Published in final edited form as:

Lang Cogn Neurosci. 2015 ; 30(6): 716–734. doi:10.1080/23273798.2015.1008524.

The Causes and Consequences Explicit in Verbs

Joshua K. Hartshorne, **Timothy J. O'Donnell**, and **Joshua B. Tenenbaum**

Department of Brain & Cognitive Sciences, Massachusetts Institute of Technology

Abstract

Interpretation of a pronoun in one clause can be systematically affected by the verb in the previous clause. Compare *Archibald angered Bartholomew because he*… (he=Archibald) with *Archibald criticized Bartholomew because he*… (he=Bartholomew). While it is clear that meaning plays a critical role, it is unclear whether that meaning is directly encoded in the verb or, alternatively, inferred from world knowledge. We report evidence favoring the former account. We elicited pronoun biases for 502 verbs from seven Levin verb classes in two discourse contexts (implicit causality and implicit consequentiality), showing that in both contexts, verb class reliably predicts pronoun bias. These results confirm and extend recent findings about implicit causality and represent the first such study for implicit consequentiality. We discuss these findings in the context of recent work in semantics, and also develop a new, probabilistic generative account of pronoun interpretation.

In the early 1970s, Catherine Garvey and Alfonso Caramazza made an intriguing discovery: By changing the verb in one clause, they could radically change the interpretation of a pronoun in a subsequent clause (Garvey & Caramazza, 1974; Garvey, Caramazza, & Yates, 1974).

(1) **a.** Archibald angered Bartholomew because he was reckless.

b. Archibald criticized Bartholomew because he was reckless.

Most individuals interpret *he* as referring to Archibald in (1a) and Bartholomew in (1b).

Subsequent investigation revealed that this effect of the verb is independent of the content of the second clause:

(2) **a.** Archibald angered Bartholomew because he…

b. Archibald criticized Bartholomew because he…

Importantly, this effect is a bias that can be overcome if it conflicts with material that comes after the pronoun (*Archibald angered Bartholomew because he is irritable*), albeit with a significant processing cost (Garvey et al., 1974).

Numerous verbs have been identified that bias listeners to resolve the pronoun to the previous subject (e.g., *anger, delight, dare, trick*), whereas many others bias listeners to resolve the pronoun to the previous object (e.g., *criticize, judge, love, hate*). Garvey and Caramazza, who dubbed this phenomenon "implicit causality," took it to indicate listeners'

beliefs about causation: Listeners interpret the pronouns as they do because they understand Archibald to be the implicit cause of *Archibald angered Bartholomew* but understand Bartholomew to be the implicit cause of *Archibald criticized Bartholomew*.[1]

Implicit causality is not the only phenomenon in which pronoun interpretation is biased. In sentences where the second clause describes a consequence of the first, a similar effect, known as "implicit consequentiality," arises (Au, 1986; Crinean & Garnham, 2006; Pickering & Majid, 2007; Stevenson, Crawley, & Kleinman, 1994; Stewart, Pickering, & Sanford 1998):

(3) **a.** Because Archibald angered Bartholomew, he… [he=Bartholomew]

b. Because Archibald liked Bartholomew, he… [he=Archibald]

Again, some verbs bias pronoun interpretation towards the object (e.g., *anger, delight*) and others towards the subject (e.g., *like, hate*). Here, the relevant semantic distinction appears to be that one of the participants in the event is understood to have been more affected by the event, and thus is more likely to feature in consequences of the event. Note that the direction of a verb's bias is not necessarily the same as for implicit causality.

Although implicit causality and consequentiality were initially described as pronoun interpretation biases, recent empirical and theoretical results suggest that these biases are manifestations of a broader expectation about who will be mentioned next in some discourse (cf. Arnold, 2001; Arnold, Brown-Schmidt, & Trueswell, 2007; Kehler, 2002; Kehler, Kertz, Rohde, & Elman, 2008).[2] If the next noun phrase is a pronoun, this expectation will color the interpretation of that pronoun (pronoun-specific interpretation strategies may still apply). Support for this position comes from studies where participants provide continuations to sentences truncated before the pronoun (*Archibald angered Bartholomew because*…), and researchers record who is mentioned in the continuation. The likelihood that the continuation provided by participants will mention a particular character correlates at ceiling rates with pronoun interpretation biases for equivalent sentences (Garvey et al., 1974; Hartshorne, 2014; Kehler et al., 2008). We will use the more general term *re-mention bias* to refer collectively to both production and interpretation biases, whether involving pronouns or other referring expressions (cf. Hartshorne, 2014).

Note that while implicit causality and consequentiality are the most frequently discussed re-mention biases – and are the focus of this paper – they are likely just two examples of a broader phenomenon: There is some indication that verbs also have systematic effects in sentences connected by *and*, *but*, and other connectives (Erlich, 1980; Koornneef & Sanders, 2013; Featherstone & Sturt, 2010; Rigalleau, Guerry, & Granjon, in press; Stevenson, Knott, Oberlander, & McDonald, 2000). As research goes forward, it will be important to determine which results generalize to the broader class of phenomena.

---

[1] *Note that "implicit causality" is also used to refer to a task in which people make inferences about various attributes of individuals based on events they have participated in (Brown & Fish, 1983a). Though initially thought to be related, subsequent research has found little or no relationship between the two kinds of implicit causality (Hartshorne, 2014). All discussion of implicit causality in the present paper pertains to the original Garvey and Caramazza phenomenon.*

[2] An alternative is that these biases are in fact learned heuristics derived from the statistics of pronoun use itself, which are used to predict the likely reference of a pronoun (Crawley et al., 1990; Fletcher, 1984). We return to this account in the General Discussion.

World Knowledge or Language?

Most researchers agree that intuitions about the causes and effects of events drive implicit causality and consequentiality biases. An open question – and the focus of the present study – is what information underlies these intuitions.

Many researchers have argued that these intuitions must be partly or entirely inferred from general knowledge about the typical causes and effects of different kinds of events (Brown & Fish, 1983a, 1983b; Corrigan, 2001, 2002, 2003; Pickering & Majid, 2007; Semin & Fiedler, 1991). These theories draw on a distinction between what is literally entailed by a sentence and what may be inferred based on additional, extra-linguistic knowledge. For instance, consider *Archibald was born two hundred years ago*. This sentence literally conveys only that two hundred years have passed since Archibald's birth. Most listeners will also conclude that Archibald is no longer alive, but this follows from our beliefs about normal human lifespans, not from the sentence itself. An optimist who holds out hope for Archibald's continued vitality would be accused of misunderstanding the world, not the sentence.

Advocates of the *world knowledge* approach to re-mention biases argue that *Archibald angered Bartholomew* does not entail anything about the cause of the anger. Rather, listeners must infer the likely cause based on what they have learned about typical causes and explanations of anger. Thus, on the world knowledge account, re-mention biases are a probe into people's knowledge and reasoning about the world. As such, implicit causality has been used to investigate adults' and children's understanding of the causes of events (Au, 1986), the stability of these beliefs across cultures (Brown & Fish, 1983a), and people's expectations about gender roles (Ferstl, Garnham, & Manouilidou, 2011; Goikoetxea, Pascual, & Acha, 2008; Mannetti & de Grada, 1991).

