

Implicit Consequentiality

Andrew J Stewart (andrew@psy.gla.ac.uk)
Martin J Pickering (martin@psy.gla.ac.uk)
Anthony J Sanford (tony@psy.gla.ac.uk)

Department of Psychology
University of Glasgow, Glasgow, G12, Scotland, UK

Abstract

This paper examines the way in which high level semantic information influences the production and comprehension of pronouns. It reports a new type of verb semantic processing bias. We examine the effects of this bias on language comprehension.

Introduction

Experimental work on pronoun resolution places particular focus on the influence of a type of high level information referred to as implicit causality (Garnham, Traxler, Oakhill & Gernsbacher, 1996; Garvey & Caramazza, 1974). Implicit causality is a semantic bias associated with particular verbs which affects clausal integration.

- (1) John blamed Bill because he broke the window.
- (2) John fascinated Bill because he was interesting.

In Example (1) there is a preference to interpret the pronoun as referring to the character 'Bill' (Noun Phrase 2) whereas in Example (2) there is a preference to interpret the pronoun as referring to the character 'John' (Noun Phrase 1). The implicit causality congruency effect is the finding that it takes longer to read a sentence when it contains an ending inconsistent with the verb bias than an ending consistent with the verb bias (Caramazza, Grober, Garvey & Yates, 1977) (Examples (3) and (4) below).

- (3) John blamed Bill because he hated Bill.
- (4) John fascinated Bill because he was interested by John.

However, implicit causality is not the only semantic bias associated with the process of clausal integration. In this paper we examine a new type of bias which we refer to as implicit consequentiality. In the same way that events have causes, they also have consequences. Causes of an event occurring at t arise at some point prior to t . Conversely, consequences arise at some point following t . In the same way that different types of events can have different types of causes, different types of events can also have different types of consequences. Recall Example (1). The act of NP1 ('John') blaming NP2 ('Bill') is usually preceded by the character occupying the NP2 ('Bill') position doing something to bring about this outcome. This locus of cause gives rise to the implicit causality

effect. Example (5) below describes an event and the consequences following. Intuitively, the act of NP1 ('John') annoying NP2 ('Bill') is also usually followed by some consequence affecting the character occupying the NP2 ('Bill') position. We propose this is an example of an implicit consequentiality bias.

In Example (5), what follows from 'John' annoying 'Bill' is 'Bill' complaining about it.

- (5) Because John annoyed Bill, he complained to the art teacher.

But there is no restriction as to which character may experience the consequences of the event. In Example (6), the same subordinate clause is followed by a consequence focusing on the first mentioned character, 'John'.

- (6) Because John annoyed Bill, he was punished by the art teacher.

We expect to find an implicit consequentiality congruency effect similar to the implicit causality congruency effect. In other words, we expect reading time to the disambiguating region, i.e. the region following the pronoun, in (6) to be read more slowly than the equivalent region in (5) if it is the case that the verb 'annoy' possesses a strong bias to focus on the character occupying the second Noun Phrase.

Finding evidence for biases arising as a result of knowledge about consequences of events is interesting for a number of reasons. Traditional explanations of the implicit causality phenomenon have suggested that verbs possessing such biases have a tag encoded at the verb root corresponding to that particular verb's bias (Garvey & Caramazza, 1974). Finding that an individual verb possesses different types of biases depending on the context in which that verb occurs weakens this manner of explanation as it requires as many tags to encode biases as there are biases. Additionally, indication of an influence on processing of general knowledge about the temporal context in which events in the world typically occur immediately raises the question of how (and when) such knowledge comes to bear a processing influence. An alternative to an account proposing that biases are encoded as lexical tags is an account focusing on a more general mechanism informed by knowledge about typical events in the world.

One problem pervasive in the literature examining implicit causality is one of a lack of employment of relevant experimental controls. It is well documented that sentences containing continuations following the pronoun that ultimately resolve with reference to the character inconsistent with the bias take longer to read (Caramazza et al, 1977). However, it may be the case that such sentences simply describe less plausible events. In Stewart, Pickering and Sanford (1998) we addressed this question and reported that implicit causality congruency effects are independent of plausibility effects. In our examination of implicit consequentiality below we employ plausibility controls of the same type as we have used previously.

