cheered Stephen because he was winning.
- 30a Gail kicked Matthew because she was infuriated.
- 30b Gail kicked Matthew because he was infuriating.
- 31a Graeme restricted Ann because he was a guard.
- 31b Graeme restricted Ann because she was a prisoner.
- 32a Joyce helped Patrick because she was kind.
- 32b Joyce helped Patrick because he was needy.
- 33a Emily hugged Daniel because she was caring.
- 33b Emily hugged Daniel because he was upset.
- 34a Craig supported Diana because he was strong.
- 34b Craig supported Diana because she was weak.
- 35a Sebastian pushed Florence he was a bully.
- 35b Sebastian pushed Florence because she was a pain.
- 36a Terrence protected Ruth because he was a bodyguard.
- 36b Terrence protected Ruth because she was a pop star.

## Appendix 6

Materials for Experiment 9: Items 1-12 contain NP1 biasing verbs, while items 13-24 contain NP2 biasing verbs. For each item, (i) is the context sentence, and is presented first. The context sentence presents either low consensus information (a) or high consensus information (b). Two different quantifiers were used for each type of information. Following (i), pressing the appropriate button presented the target sentence (ii). The target sentence presented an explicit cause, which referred to NP1 or NP2. The target sentence was presented with one set of proper names in the first half of the experiment, but another set (the names in brackets) the second time. Another press of the button presents (iii). This is an anaphor-resolving question. The question was presented with the two proper names from the target sentence. The names appeared either on the left-hand side of the screen, or on the right-hand side of the screen.

- 1ia (Very) Few people apologised to Kirsty (Joanne).
- 1ib Many (Most) people apologised to Kirsty (Joanne).
- 1iia Edward (Daniel) apologised to Kirsty (Joanne) because he had been behaving selfishly.
- 1iib Edward (Daniel) apologised to Kirsty (Joanne) because she didn't deserve the criticism.
- 1iiia Who had been behaving selfishly?
- 1iiib Who didn't deserve the criticism?
  
- 2ia (Very) Few people confessed to Johnny (Philip).
- 2ib Many (Most) people confessed to Johnny (Philip).
- 2iia Daphne (Tracey) confessed to Johnny (Philip) because she had stolen the money.
- 2iib Daphne (Tracey) confessed to Johnny (Philip) because he would not be judgmental.
- 2iiia Who had stolen the money?
- 2iiib Who would not be judgmental?
  
- 3ia (Very) Few people fascinated Paula (Lydia).
- 3ib Many (Most) people fascinated Paula (Lydia).
- 3iia Kevin (Barry) fascinated Paula (Lydia) because he performed magic tricks.
- 3iib Kevin (Barry) fascinated Paula (Lydia) because she was easily entertained.
- 3iiia Who performed magic tricks?
- 3iiib Who was easily entertained?

- 4ia (Very) Few people infuriated Tom (Ben).  
 4ib Many (Most) people infuriated Tom (Ben).  
 4iia Jan (Ann) infuriated Tom (Ben) because she had broken the promise.  
 4iib Jan (Ann) infuriated Tom (Ben) because he hated being deceived.  
 4iia Who had broken the promise?  
 4iib Who hated being deceived?
- 5ia (Very) Few people disappointed Beth (Rose).  
 5ib Many (Most) people disappointed Beth (Rose).  
 5iia Bill (John) disappointed Beth (Rose) because he failed to appear.  
 5iib Bill (John) disappointed Beth (Rose) because she had high standards.  
 5iia Who failed to appear?  
 5iib Who had high standards?
- 6ia (Very) Few people troubled Roy (Tim).  
 6ib Many (Most) people troubled Roy (Tim).  
 6iia Kay (Amy) troubled Roy (Tim) because she was starting to behave rather strangely.  
 6iib Kay (Amy) troubled Roy (Tim) because he hated seeing others feeling very sad.  
 6iia Who was starting to behave rather strangely?  
 6iib Who hated seeing others feeling very sad?
- 7ia (Very) Few people inspired Sarah (Carol).  
 7ib Many (Most) people inspired Sarah (Carol).  
 7iia David (Henry) inspired Sarah (Carol) because he had managed to beat the odds.  
 7iib David (Henry) inspired Sarah (Carol) because she needed someone to look up to.  
 7iia Who had managed to beat the odds?  
 7iib Who needed someone to look up to?
- 8ia (Very) Few people telephoned Joel (Luke).  
 8ib Many (Most) people telephoned Joel (Luke).  
 8iia Nora (Nancy) telephoned Joel (Luke) because she wanted to ask a favour.  
 8iib Nora (Nancy) telephoned Joel (Luke) because he wouldn't remember to call.  
 8iia Who wanted to ask a favour?  
 8iib Who wouldn't remember to call?
- 9ia (Very) Few people amused Katie (Donna).  
 9ib Many (Most) people amused Katie (Donna).  
 9iia Keith (James) amused Katie (Donna) because he performed hilarious impressions.  
 9iib Keith (James) amused Katie (Donna) because she was very easily entertained.  
 9iia Who performed hilarious impressions?  
 9iib Who was very easily entertained?

- 10ia (Very) Few people concerned David (Terry).
- 10ib Many (Most) people concerned David (Terry).
- 10iia Trudy (Susan) concerned David (Terry) because she was starting to behave erratically.
- 10iib Trudy (Susan) concerned David (Terry) because he hated seeing friends in trouble.
- 10iiia Who was starting to behave erratically?
- 10iiib Who hated seeing friends in trouble?
  