The alternative *semantic structure account* is that causes and consequences are not implicit but are actually part of the literal meaning of the verb (Arnold, 2001; Brown & Fish, 1983b; Crinean & Garnham, 2006; Garvey & Caramazza, 1974; Hartshorne & Snedeker, 2013; Rudolph & Forsterling, 1997; Stevenson et al., 2000).[3] Garvey and Caramazza (1974) suggested that verbs mark their subject or object (or neither) as the cause of the event. Subsequent researchers tried to reduce this "implicit cause" feature to aspects of verb lexical semantics, such as thematic roles (see especially Brown & Fish, 1983b; Crinean & Garnham, 2006). An action verb (*kick*, *paint*, *break*, *throw*) involves an AGENT (the subject) effecting some change on the PATIENT (the object). Because AGENTS are by definition causal actors and PATIENTS by definition suffer some consequence as the result of the event, such verbs should be subject-biased in implicit causality and object-biased in implicit consequentiality. Thus, on this account, re-mention biases are not probes into people's knowledge about the world in a language-independent manner but rather are probes into the semantic representations underlying basic linguistic processes. As such, the causality and

[3]As indicated in the citations, Brown & Fish (1983b) have been influential in the development of both positions and can be read as supporting either one.

consequentiality, etc., in verbs is not so much *implicit* as *explicit* (hence the title of this paper).

Thus while the *world knowledge* account predicts that changing facts about the world changes the IC bias of a verb, on the *semantic structure* account, facts about the world are relevant only in that they cause speakers to coin and use verbs that encode specific information about causality and affectedness.

Two lines of evidence have been used to distinguish these accounts. Until recently, both favored the world knowledge account. First, preliminary evidence suggested that re-mention biases were affected by not just the verb but also knowledge about the actors (i.e., Archibald's and Bartholomew's genders, occupations, relative social status, etc.), supporting the conclusion that re-mention biases are calculated over a rich representation of the event that incorporates substantial knowledge about the world (cf. Pickering & Majid, 2007). However, the preliminary evidence that these additional factors modulated re-mention biases largely failed to hold up under more systematic investigation (Goikoetxea et al., 2008; Hartshorne, 2014; see also Ferstl et al., 2011).

Second, proponents of the semantic structure account have long had difficulty in finding a semantic characterization of verbs that correctly predicts re-mention biases (for review, see Hartshorne & Snedeker, 2013). For instance, on early accounts, agent-patient verbs were predicted to be subject-biased in implicit causality sentences, whereas in fact many are object-biased (e.g., *criticize*). This lead researchers to posit a new semantic role (EVOCATOR), which was circularly defined as a PATIENT that nonetheless attracts implicit causality bias (Au, 1986; Rudolph & Forsterling, 1997). The failure to formulate an accurate predictive theory has been seen as a challenge to the semantic structure account. Here, too, more recent evidence has begun challenge this conclusion, as we discuss in the following section.

Before discussing this evidence, we note one potential liability for the world knowledge account. Numerous studies have now shown that at least in some cases implicit causality bias can affect pronoun interpretation within about half a second of encountering the pronoun, if not earlier (Cozijn, Commandeur, Vonk, & Noordman, 2011; Featherstone & Sturt, 2010; Koornneef & Sanders, 2012; Koornneef & van Berkum, 2006; Pyykkonen & Jarvikivi, 2010). This raises the question of whether this is sufficient time for complex inferences based on unconstrained amounts of world knowledge. In contrast, it has long been established that listeners can use the lexical semantics of verbs to predict upcoming words within a few hundred milliseconds of encountering the verb (Altmann & Kamide, 1999).[4]

Finer-Grained Semantic Representations

Hartshorne and Snedeker (2013) – henceforth H&S – provided new evidence that semantic structure drives implicit causality. They noted that contemporary work in verb semantics has found it necessary to invoke much richer, finer-grained semantic representations than those considered in the re-mention literature. For instance, Crinean and Garnham (2006) invoke

[4]We thank Andrew Stewart for pointing this out.

only 5 semantic roles, whereas contemporary semantic role theories may have several dozen (Kipper, Korhonen, Ryant, & Palmer, 2006; Schuler, 2005; for review, see Levin & Rappaport Hovav, 2005). Many theorists have also argued that treating semantic roles as primitives (as is done in the re-mention literature) results in too brittle a theory, and they have argued for more articulated representations (for review, see Levin & Rappaport Hovav, 2005). Thus, since semantic structure accounts have invoked primitive semantic roles that are too coarse-grained, it is perhaps not surprising that they over-generalized. H&S also noted that all contemporary theories of verb semantics incorporate notions of causation and affectedness as core components of meaning; it would be surprising if these did not play a role in implicit causality and consequentiality, respectively.

How to best characterize verb meanings remains an area of active research (Levin and Rappaport Hovav, 2005). H&S abstracted away from specific proposals by focusing on the verb classes provided in the comprehensive and authoritative classification of Levin (1993), as modified and extended in VerbNet (Kipper et al, 2006; Schuler, 2005). Levin verb classes result from categorizing verbs according to the syntactic frames in which they can appear: While *break*, *roll*, *hit*, and *push* can all be used in transitive frames (*Archibald broke/ rolled/hit/pushed the vase*), only the first two can be used intransitively (*The vase broke/ rolled/*hit/*pushed*). Substantial evidence has amassed indicating that which verbs can appear in which syntactic frames is largely or entirely a function of core semantic features, including causation and affectedness (Ambridge, Pine, Rowland, Jones, & Clark, 2009; Croft, 2012; Goldberg, 2003; Jackendoff, 1990; Levin & Rappaport Hovav, 2005; Pinker, 1989; Tenny & Pustejovsky, 2000; *inter alia*).[5]

Thus, H&S asked whether verbs in the same syntactic class would share the same implicit causality bias. If so, that suggests that the same underlying semantic structure that drives Levin verb classes also drives re-mention biases. H&S focused on five verb classes, finding just such a pattern (Figure 1). Moreover, which classes were biased in which direction matched previous suggestions as to how causality is (or is not) encoded by verbs in those classes.

Motivation for Current Study

The experiment presented in this paper has three main goals. First, H&S showed that, at least in some cases, implicit causality bias is systematically predictable from Levin/VerbNet verb class, as is predicted by the semantic structure account. Below, we test whether this finding extends to implicit consequentiality. There is some preliminary evidence that it may: Stewart and colleagues (1998) elicited result biases for 49 verbs (reported in Crinean & Garnham, 2006), most of which are members of VerbNet classes 31.1 (*amuse*), 31.2 (*admire*), or 33 (*acclaim*). All of the class 31.1 verbs (N=11) were object-biased, all of the class 31.2 verbs (N=18) were subject-biased, and all but one of the class 33 verbs (N=9) were object-biased, with the one exception (*honor*) being only slightly subject-biased.

[5]For instance, the reason that both *Agnes broke the vase* and *The vase broke* are grammatical but *Beatrice hit the vase* is grammatical while **The vase hit* is not is that *break* describes an externally caused event, whereas *hit* does not (cf. Levin & Rappaport Hovav, 2005).

Second, of the five syntactic classes investigated by H&S, only two were tested exhaustively, with only a sample of verbs tested in the other three classes. In this study, we exhaustively test implicit causality and consequentiality biases for seven verb classes – including the five for which H&S investigated implicit causality bias.