So, there are two main goals motivating this paper. The first is one of examining whether biases other than the implicit causality bias documented in the literature exist. The second is concerned with the time course of the influence of the semantic bias on reading. This is motivated by general issues concerning the structure of the cognitive architecture and focuses on the time course of biases influencing processing. Much work on other aspects of language processing has been guided by Fodor's modularity thesis (Fodor, 1983). The basic proposal of modularity is that initial processing proceeds in a restricted manner informed by only a subset of all the information that will ultimately come to bear a weight on the final processing outcome. In parsing it is considered that some degree of processing informed purely by syntactic information precedes processing informed by semantic factors (Frazier, 1987). Drawing an explicit analogy between the parsing literature and the processing of pronouns, we propose that if the part of the language system associated with interpreting pronouns behaves in a modular or restricted manner we would expect to find a delayed influence of high level semantic factors such as implicit consequentiality. Under a restricted processing account initial processing will proceed with reference only to low level information. Before an answer to this question can be broached, however, we need to satisfy ourselves that implicit consequentiality does exist as a bias which influences language processing. Experiment 1 following examines subjects sensitivity to implicit consequentiality information when engaged in a language production task.

Experiment 1

Our first experiment was performed to insure that implicit consequentiality information does influence language processing in the manner that we suggest. Specifically, we propose that some verbs focus on the character occupying the first Noun Phrase, while others lead to a focus on the character occupying the second Noun Phrase.

Method

The task was an off-line production task requiring subjects to complete fragments of the sort 'Because John loved Mary,...'. The two characters were differentiated by gender to allow us to unambiguously score whether pronominal reference was to the first or second mentioned character.

Subjects

Twenty-two English speaking subjects participated.

Stimuli

Experimental booklets were constructed containing 110 sentence fragments. We examined the same 50 verbs whose implicit causality biases had been examined by Stewart (1998). Each verb appeared in a sentence of the form :

Because Sue VERBED John, ...

Names were varied. Each booklet contained 60 filler materials of a similar structure to the above fragment. Each verb appeared only once. The gender of the two characters in each sentence were different.

Procedure

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

Results

The endings for each experimental sentence were scored on the basis of whether the main clause initially contained an anaphor referring to the first or second noun phrase. Other types of continuations were scored as 'Others'. The table below contains the 16 most strongly biasing verbs (the full list of 50 verbs can be obtained from the authors).

Table 1 : Completion type for 16 biasing verbs

Verb	NP1	NP2	Others
dreaded	17	1	4
appreciated	17	3	2
adored	21	1	0
admired	22	0	0
loathed	19	2	1
resented	18	2	2
despised	18	2	2
liked	19	1	2
aggravated	0	20	2
annoyed	6	15	1
scolded	0	21	1
confused	0	20	2
punished	1	21	0
flattered	0	19	3
thanked	0	22	0
congratulated	2	19	1

Discussion

The data gathered by Experiment 1 certainly indicate that subjects are sensitive to implicit consequentiality biases in language production. A particular subset of verbs possess implicit consequentiality biases which focus on the character occupying the first Noun Phrase and a particular subset possess biases focusing on the character occupying the second Noun Phrase. Our initial question of whether we could find semantic biases which influenced production

in the same way that implicit causality influences production has been answered.

We now turn to the second question addressing the issue of the manner in which such biases exert an influence on processing.

Experiment 2

Experiment 1 provided evidence for implicit consequentiality biases influencing language production. In Experiment 2 we examine the way in which such biases influence language comprehension. Assuming that the language system has developed to take advantage of the statistical regularities in the linguistic environment, we expect to find evidence that comprehenders are sensitive to implicit consequentiality biases in a manner similar to language producers. The bias acts as a potentially informative cue with respect to which character the pronoun refers. However, the bias is not as strong as a gender cue which effectively rules out pronominal reference to a potential antecedent differing in gender from a gender marked pronoun.