- 11ia (Very) Few people amazed Kathleen (Florence).
- 11ib Many (Most) people amazed Kathleen (Florence).
- 11iia Joseph (Jonathan) amazed Kathleen (Florence) because he passed the exam.
- 11iib Joseph (Jonathan) amazed Kathleen (Florence) because she was easily impressed.
- 11iiia Who passed the exam?
- 11iiib Who was easily impressed?
  
- 12ia (Very) Few people called Robert (Brian).
- 12ib Many (Most) people called Robert (Brian).
- 12iia Susan (Diana) called Robert (Brian) because she had found the telephone number.
- 12iib Susan (Diana) called Robert (Brian) because he couldn't make outgoing calls.
- 12iiia Who had found the telephone number?
- 12iiib Who couldn't make outgoing calls?
  
- 13ia (Very) Few people congratulated Gail (Rita).
- 13ib Many (Most) people congratulated Gail (Rita).
- 13iia Tony (Fred) congratulated Gail (Rita) because he was very impressed.
- 13iib Tony (Fred) congratulated Gail (Rita) because she had won the championship.
- 13iiia Who was very impressed?
- 13iiib Who had won the championship?
  
- 14ia (Very) Few people appreciated Douglas (Richard).
- 14ib Many (Most) people appreciated Douglas (Richard).
- 14iia Gillian (Kathryn) appreciated Douglas (Richard) because she needed the extra help.
- 14iib Gillian (Kathryn) appreciated Douglas (Richard) because he had offered to help.
- 14iiia Who needed the extra help?
- 14iiib Who had offered to help?
  
- 15ia (Very) Few people detested Hilary (Phoebe).
- 15ib Many (Most) people detested Hilary (Phoebe).
- 15iia Norman (Trevor) detested Hilary (Phoebe) because he hated being taken advantage of.
- 15iib Norman (Trevor) detested Hilary (Phoebe) because she was completely unreliable.
- 15iiia Who hated being taken advantage of?
- 15iiib Who was completely unreliable?



- 16ia (Very) Few people liked Alan (Tony).  
 16ib Most people liked Alan (Tony).  
 16iia Mona (Mary) liked Alan (Tony) because she was made to feel quite at home.  
 16iib Mona (Mary) liked Alan (Tony) because he was full of incredibly helpful advice.  
 16iiia Who was made to feel quite at home?  
 16iiib Who was full of incredibly helpful advice?
- 17ia (Very) Few people despised Kimberly (Helena).  
 17ib Many (Most) people despised Kimberly (Helena).  
 17iia Sebastian (Gregory) despised Kimberly (Helena) because he had felt very let down.  
 17iib Sebastian (Gregory) despised Kimberly (Helena) because she seemed to lie constantly.  
 17iiia Who had felt very let down?  
 17iiib Who seemed to lie constantly?
- 18ia (Very) Few people thanked Derek (Harry).  
 18ib Many (Most) people thanked Derek (Harry).  
 18iia Eliza (Cathy) thanked Derek (Harry) because she had appreciated the present.  
 18iib Eliza (Cathy) thanked Derek (Harry) because he had brought the present.  
 18iiia Who had appreciated the present?  
 18iiib Who had brought the present?
- 19ia (Very) Few people loathed Jill (Lisa).  
 19ib Many (Most) people loathed Jill (Lisa).  
 19iia Neil (Mick) loathed Jill (Lisa) because he was starting to feel upstaged.  
 19iib Neil (Mick) loathed Jill (Lisa) because she had very little integrity.  
 19iiia Who was starting to feel upstaged?  
 19iiib Who had very little integrity?
- 20ia (Very) Few people praised Ian (Bob).  
 20ib Many (Most) people praised Ian (Bob).  
 20iia Kay (Sue) praised Ian (Bob) because she was impressed by the project.  
 20iib Kay (Sue) praised Ian (Bob) because he behaved very courageously.  
 20iiia Who was impressed by the project?  
 20iiib Who behaved very courageously?
- 21ia (Very) Few people scolded Joy (Liz).  
 21ib Many (Most) people scolded Joy (Liz).  
 21iia Guy (Ray) scolded Joy (Liz) because he was aware of the potential danger.  
 21iib Guy (Ray) scolded Joy (Liz) because she had damaged the mahogany table.  
 21iiia Who was aware of the potential danger?  
 21iiib Who had damaged the mahogany table?

- 22ia (Very) Few people noticed Linda (Anna).  
 22ib Many (Most) people noticed Linda (Anna).  
 22iia Simon (Billy) noticed Linda (Anna) because he was always exceedingly observant.  
 22iib Simon (Billy) noticed Linda (Anna) because she wore a remarkably colourful dress.  
 22iia Who was always exceedingly observant?  
 22iib Who wore a remarkably colourful dress?
- 23ia (Very) Few people punished Roger (Paul).  
 23ib Many (Most) people punished Roger (Paul).  
 23iia Chloe (Gemma) punished Roger (Paul) because she had had enough.  
 23iib Chloe (Gemma) punished Roger (Paul) because he had been very trying.  
 23iia Who had had enough?  
 23iib Who had been very trying?
- 24ia (Very) Few people admired Mike (Grant).  
 24ib Many (Most) people admired Mike (Grant).  
 24iia Irene (Emily) admired Mike (Grant) because she needed a role model.  
 24iib Irene (Emily) admired Mike (Grant) because he was very motivated.  
 24iia Who needed a role model?  
 24iib Who was very motivated?