Finally, as noted by both Rudolph and Forsterling (1997) and H&S, there has been a tendency in the re-mention literature to focus on a small set of about four dozen verbs drawn from Garvey et al. (1974) and Brown and Fish (1983b), resulting in theories being overfit to this narrow sample. In the case of implicit causality, this issue has lately been ameliorated by the publication of several large datasets involving hundreds of verbs (Ferstl et al., 2011; Goikoetxea et al., 2008; Hartshorne & Snedeker, 2013). Nevertheless, this remains an issue for implicit consequentiality. To address this issue, we report implicit consequentiality biases for 502 verbs.

Experiment

We tested implicit causality and consequentiality biases for a comprehensive sample of verbs in each of seven syntactic classes (Table 1). These were the five classes tested by Hartshorne and Snedeker (2013), plus two additional classes (31.3 and 31.4), which consist of the emotion verbs[6] that take indirect objects rather than direct objects (classes 31.1 and 31.2). Classes 31.1 (*amuse*), 31.2 (*admire*), and 33 (*acclaim*) are three classes of verbs frequently discussed in the re-mention bias literature: experiencer-object emotion verbs, experiencer-subject emotion verbs, and judgment verbs (cf. Crinean & Garnham, 2006; Hartshorne & Snedeker, 2013; Rudolph & Forsterling, 1997).

Note that our present focus is whether the class of each verb predicts its re-mention bias. In the General Discussion, we discuss the results in terms of current theories of verb semantics.

Method

Subjects

1,638 native English-speaking volunteers (ages 18–81, M=32, SD=14; 991 female) who reported no history of dyslexia or psychiatric disorders and who were aged 18 to 81 completed the experiment through the Web portal gameswithwords.org. Children under 18 were excluded because they potentially do not know low-frequency words. Two participants aged 96 and 100 were excluded for having implausible ages. An additional 201 participants were excluded for missing more than one of the control questions (see below). Varying the exclusion criteria had little effect on the results.

Materials and Procedure

An archived version of the experiment may be viewed at www.gameswithwords.org/PronounSleuth. Each participant judged fourteen implicit causality sentences (4) and

[6]In the psycholinguistic literature, emotion verbs have often been grouped together with propositional attitude verbs (*think, believe*) and education verbs (*teach, learn*). Collectively, these verbs are known as “psych verbs”. Because these verbs appear in distinct verb classes, we discuss them separately. Moreover, H&S found little evidence that these different types of psych verbs pattern similarly with regards to re-mention biases.

fourteen implicit consequentiality sentences (5). For each participant, the order of the sentence types was randomized.

(4) Sally VERBed Mary because she daxed.

(5) Because Sally VERBed Mary, she daxed.

Participants were asked "Who do you think daxed?" and asked to choose the subject (*Sally*) or the object (*Mary*). In order to account for any bias to choose the name based on position (left or right), the order of the clickable options was randomized on each trial.

The names of the characters (*Sally*, *Mary*) and the novel verb (*daxed*, *gorped*) were sampled without replacement from a set of 70 common and unambiguously female names chosen from the (USA) Social Security Administration database and 32 novel verbs, respectively. Participants were told that some real words had been replaced with novel words (e.g., *daxed*) in order to make the task more challenging; this manipulation was intended to ensure that sentence content after the pronoun could have little effect on pronoun interpretation (*cf* Hartshorne & Snedeker, 2013; Hartshorne, Sudo & Uruwashi, 2013). The use of randomly-chosen common names had a similar purpose.

Each participant was thus tested on 28 verbs, which were chosen randomly without replacement from the total set of 502. These represented nearly all the verbs in seven VerbNet classes (see Table 1), excluding those which can appear in another class as a simple transitive with two animate arguments and thus are polysemous. One class 31.1 verb (*dumbfounded*) was misspelled and thus is excluded from analyses. Examples of each class are given in Table 1.

Because we aimed for an exhaustive test of the semantic-structure hypothesis, we included all relevant verbs. However, we note that some of the verbs were only marginally acceptable in the transitive construction (*Archibald coaxed Bill*). Because we presented the verbs in transitive sentences, this lower acceptability could result in participants having weaker intuitions about these sentences, increasing noise. Note that this should work against our hypothesis, attenuating any relationship between verb class re-mention bias. Any evidence we find of such a relationship would be that much more compelling.[7]

Four catch trials were included where the pronoun was disambiguated by gender (e.g., *Sally thanked John because he daxed*). Participants who made more than one error on these catch trials were excluded.

Results

Description of analyses

Each verb was judged in an average of 40 implicit causality and 40 implicit consequentiality sentences (Range=24–69, SD=7). Because generalization across items is of primary interest, and because the nature of the data collection does not readily permit by-subjects analyses, analyses below are conducted by item only.

[7]We thank Jennifer Arnold for raising this point.

Although re-mention biases are usually reported in terms of the percentage of participants resolving the pronoun to the previous subject, the use of this non-linear scale distorts effect sizes, violates the assumptions of standard analysis techniques (e.g., ANOVA, regression), and generally complicates interpretation.[8] Thus, we linearized the results with the empirical logit transformation (Haldane, 1955). As a result, positive values indicate a subject bias, negative values indicate an object bias, and zero indicates no bias. Absolute numerical results for all 502 verbs are listed the appendix.

Implicit Consequentiality

Across verbs, we observed a wide range of implicit consequentiality biases, with a slight overall object bias (Figure 2). However, as shown in Figure 3, these biases were a systematic function of verb class. All seven classes showed a significant bias in two-tailed t-tests (ps<.05), though the result for class 31.4 (which consists of only three verbs) would not have survived correction for multiple comparisons ($t(2)=4.8$, $p=.04$). The two experiencer-subject classes (31.2 and 31.3) were subject-biased. The other five classes were object-biased. These findings replicate and extend Stewart and colleagues' (1998) observations that class 31.1 verbs are more object-biased in implicit consequentiality than are classes 31.2 and 33.

We conducted pairwise comparisons of classes. Only comparisons that are significant below the Sidak-correction alpha of 0.002 (adjusted for 21 comparisons) are considered significant. As such, the two experiencer-subject classes (31.2, and 31.3) were not significantly different from one another ($t(97)=2.8$, $p=.006$), but both were significantly different from all other classes (ps<.002) with the exception of the under-powered comparisons with classes 31.3 and 31.4 ($t(62)=2.0$, $p=.05$). The other five classes did not differ significantly from one another, with the exception of classes 31.1 and 45.4 ($t(283)=3.6$, $p=.0003$).[9]

Implicit Causality—Across all verbs, we observe a wide range of biases, again with a slight overall object bias (Figure 4). As shown in Figure 5, these biases were a systematic function of verb class. The overall distribution of results for classes 31.1 (*amuse*), 31.2 (*admire*), 33 (*acclaim*), 45.4 (*age*), and 59 (*coax*) replicate what was observed by H&S, shifted in the direction of an object bias. We consider the source of this object shift in the Discussion. Oblique emotion verbs in classes 31.3 and 31.4 – which were not tested by H&S – behaved like their transitive counterparts (classes 31.2 and 31.1, respectively).