We examine the behaviour of implicit consequentiality information in the context of the presence and absence of a gender cue. In the presence of a gender cue, we propose the relative influence of implicit consequentiality will be reduced as the pronominal antecedent can be identified using gender information (Example 5).

(5) Because John punished Sue, she started to cry.

However, when such a gender contrast is lacking we argue the degree to which implicit consequentiality influences processing will be maximised as it provides a strong cue as to which character is referred to by the pronoun (Example 6).

(6) Because John punished Bill, he started to cry.

So, we have defined the conditions under which we expect to find implicit consequentiality information exerting a processing influence. Precisely when this influence will occur depends on when the reader interprets the pronoun. If the reader tries to interpret the pronoun as soon as it is encountered we expect to find evidence for the implicit consequentiality congruency effect on the reading time data associated with the fragment ending in the pronoun. If its influence is delayed we expect to find the implicit consequentiality congruency effect on reading time second fragment of the sentence.

Method

Subjects

Thirty two English speaking subjects participated.

Stimuli

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

NP1 or NP2 biasing. Referent could either be first or second mentioned character.

There were 32 sets of experimental materials.

Unambiguous pronoun / NP1 verb / Referent character1
Because Harold dreaded Joanne, he steadfastly refused to go back to school.

Unambiguous pronoun / NP1 verb / Referent character2
Because Joanne dreaded Justin, he was told to try acting less aggressively.

Ambiguous pronoun / NP1 verb / Referent character1
Because Harold dreaded Justin, he steadfastly refused to go back to school.

Ambiguous pronoun / NP1 verb / Referent character2
Because Harold dreaded Justin, he was told to try acting less aggressively.

There were also 96 filler items. There were 32 experimental items created using the 16 verbs which Experiment 1 indicated exhibited strong implicit consequentiality biases. Each subject saw each verb twice. Pretensing ensured that regardless of whether the sentence resolved in a manner consistent or inconsistent with the implicit consequentiality bias, the plausibility of the described events were equivalent. This control ensures that a simply plausibility difference between each condition cannot account for any reading time differences we may find. The materials were also carefully controlled for length. The experiment was divided into two halves with a break halfway through. The first occurrence of each verb was in the first half of the study and the second occurrence in the second half.

Procedure

The experiment was run on an Apple Macintosh computer using the PsyScope experimental software (Cohen, MacWhinney, Flatt & Provost, 1993). A button box was connected to the computer which recorded subjects' responses with millisecond accuracy. A fixation point appeared on the left hand side of the screen. Subjects pressed a button on the button box and the fixation point was followed by presentation of the first fragment of the sentence. This fragment included the pronoun, e.g. 'Because John blamed Mary, she'. Once subjects read the sentence they pressed the button again. The second fragment then appeared. Pressing the button again resulted in a question appearing after every trial. A break of a minimum duration of 30 seconds occurred halfway through each experiment.

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

Results

First Fragment Analysis (e.g. 'Because John annoyed Bill, he')

We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors. No main effects or interactions approached significance. We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors for just the unambiguous pronouns. No main effects or interactions approached significance ($F_s < 2$).

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

Verb / Referent	Ambig Pro	Unambig Pro
NP1 / Char1	2332	2159
NP1 / Char2	2174	2201
NP2 / Char1	2391	2230
NP2 / Char2	2334	2355

Second Fragment Analysis (e.g. 'complained to the maths teacher.')