All classes showed significant biases (ps<.001) with the exception of 31.1 and 31.4 (ps>.1). The only non-significant pairwise comparisons between classes, after Sidak correction for 21 comparisons (alpha=.002), are as follows: 31.1 & 31.4 ($t(191)<1$), 31.4 & 59 ($t(42)=1.4$, $p=.17$), 31.2 & 33 ($t(109)<1$), and 45.4 & 59 ($t(134)=2.7$, $p=.007$).

[8] A scale is linear if a change of d units has equal value at any point in the scale. Percentages violate this rule: The difference between 50% and 51% (d=1%) is less meaningful than the difference between 99% and 100% (d=1%). This complicates comparison of results, since large numerical differences can actually be "smaller" than small numerical differences. For demonstration and discussion, see Jaeger (2008).

[9] If not correcting for multiple comparisons, classes 45.4 and 59 are also significantly different from one another ($t(134)=2.5$, $p=.01$).

Combined Analysis

We followed up the above analyses by quantifying how much information verb class provides about the bias of a given verb. We fit the raw data to a hierarchical logistic regression with discourse type (causality or consequentiality) as a fixed effect, and with random intercepts and slopes for each verb. The standard deviation of the random intercepts (1.18) and slopes (1.62) quantifies uncertainty about verb bias if all one knows is the discourse type of the sentence. Adding a fixed effect of verb class greatly reduced this uncertainty (SDs = 0.70, 0.79, respectively; *p*s < .001), indicating that verb class is highly informative about implicit causality and consequentiality bias.[10]

Discussion

We replicated and extended Hartshorne & Snedeker's (2013) finding that implicit causality bias is a systematic function of Levin verb class, and we extended these findings to implicit consequentiality. Importantly, although Levin verb classes are defined syntactically, we do not interpret this as an effect of syntax *per se*: Rather, both syntactic class and re-mention biases are likely a function of whether and how verbs specify causality, affectedness, and other semantic features (see footnote [3]; Levin 1993; Levin & Rappaport Hovav, 2005).

These results are predicted by the semantic structure account, on which the semantic information responsible for re-mention biases is read off the linguistic signal itself, rather than being inferred from general knowledge about the world. On this view, implicit causality and consequentiality are quite explicit in language.

The world knowledge account was formulated to explain how people infer the causes and consequences of linguistically-encoded events (e.g., *Archibald angered Bartholomew*). However, linguistic meaning appears to be sufficient to explain much of the phenomenon, and it is not currently clear whether world knowledge plays any additional role. Further development of the world knowledge account would require meeting three challenges. First, one would need to explain why listeners infer causality and affectedness from world knowledge rather than simply using the information about causality and affectedness conveyed by the verb. Second, one would need to actually show that people's (nonlinguistic) beliefs about the likely causes and results actually predict re-mention bias, a core prediction that surprisingly has not yet been tested. Finally, one would need to show that world knowledge-based inferences could be calculated quickly enough to match the speed with which re-mention biases appear online (cf. Koornneef & van Berkum, 2006, *inter alia*).

Note that despite the role of verb semantics, re-mention biases remain a pragmatic inference, albeit a semantically-derived one. For convenience, researchers often refer to *the* cause of an event, though in fact any event has multiple causes and thus can be explained in multiple ways; at best, listeners can only make reasonable guesses as to what cause is most likely to be discussed (cf. Pickering & Majid, 2007).[11] The fact that listeners expect explanations to

[10]We assessed significance with a permutation analysis: We randomly reassigned verbs to verb class with the constraint that the number of verbs in each class remain the same, and then refit the model. We repeated this process 1,000 times. In all 1,000 iterations, the standard deviation of the random intercepts and slopes was much greater than the true model, never dropping below 1.15 and 1.58, respectively.

refer to the cause highlighted by the verb and consequences to refer to the entity marked as affected by the verb remains an inference and is sometimes incorrect (Caramazza et al., 1977).

Our conclusions here might seem a disappointment to those who see re-mention biases as an example of complex world knowledge inferences rapidly guiding sentence processing. However, this does not make re-mention biases less interesting. To the contrary, re-mention biases highlight the communicative power of language: Rather than requiring listeners to infer causes and consequences using general reasoning skills, language specifies (some of the) causes and consequences of events. An interesting question for future research is whether speakers strategically choose which verb to use to describe an event based on what aspects of event structure they wish to highlight.

In the remainder of this section, we address three issues. First, we address the fact that we observed stronger object biases in implicit causality in the present experiment than did H&S. Second, we discuss the re-mention biases observed for different verb classes in the context of prior semantics research. Finally, we address the question of why and how listeners make use of semantics to guide re-mention biases.

The Object Shift

Overall, in the present study the implicit causality biases were shifted in the direction of the object relative to the results in H&S. Although this unexpected result is orthogonal to our primary question (*are re-mention biases a function of verb class?*), we nonetheless considered what might have driven this result. Our methods differed from those of H&S in two salient respects. H&S a) use present tense and b) concluded with a novel noun (*Sally frightens Mary because she is a dax*). However, because this resulted in somewhat unnatural implicit consequentiality sentences, in the present study we used past tense and concluded with a novel verb (*Sally frightened Mary because she daxed*).

In a follow-up experiment, we ruled out a role for tense. Forty-eight native English speakers (26 female; 18–58 y.o., M=33, SD=11) were presented with 20 randomly-selected monosemic verbs from class 59 in explanation discourses truncated after the connective *because* (*John cheated Mary because*…) and were asked to decide whether the next word should be *he* or *she*.[12]

As Figure 6 shows, tense had little effect: In contrast with our results above but similar to those of H&S, the overall results were numerically subject-biased regardless of tense. We confirmed this result with a binomial mixed effects model, which found no effect of tense (*Wald's* z<1).[13] Although the numeric subject bias did not reach significance (*Wald's z*=1.5,

[11]Causes of an event (*Archibald angered Bartholomew*) include both the power of the agent to realize some effect (causing anger) and the liability of the patient to be affected (the ability to be angry) (cf. White, 1989). Moreover, there are proximate causes (*Archibald punching Bartholomew in the nose*) and prior causes (*Archibald's troubled upbringing, which made him violent*).

[12]The study was conducted on Amazon Mechanical Turk and participants received monetary compensation. An additional five participants were excluded for answering all four unambiguous filler trials incorrectly, nine were excluded for failing to answer every question, and three were excluded for not being native English speakers. Whether the subject was male and the object was female or *vice versa* was counter-balanced within and between participants. Two orders of stimuli were used, one of which was the reverse of the other. Half the participants saw the verbs in present tense, half in past tense.

p=.14), verbs in this task were significantly more subject-biased than in the main experiment, whether treating each tense separately or combining across them (*p*s<.00001).

Since tense cannot explain the object shift, we suspect that the culprit was our use of a novel verb in the sentence continuation, rather than a novel noun. It is long established that content after the pronoun can result in reinterpretation of the pronoun (*Archibald angered Bartholomew because he was reckless/irritable*), though with some processing cost (Caramazza et al., 1977). Because the focus in the re-mention literature has been the early re-mention biases rather than these post-pronoun revision processes, most studies are designed to minimize the effect of content after the pronoun. This is true of both our study and that of H&S: The novel nouns (H&S) and verbs (present study) were meant to provide content-free continuations that would not affect pronoun interpretation and thus would provide a clean assay of the verb- and connective-driven re-mention biases.

However, explanations involving predicate nouns (*because he is a dax*) and explanations involving verbs (*because he daxed*) are not semantically equivalent. The former suggest a stable property of one of the event-participants (*he is a swimmer/psychologist/shouter/*etc.), whereas the latter suggest an event or state (*he shouted/fell/died/*etc.). The data suggest that knowing that the explanation involves a stable trait did not affect people's intuitions about who the explanation referred to; implicit causality biases elicited using novel nouns correlate at ceiling rates with those elicited in sentence continuation experiments in which there is no material after the pronoun (Hartshorne, 2014). In contrast, it appears that people are much more likely to ascribe explanations involving an action or state to object of the antecedent. Further research is necessary to determine why this is the case.

Kinds of Causes & Kinds of Explanations

On most semantic analyses (e.g., Kipper et al., 2006), verbs in classes 31.1 (*amuse*), 31.4 (*appeal to*), 33 (*acclaim*), 45.4 (*age*), and 59 (*coax*) mark their objects as affected, explaining why these verbs are object-biased in implicit consequentiality. In contrast, the subject of the verbs in classes 31.2 (*admire*) and 31.3 (*agonize over*) is the experiencer of a mental state and thus most affected by the situation, accounting for their implicit consequentiality subject biases.

Likewise, several previous semantic analyses have suggested that the verbs in classes 31.1 (*amuse*) and 31.4 (*appeal to*) highlight the causal role of their subjects whereas the verbs in classes 31.2 (*admire*), 31.3 (*agonize over*), and 33 highlight the causal role of their object, consistent with the observed implicit causality biases, modulo the object shift discussed in the previous section (cf. Kipper et al., 2006; McKoon, Greene, & Ratcliff, 1993).

Interpreting the implicit causality findings for the verbs in classes 45.4 (*age*) and 59 (*coax*) is more complex, in part because these verbs have not been extensively studied. Intuitively, the subjects of these verbs are causal agents, and indeed the semantic structures provided by VerbNet highlights the causal role of their subjects. Consistent with this analysis, H&S find

[13]Participants and verbs were random effects, while tense and whether the male character was the subject or object were fixed effects (the latter was also non-significant, *Wald's z*=1.65, *p*=.10). Maximal random effects structure was used.

a slight but significant implicit causality subject-bias for these verbs relative to the grand mean across verbs. This matches what we found in our follow-up experiment on class 59 in the previous section (the overall object shift in our main experiment makes that experiment less informative for this issue). However, this subject bias is much weaker than the one observed for classes 31.1 (*amuse*) and 31.4 (*appeal to*). While this finding is perfectly consistent with our broader conclusion that re-mention biases are a systematic function of verb class, it is intriguing. We consider possible explanations below.

One interesting possibility, highlighted in recent work by Bott and Solstad (in press), is that different verb classes may encode different kinds of causes (see also Bittner & Dery, 2014). Bott and Solstad distinguish between simple causes (*John disturbed Mary because he was making a lot of noise*), externally anchored reasons (*John disturbed Mary because she had damaged his bike*), and internally anchored reasons (*John disturbed on purpose Mary because he was angry at her*) (for related discussion, see Lombrozo, 2010). They argue that different verb classes differ in their compatibility with these different types of causes, an intuition echoed in VerbNet's semantic representations, which distinguishes several types of causation, including *cause*, *force*, and *in reaction to*.

It may be that some types of causes make better explanations and thus, by extension, better targets for implicit causality bias, explaining the weaker effects for classes 45.4 and 59. Such a result would have broad implications beyond re-mention biases, since with a few exceptions (e.g., Verbnet), semanticists have generally assumed a unitary notion of causation. Thus, this may be a profitable avenue for future research.

Towards a Generative Account of Re-mention Biases

While evidence for the semantic structure account is not yet conclusive – there are more verb classes to be tested in more discourse contexts (e.g., sentences involving *and*, *but*, or *although*), and independently-verified semantic analyses are needed in many cases – the above results, coupled with previous work, are highly promising (cf. Arnold 2001; Crinean and Garnham, 2006; Hartshorne, 2014; Hartshorne et al., 2013; Hartshorne & Snedeker, 2013). This leaves open a deep question about re-mention biases that has received too little attention: How and why do listeners draw on semantic information to interpret pronouns in the first place?

Three possibilities have been explored in some detail. A number of researchers explained implicit causality biases by arguing that causes are salient (Garnham, Traxler, Oakhill, & Gernsbacher, 1996; McDonald & MacWhinney, 1995; Song & Fisher, 2005), assimilating the phenomenon into earlier theories on which pronouns refer to the most salient antecedent (Evans, 1980; Fletcher, 1984; Gundel, Hedberg, & Zacharski, 1993; Song & Fisher, 2005, 2007; van Dijk & Kintsch, 1983; van Rij, van Rij, & Hendriks, 2013). This account suffers in that it fails to explain implicit consequentiality. Moreover, it has proven difficult to establish that causes really are salient (cf. Garnham et al., 1996; McDonald & MacWhinney, 1995).

Other researchers have suggested that re-mention biases arise from heuristics learned from the statistics of pronoun use: People learn that in certain contexts, pronouns usually refer to

the previous cause, whereas in other contexts, pronouns usually refer to the previous affected entity, etc. (Crawley et al., 1990; Fletcher, 1984; see also Stevenson et al., 1994). Some support for this possibility comes from that fact that even infants can learn arbitrary statistical predictors of upcoming components of the speech stream (Saffran, Aslin, & Newport, 1996).

While the flexibility of this approach has some appeal, it risks presupposing the solution to the problem it is meant to solve. Third-person pronouns are always potentially ambiguous. In principle, the *he* in *Archibald frightens Bartholomew because he is scary* could refer to Archibald, Bartholomew, or any other male entity. Heuristics are invoked to explain how people nonetheless converge on an interpretation. But is difficult to see how we can use heuristics to determine pronoun reference if we must use pronoun reference to learn the heuristics in the first place.

A third, related account avoids this circularity is the *expectancy hypothesis* (Arnold, 1998, 2001; Arnold et al., 2007). Listeners learn features of the input that predict what will be mentioned next. These expectations can be learned from cases in which reference was unambiguous (*Archibald frightens Bartholomew because Archibald is scary*) and then generalized to sentences involving prepositions.

Arnold and colleagues have discussed two ways in which these expectations might be built. One is that listeners may learn heuristics from statistical correlations between features of the speech stream and subsequent referents (Arnold et al., 2007, p. 531). In this, the expectancy hypothesis operates much like the heuristics account described above, except that heuristics predict next-mention rather than pronoun reference *per se*. One potential limitation of this approach is that the relationship between causes and explanations and the relationship between affected entities and consequences is left unmotivated. People learn this relationship because it is a feature of the input. However, if the statistics of the input justified it, they should be just as happy to learn that explanations usually refer to the affected entity rather than the cause. One might wonder whether rote statistical learning is necessary to learn principled, motivated relationship such as the one between causation and explanation.

Arnold and colleagues also mention an alternative: Listeners use features of the discourse to build rich models of the speaker's intentions and goals, using this to generate expectations about what the speaker might refer to next (Arnold et al., 2007, p. 532). Although intriguing, no further detail is provided about the nature of these models, how the input is used to construct the models, or how inferences are derived from the models.