We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors. We found a main effect of ambiguity ($F(1,31)=24.64$, $p < 0.0001$; $F(1,30)=29.37$, $p < 0.0001$), reflecting faster reading times following unambiguous reference. We found a main effect of verb bias ($F(1,31)=23.09$, $p < 0.0001$; $F(1,30)=3.94$, $p < 0.06$), reflecting faster reading time following fragments containing NP2 biasing verbs. We found an interaction between ambiguity and verb bias that was marginal by items ($F(1,31)=4.55$, $p < 0.05$; $F(1,30)=2.96$, $p < 0.1$). We found an ambiguity x referent interaction that was significant by subjects only ($F(1,31)=4.87$, $p < 0.05$; $F(1,30)=2.96$, $p < 0.1$). We also found a verb bias x referent interaction which again was significant by subjects only ($F(1,31)=5.39$, $p < 0.05$; $F(1,30) < 1$). The 3-way interaction of verb bias x referent x ambiguity was significant by both subjects and items ($F(1,31)=4.26$, $p < 0.05$; $F(1,30)=4.50$, $p < 0.05$). We examine this 3-way interaction in detail below.

Table 3 : Reading times for second fragment, subject means for each condition with all times in msec.

Verb / Referent	Ambig Pro	Unambig Pro
NP1 / Char1	2210	2208
NP1 / Char2	2566	2055
NP2 / Char1	2355	1809
NP2 / Char2	2241	1789

We explored the nature of the 3-way interaction by performing analyses on the Ambiguous and Unambiguous

Pronoun conditions separately. We shall firstly focus on the Ambiguous Pronoun conditions.

Ambiguous Pronoun Conditions

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors. We found a significant verb bias x referent interaction that was marginal by items ($F(1,31)=7.56$, $p < 0.01$; $F(1,30)=3.09$, $p < 0.1$).

Unambiguous Pronoun Conditions

We performed 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors. We found a main effect of verb bias ($F(1,31)=23.74$, $p < 0.0001$; $F(1,30)=6.09$, $p < 0.05$) corresponding to faster reading time following an NP2 biasing verb. There was no evidence of a verb bias x referent interaction (both $F_s < 1$).

We can interpret the 3-way interaction as reflecting a difference in the implicit consequentiality congruency effect between the Ambiguous and Unambiguous Pronoun conditions. We found a congruency effect in the Ambiguous Pronoun conditions (marginal by items) but failed to find one in the Unambiguous Pronoun conditions.

Question Response Time Analysis

We performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent) ANOVAs with both subjects and items as random factors. We found a main effect of ambiguity ($F(1,31)=55.83$, $p < 0.0001$; $F(1,30)=49.48$, $p < 0.0001$) indicating faster question response times following a sentence containing unambiguous reference. We found a main effect of verb bias ($F(1,31)=8.57$, $p < 0.01$; $F(1,30)=6.29$, $p < 0.05$), indicating faster response times following a sentence containing an NP2 biasing verb. We also found a main effect of referent that was significant by subjects only ($F(1,31)=5.98$, $p < 0.05$; $F(1,30) < 2$) reflecting faster response times following reference to the first mentioned character. We found an ambiguity x referent interaction ($F(1,31)=12.26$, $p < 0.005$; $F(1,30)=9.39$; $p < 0.005$) and an ambiguity x verb bias x referent interaction significant by subjects only ($F(1,31)=4.43$, $p < 0.05$; $F(1,30)=2.33$, $p < 0.14$).

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

Verb / Referent	Ambig Pro	Unambig Pro
NP1 / Char1	2323 (80.5%)	2083 (91.3%)
NP1 / Char2	2910 (50.4%)	1824 (92.9%)
NP2 / Char1	2229 (90.5%)	1730 (96.0%)
NP2 / Char2	2508 (74.0%)	1727 (93.7%)

Question Response Accuracy Analysis

For the question response accuracy data we also performed 2 (Unambiguous pronoun vs Ambiguous pronoun) x 2 (NP1 verb vs NP2 verb) x 2 (Char1 vs Char2 referent)