Building on the insights of the expectancy hypothesis as well as work by Kehler and colleagues (Kehler, 2002; Kehler et al., 2008) and Sagi and Rips (in press), we introduce the *most probable message (MPM) account*. On the MPM account, re-mention biases are a manifestation of a much broader account of language comprehension. Listeners have an intuitive model of the generative[14] processes that give rise to speaker behavior: What the

[14]Here we use the term *generative model* in the original sense of a model that explains how some observable behavior was generated, and not in reference to a particular tradition of syntactic theory (cf. Tenenbaum, Kemp, Griffiths, and Goodman, 2011).

speaker says depends on the speaker's goals, intentions, knowledge state, and – most importantly – what message the speaker wishes to convey. In this, we build on recent computational models of pragmatics (Frank and Goodman, 2012; Goodman and Stuhlmuller, 2013) which are themselves rooted in earlier observations by Grice (1989).

Like many current theories, MPM assumes incremental, predictive processing (Altmann & Kamide, 1999; Arnold et al., 2001; Arnold et al., 2007; Huang & Snedeker, 2011; Levy, 2008; Kutas, DeLong, & Smith, 2011; Snedeker, 2009; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995; van Berkum, 2008). In particular, the listener tries to infer the speaker's message based on whatever information is currently available, which may be less than the entire utterance.

For the moment, though, we consider how the MPM account handles pronouns in the context of complete sentences. We return to predictive re-mention biases later. Focusing just on the pronoun, the problem is to determine its intended reference in sentences like (6), which reduces to determining which sentence in (7) has the same meaning as (6):

(6) Archibald angered Bartholomew because he is reckless.

(7)

a. Archibald angered Bartholomew because **Archibald** is reckless.

b. Archibald frightened Bartholomew because **Bartholomew** is reckless.

c. Archibald frightened Bartholomew because **Cameron** is reckless.

d. Archibald frightened Bartholomew because **Dionysus** is reckless.

Etc.

As noted by Kehler et al. (2008), the probability the speaker meant (7a), for instance, is proportional to the probability that she would say (6) given that she meant (7a) times the probability that she wanted to communicate (7a) in the first place. That is, following Bayes' rule:

(8) $P(message \mid utterance) \propto P(utterance \mid message)\, P(message)$

Note that speakers typically only use pronouns to refer to a recently-mention entity. As a result, the probability of uttering (6) – which has a pronoun – when intending to convey (7c) or (7d), etc., is very low, since Cameron and Dionysus, etc., have not been mentioned previously. By equation (8), (6) is unlikely to mean anything other than (7a) or (7b), as desired.

It remains to explain why (7a) is preferred to (7b). Again, expectations about the speaker's behavior play a critical role. By hypothesis, the probability that the speaker would want to convey a particular message depends on whether that message is true, whether the speaker knows it is true, and what the speaker is motivated to communicate (Grice, 1989). Note that in all the examples described in this paper, we know nothing about the speaker's intentions or her knowledge of the world. We do, however, know what is likely to be true: Archibald being reckless is much more likely to result in Archibald angering Bartholomew than is Bartholomew being reckless, and so the speaker is more likely *a priori* to want to convey (7a) than (7b). Thus, by equation (8), the most likely interpretation of (6) is (7a), as desired.

A similar analysis applies straightforwardly to implicit consequentiality sentences. *Archibald angered Bartholomew so he plotted revenge* most likely refers to Bartholomew plotting revenge because that results in an interpretation of the sentence that is more likely to actually be true of the world.[15] Note that this captures the widely-shared intuition that linguistic ambiguity is often resolved in favor of the more plausible interpretation (Frank & Goodman, 2012; Garnsey, Pearlmutter, Myers, & Lotocky, 1997; Grice, 1989; Hobbs, Stickel, Appelt, & Martin, 1993; McRae, Spivey-Knowlton, & Tanenhaus, 1998).

Now we turn to verb bias – the topic of the present work. Again, we would like to know which interpretation of a pronoun (10) is meant to be conveyed by some sentence (9).

(9) Archibald angered Bartholomew because he…

(10)
- **a.** Archibald angered Bartholomew because **Archibald**…
- **b.** Archibald angered Bartholomew because **Bartholomew**…
- **c.** Archibald angered Bartholomew because **Cameron**…
- **d.** Archibald angered Bartholomew because **Dalton**…
- …

Unlike in the previous case, we have not yet heard the entire sentence. For this reason, we cannot directly compare the prior probabilities that (10a) and (10b) are true for the simple reason that we do not know what events these partial sentences refer to. One possibility would be to sum over all possible continuations of the sentence. Unfortunately, there are an infinite number of possible continuations, making this a difficult computational problem, particularly since listeners can calculate implicit causality biases within a few hundred milliseconds of encountering the pronoun (Koornneef and van Berkum, 2006; Pyykkonen and Jarvikivi, 2010; *inter alia*).

The semantic structure account suggests a possible fast approximation. From the meaning of the verb *anger*, we know that there is at least one cause of Archibald angering Bartholomew that involves Archibald. There may be other causes as well, and they may involve Archibald, Bartholomew, or even other people. However, all else being equal, it is more likely that the explanation of Archibald angering Bartholomew will involve Archibald for the simple reason that we know at least one such explanation exists. Thus, interpreting (9) such that *he* refers to Archibald results in a proposition that is more likely to actually be true than interpretations where *he* refers to Bartholomew. Listeners who are sensitive to this fact would exhibit the standard re-mention bias.

This account has several key features that differentiate it from other accounts. Unlike salience accounts, it incorporates discourse structure. Unlike heuristic and expectancy accounts, re-mention biases are derived from the structure of language and thought rather than learned from correlations in the input. There is no explicit salience hierarchy or

[15]We are indebted to Sagi & Rips (in press) for the suggestion of using the probabilities of the events themselves. Among the many differences between our approach and theirs is that rather than invoking Gricean reasoning, Sagi & Rips embed their theory in the notion of causal identity: What makes *he* in likely to refer to Archibald in (6) is the fact that *he* and Archibald are likely to be causally related. We are currently attempting to tease apart these two accounts in ongoing research.

expectedness hierarchy; re-mention biases are not expectations about reference but rather manifestations of current beliefs about what the speaker is saying, based on the best available evidence. There is much left to be done to flesh out and test this account, a process currently underway (Hartshorne, O'Donnell, and Tenenbaum, in prep). However, we believe that it provides a potentially promising avenue for further research and consideration.

Conclusion

Above, we present data from 502 verbs, showing that implicit causality and consequentiality biases are a systematic function of Levin verb class. This represents by far the largest survey of implicit consequentiality biases to date, and thus these data should serve to test and constrain all theories. We argue that in particular these results support a theory on which the semantic information underlying re-mention biases is encoded in the structure of language, rather than inferred from general knowledge about the world.

Acknowledgments

Thanks are due to Jesse Snedeker, Jess Sullivan, Hugh Rabagliati, Pierina Cheung, Jennifer Arnold, Jennifer Rodd, Mahesh Srinivasan, & Andrew Stewart for comments and suggestions, to John H. Krantz for assistance in subject recruitment, and to the volunteers who participated in this study. The first author was supported by an NIH Ruth L. Kirschstein National Research Service Award (5F32HD072748) and by the NSF Graduate Research Fellowship Program award.

References

Altmann GTM, Kamide Y. Incremental interpretation at verbs: Restricting the domain of subsequent reference. Cognition. 1999; 73:247–264. [PubMed: 10585516]

Ambridge B, Pine JM, Rowland CF, Jones RL, Clark V. A semantics-based approach to the "no negative evidence problem". Cognitive Science. 2009; 33(7):1301–1316. [PubMed: 21585506]

Arnold JE. Reference form and discourse patterns (Doctoral dissertation, Stanford University). Dissertation Abstracts International. 1998; 59:2950.

Arnold JE. The Effect of Thematic Roles on Pronoun Use and Frequency of Reference Continuation. Discourse Processes. 2001; 31(2):137–162.

Arnold JE, Brown-Schmidt S, Trueswell JC. Children's use of gender and order-of-mention during pronoun comprehension. Language and Cognitive Processes. 2007; 22(4):527–565.

Au TK-F. A verb is worth a thousand words: the causes and consequences of interpersonal events implicit in language. Journal of Memory and Language. 1986; 25:104–122.

Bittner, D.; Dery, JE. Emergence of implicit causality in German-learning children; Paper presented at the International Association for the Study of Child Language (IASCL); Amsterdam, Netherlands. 2014.

Bott, O.; Solstad, T. From verbs to discourse: A novel account of implicit causality. In: Hemforth, B.; Mertins, B.; Fabricius-Hansen, C., editors. Psycholinguistic approaches to meaning and understanding across languages. (in press).

Brown R, Fish D. Are there universal schemas of psychological causality? Archives de Psychologie. 1983a; 51:145–153.

Brown RW, Fish D. The psychological causality implicit in language. Cognition. 1983b; 14:237–273. [PubMed: 6686101]

Caramazza A, Grober EH, Garvey C, Yates J. Comprehension of Anaphoric Pronouns. Journal of Verbal Learning and Verbal Behavior. 1977; 16:601–609.

Crinean M, Garnham A. Implicit causality, implicit consequentiality and semantic roles. Language and Cognitive Processes. 2006; 21(5):636–648.

Corrigan R. Implicit causality in language: event participants and their interactions. Journal of Language and Social Psychology. 2001; 20:285–320.

Corrigan R. The influence of evaluation and potency on perceivers' causal attributions. European Journal of Social Psychology. 2002; 32:363–382.

Corrigan R. Preschoolers' and adults' attributions of who causes interpersonal events. Infant and Child Development. 2003; 12:305–328.

Cozijn R, Commandeur E, Vonk W, Noordman LGM. The time course of the use of implicit causality information in the processing of pronouns: A visual world paradigm study. Journal of Memory and Language. 2011; 64:381–403.

Ehrlich K. Comprehension of pronouns. The Quarterly Journal of Experimental Psychology. 1980; 32(2):247–255.

Evans G. Pronouns. Linguistic Inquiry. 1980; 11(2):337–362.

Featherstone CR, Sturt P. Because there was a cause for concern: An investigation into a word-specific prediction account of the implicit-causality effect. The Quarterly Journal of Experimental Psychology. 2010; 63(1):3–15. [PubMed: 19691006]

Ferstl EC, Garnham A, Manouilidou C. Implicit causality bias in English: a corpus of 300 verbs. Behavior Research Methods. 2011; 43(1):124–135. [PubMed: 21287124]

Fletcher CR. Markedness and topic continuity in discourse processing. Journal of Verbal Learning and Verbal Behavior. 1984; 23(4):487–493.

Frank MC, Goodman ND. Pragmatic reasoning in language games. Science. 2012; 336(6084):998. [PubMed: 22628647]

Garnham A, Traxler M, Oakhill J, Gernsbacher MA. The Locus of Implicit Causality Effects in Comprehension. Journal of Memory and Language. 1996; 35(4):517–543.

Garnsey SM, Pearlmutter NJ, Myers E, Lotocky MA. The contributions of verb bias and plausibility to the comprehension of temporarily ambiguous sentences. Journal of Memory and Language. 1997; 37(1):58–93.

Garvey C, Caramazza A. Implicit causality in verbs. Linguistic Inquiry. 1974; 5(3):459–464.

Garvey C, Caramazza A, Yates J. Factors influencing assignment of pronoun antecedents. Cognition. 1974; 3(3):227–243.

Goikoetxea E, Pascual G, Acha J. Normative study of the implicit causality of 100 interpersonal verbs in Spanish. Behavior Research Methods. 2008; 40(3):760–772. [PubMed: 18697672]

Grice, P. Studies in the Way of Words. Cambridge, MA: Harvard University Press; 1989.

Gundel JK, Hedberg N, Zacharski R. Cognitive status and the form of referring expressions in discourse. Language. 1993:274–307.

Hartshorne JK. What is implicit causality? Language, Cognition, and Neuroscience. 2014; 29(7):804–824.

Hartshorne JK, O'Donnell TJ, Tenenbaum JB. The Most Probable Message. (in prep).

Hartshorne JK, Snedeker J. Verb argument structure predicts implicit causality: The advantages of finer-grained semantics. Language and Cognitive Processes. 2013

Hartshorne JK, Sudo Y, Uruwashi M. Are implicit causality pronoun resolution biases consistent across languages and cultures? Experimental Psychology. 2013; 60(3):179–196. [PubMed: 23261949]

Huang Y, Snedeker J. 'Logic & Conversation' revisited: Evidence for a division between semantic and pragmatic content in real time language comprehension. Language and Cognitive Processes. 2011; 26:1161–1172.

Hobbs JR, Stickel ME, Appelt DE, Martin P. Interpretation as abduction. Artificial Intelligence. 1993; 63(1–2):69–142.

Kehler A, Kertz L, Rohde H, Elman JL. Coherence and coreference revisited. Journal of Semantics. 2008; 25:1–44. [PubMed: 22923856]

Kipper K, Korhonen A, Ryant N, Palmer M. Extending VerbNet with novel verb classes. Presented at the Proceedings of LREC. 2006

Koornneef AW, Sanders TJM. Establishing coherence relations in discourse: The influence of implicit causality and connectives on pronoun resolution. Language and Cognitive Processes. 2012 Sep.:1–38.

Koornneef AW, Van Berkum JJA. On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye tracking. Journal of Memory and Language. 2006; 54(4):445–465.

Kutas, M.; DeLong, KA.; Smith, NJ. A look at what lies ahead: Prediction and predictability in language processing. In: Bar, M., editor. Predictions in the brain: Using our past to generate a future. p. 190-207.

Levin, B. English Verb Classes and Alternations. Chicago, IL: University Of Chicago Press; 1993.

Levy R. Expectation-based syntactic comprehension. Cognition. 2008; 106(3):1126–1177. [PubMed: 17662975]

Lombrozo T. Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. Cognitive Psychology. 2010; 61(4):303–332. [PubMed: 20801434]

McKoon G, Greene SB, Ratcliff R. Discourse Models, Pronoun Resolution, and the Implicit Causality of Verbs. Journal of Experimental Psychology: Learning, Memory, and Cognition. 1993; 19:1040–1052.