ANOVAs for both subjects and items as random factors. We found a main effect of ambiguity ($F(1,31)=103.51$, $p<0.0001$; $F(1,30)=51.79$, $p<0.0001$) reflecting greater response accuracy following a sentence containing unambiguous reference. We found a main effect of verb bias ($F(1,31)=26.78$, $p<0.0001$; $F(1,30)=8.41$, $p<0.01$) reflecting greater response accuracy following a sentence containing an NP2 biasing verb. We also found a main effect of referent ($F(1,31)=34.74$, $p<0.0001$; $F(1,30)=21.25$, $p<0.0001$) reflecting greater response accuracy following a sentence containing reference to the first mentioned character. The ambiguity x verb bias interaction was significant ($F(1,31)=14.72$, $p<0.001$; $F(1,30)=6.56$, $p<0.05$). The ambiguity x referent interaction was also significant ($F(1,31)=24.04$, $p<0.0001$; $F(1,30)=16.27$, $p<0.0005$). The 3-way interaction of ambiguity x verb bias x referent was significant only by subjects ($F(1,31)=4.90$, $p<0.05$; $F(1,30)=2.40$, $p<0.14$).

Discussion

The lack of an interaction between referent and verb bias for the unambiguous pronoun conditions in the fragment 1 reading time data indicates that implicit consequentiality does not influence processing as soon as a pronoun is read. Either processing of the pronoun is delayed until information sufficient for antecedent identification is encountered or an initial commitment to a possible antecedent informed by low level factors is made.

Our 3-way interaction found in the fragment 2 reading time data indicates that the point at which implicit consequentiality first exerts an influence of processing is during integration. Whether such an influence occurs is a function of the presence or absence of additional information pertinent to the goal of interpreting the pronoun. The 3-way interaction corresponds to a difference in the nature of the implicit consequentiality congruency effect between the Unambiguous and Ambiguous Pronoun conditions. We found evidence for a congruency effect for the Ambiguous Pronoun conditions (marginal on F2s). No hint of a congruency effect was found in the case of the Unambiguous Pronoun conditions. The most obvious explanation for our 3-way interaction is that under circumstances where a pronoun's antecedent can be identified using information at a level lower than that of semantics (e.g. gender information), identification will proceed informed solely by this morphosyntactic information. When such unequivocal information is available to the reader it seems reasonable that other factors that would be informative in the absence of such a cue are not called upon to aid processing of the anaphor. However, in cases where there is not a restriction as to which antecedent is appropriate the language processor takes advantage of what other cues are available in the input. Under circumstances of this type, the implicit consequentiality bias is used by the system.

The question response time and question accuracy data indicate that subjects responded more slowly and with reduced accuracy to those questions following sentences which contained continuations referring to the second mentioned character inconsistent with the NP1 biasing

verb. We didn't find equivalent difficulty for the incongruent conditions where reference was to the first mentioned character in the context of an NP2 biasing verb. We suggest this is due to the first mention privilege which results in the first mentioned character always being easy to access (Gernsbacher & Hargreaves, 1988; Neath & Knoedler, 1994).

General Discussion

The basic question we posed at the outset of this paper was whether we could find verb bias effects comparable to implicit causality effects but corresponding to implied consequences rather than implied causes. We have been successful in answering this. The production data from Experiment 1 indicates that subjects are sensitive to the consequences of events described by particular verbs and use this information to focus on one participant rather than the other. In language production implicit consequentiality exerts an influence in a manner similar to implicit causality.

The second question we raised focused on the point in time at which implicit consequentiality information exerts an influence on processing. Consider Example (7) which contains an ending inconsistent with the verb bias :

(7) Because Michael annoyed Kathryn, he was punished by the art teacher.

Our production experiment indicated that the verb 'annoy' contains an implicit consequential bias focusing on the second character (i.e. 'Kathryn'). However, the lack of evidence of readers detecting the anomaly in reading the first fragment (up to and including the pronoun) in Experiment 2 indicates that implicit consequentiality isn't used by the system to inform processing as soon as a pronoun is encountered. When the pronoun is encountered it indicates that the sentence is going to continue in a manner inconsistent with the verb bias. If verb bias information was being used at this point we would have expected to have found evidence for an implicit consequentiality congruency effect in the fragment 1 reading time data. There was no evidence for this. The earliest point at which we found the congruency effect was in the pattern of data associated with reading the second fragment (i.e. following the pronoun). With respect to the time course of the influence of this information, finding a delayed influence is consistent with the literature examining the manner in which implicit causality influences processing (Garnham et al, 1996; Stewart, 1998). This seems to suggest that either initial processing of a pronoun proceeds in a restricted manner informed solely by low level factors or processing is delayed.

The 3-way interaction suggests that implicit consequentiality information only exerts an influence when additional cues that would unambiguously identify the pronominal referent are lacking. The implicit consequentiality congruency effect was present in the ambiguous pronoun conditions (although marginal by items) but absent in the unambiguous pronoun conditions. When gender information which would otherwise have

uniquely identified the pronominal antecedent is lacking, implicit consequentiality is a strong cue which the system employs to facilitate processing. When such information is present, implicit consequentiality information is not used by the system. When implicit consequentiality information is used, a persistence of the effect in both the question answering time and response accuracy data suggests that not only may the relative incoherence present in sentences containing continuations which clash with a verbs bias lead to a processing difficulty, but the lack of coherence may also carry over to influence the way in which the reader represents the described event. Responses querying a representation lacking strong coherence are slower and exhibit reduced accuracy compared to those querying a representation possessing a stronger cohesive structure. So, we propose that both reading and representational problems arise as a result of a reader encountering sentences containing information inconsistent with a particular verb's implicit consequential bias.

The pattern of production data in Experiment 1 is strongly similar to the set of data found in an equivalent study examining the way in which implicit causality information is used in language production (Stewart, 1998). Comparing the verbs in Experiment 1 with those verbs reported in Stewart (1998) it is clear that a particular verb's implicit causality bias cannot be used to predict that verb's implicit consequentiality bias. There is no obvious relationship between the two sets of biases. We strongly believe that biases arising as a result of focusing on either causes or consequences of particular events should only be seen as two examples drawn from a larger set of possible biases. In this paper we have demonstrated the influence of an implicit consequentiality bias both in language production and comprehension. At its most general we can define a processing bias as anything that directs the language producer's or comprehender's focus of attention. An account proposing that such biases are encoded as lexical tags appears inadequate to account for a much more pervasive focus effect. We believe that only through reference to general knowledge and readers' awareness of the typical temporal contexts in which events occur can an explanation capable of accounting for general processing biases be forthcoming.

Acknowledgements

This research was supported by a UK BBSRC studentship awarded to the first author.

References

- Caramazza, A., Grober, E., Garvey, C. & Yates, J. (1977) Comprehension of Anaphoric Pronouns. *Journal of Verbal Learning and Verbal Behaviour*, 16, 601-609.
- 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.
- Fodor, J.A. (1983) *The Modularity of Mind*. Cambridge, MA: Bradford.

- Frazier, L. (1987) Sentence processing: a tutorial review. In M. Coltheart (ed.) *Attention and Performance XII: The psychology of reading*. Hillsdale, N.J.: Erlbaum.
- Garnham, A., Traxler, M.J., Oakhill, J. and Gernsbacher, M.A. (1996) The locus of implicit causality effects in comprehension. *Journal of Memory and Language*, 4, 517-543.
- Garvey, C. and Caramazza, A. (1974) Implicit causality in verbs. *Linguistic Inquiry*, 5, 459-464.
- Gernsbacher, M.A. & Hargreaves, D. (1988) Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language*, 27, 699-717.
- Neath, I & Knodler, A.J. (1994) Distinctiveness and serial position effects in recognition and sentence processing. *Journal of Memory and Language*, 33, 776-795.
- Stewart, A.J. (1998) The time course of the influence of implicit causality information on resolving anaphors. Unpublished PhD thesis, University of Glasgow.
- Stewart, A.J., Pickering, M.J. & Sanford, A.J. (1998) The processing consequences of implicit causality. *Paper presented at : The Experimental Psychology Society, London meeting, January, 1998.*