McRae K, Spivey-Knowlton MJ, Tanenhaus MK. Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. Journal of Memory and Language. 1998; 38:283–312.

Pickering MJ, Majid A. What are implicit causality and consequentiality? Language and Cognitive Processes. 2007; 22(5):780–788.

Pyykkonen P, Jarvikivi J. Activation and persistence of implicit causality information in spoken language comprehension. Experimental Psychology. 2010; 57:5–16. [PubMed: 20176549]

Rigalleau F, Guerry M, Granjon L. The effect of connectives on the selection of arguments: Implicit consequentiality bias for the connective ‘but’. Journal of Psycholinguistic Research. (in press).

Rudolph U, Forsterling F. The Psychological Causality Implicit in Verbs: A Review. Psychological Bulletin. 1997; 121(2):192–218.

Saffran JR, Alsin RN, Newport EL. Statistical learning by 8-month-old infants. Science. 1996; 274(5294):1926–1928. [PubMed: 8943209]

Sagi E, Rips LJ. Identity, causality, and pronoun ambiguity. Topics in Cognitive Science. (in press).

Schuler, KK. “VerbNet: A broad-coverage, comprehensive verb lexicon” by Karin Kipper Schuler. University of Pennsylvania; 2005.

Snedeker, J. Children’s sentence processing. In: Bavin, E., editor. The handbook of child language. Cambridge, UK: Cambridge University Press; 2009. p. 331-338.

Song H-J, Fisher C. Who”s ‘she’? Discourse prominence influences preschoolers” comprehension of pronouns. Journal of Memory and Language. 2005; 52(1):29–57.

Song H-J, Fisher C. Discourse prominence effects on 2.5-year-old children’s interpretations of pronouns. Lingua. 2007; 117(11):1959–1987. [PubMed: 18978930]

Stevenson RJ, Crawley RA, Kleinman D. Thematic roles, focus and the representation of events. Language and Cognitive Processes. 1994; 9(4):519–548.

Stevenson RJ, Knott A, Oberlander J, McDonald S. Interpreting pronouns and connectives: Interactions among focusing, thematic roles and coherence relations. Language and Cognitive Processes. 2000; 15(3):225–262.

Stewart AJ, Pickering MJ, Sanford AJ. Implicit Consequentiality. Proceedings of the 20th Annual Conference of the Cognitive Science Society. 1998:1–9.

Tanenhaus MK, Spivey-Knowlton MJ, Eberhard KM, Sedivy JE. Integration of visual and linguistic information in spoken language comprehension. Science. 1995; 268:1632–1634. [PubMed: 7777863]

Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND. How to grow a mind: Statistics, structure, and abstraction. Science. 2011; 331(6022):1279–1285. [PubMed: 21393536]

Van Berkum JJA. Understanding sentences in context: What brain waves can tell us. Current Directions in Psychological Science. 2008; 17(6):376–380.

Van Dijk, T.; Kintsch, W. Strategies of discourse comprehension. New York: Academic Press; 1983.

van Rij J, van Rijn H, Hendriks P. How WM Load Influences Linguistic Processing in Adults: A Computational Model of Pronoun Interpretation in Discourse. Topics in Cognitive Science. 2013 n/a–n/a.

White PA. A theory of causal processing. British Journal of Psychology. 1989; 80:431–454.

Appendix

Results by Verb (Chose Subj / N)

Verb	class	Result	Explanation
abashed	31.1	14/46	15/41
affected	31.1	16/44	22/42
afflicted	31.1	7/35	30/61
affronted	31.1	10/49	15/42
aggravated	31.1	11/36	31/51
aggrieved	31.1	14/48	25/43
agitated	31.1	9/48	25/45
agonized	31.1	12/31	21/46
alarmed	31.1	7/44	37/48
alienated	31.1	22/45	17/46
amazed	31.1	15/38	30/38
amused	31.1	12/39	41/54
angered	31.1	14/39	34/52
antagonized	31.1	11/33	12/41
appalled	31.1	15/38	25/46
appeased	31.1	15/43	19/54
aroused	31.1	10/45	36/58
assuaged	31.1	17/42	8/44
astonished	31.1	9/33	34/47
astounded	31.1	16/40	30/42
awed	31.1	23/53	23/47
baffled	31.1	8/39	30/42
befuddled	31.1	14/43	22/39
beguiled	31.1	16/43	18/42
bewildered	31.1	10/34	33/51
bewitched	31.1	16/51	15/50
boggled	31.1	17/48	17/38
bored	31.1	15/49	26/35
bugged	31.1	6/59	20/35
calmed	31.1	8/35	12/46
captivated	31.1	14/37	34/49
chagrined	31.1	16/45	16/64
charmed	31.1	14/38	31/48
cheered	31.1	12/51	5/42
comforted	31.1	8/34	4/51

Verb	class	Result	Explanation
confounded	31.1	15/38	29/54
confused	31.1	12/32	32/42
consoled	31.1	18/44	8/50
contented	31.1	14/44	19/56
cowed	31.1	17/43	12/58
daunted	31.1	16/51	21/46
dazed	31.1	15/39	30/42
dazzled	31.1	13/44	38/54
dejected	31.1	13/35	12/57
delighted	31.1	11/43	34/47
demoralized	31.1	10/46	16/69
depressed	31.1	11/41	27/48
disappointed	31.1	16/34	49/60
discombobulated	31.1	9/42	20/49
discomfitted	31.1	17/42	23/46
discomposed	31.1	20/59	14/46
disconcerted	31.1	22/49	30/49
disgraced	31.1	12/36	16/42
disgruntled	31.1	17/48	27/58
disgusted	31.1	19/41	29/51
disheartened	31.1	11/36	25/48
disillusioned	31.1	12/38	31/52
dismayed	31.1	14/33	30/64
dispirited	31.1	12/34	26/56
displeased	31.1	12/39	28/46
disquieted	31.1	21/54	30/65
dissatisfied	31.1	16/34	33/48
distracted	31.1	15/48	29/47
distressed	31.1	15/47	38/53
disturbed	31.1	18/48	22/33
diverted	31.1	12/39	13/47
elated	31.1	10/29	23/42
electrified	31.1	12/38	18/51
embarrassed	31.1	14/45	24/41
emboldened	31.1	9/41	30/52
enchanted	31.1	12/39	29/55
enervated	31.1	15/46	23/50
engrossed	31.1	12/41	25/38
enlightened	31.1	12/35	17/52
enlivened	31.1	11/37	24/44
enraged	31.1	10/37	30/44
enraptured	31.1	16/32	17/38

Verb	class	Result	Explanation
entertained	31.1	8/37	13/42
enthralled	31.1	10/31	33/51
enthused	31.1	15/47	28/47
entranced	31.1	9/43	27/51
exasperated	31.1	11/48	25/37
excited	31.1	8/39	29/49
exhausted	31.1	9/44	26/38
exhilarated	31.1	17/48	30/51
fascinated	31.1	18/47	36/53
fatigued	31.1	5/28	30/47
fazed	31.1	7/39	20/48
flabbergasted	31.1	12/36	30/42
floored	31.1	20/56	16/49
flustered	31.1	12/47	34/57
frustrated	31.1	12/41	28/42
galled	31.1	19/38	18/49
galvanized	31.1	14/37	15/57
gratified	31.1	21/51	14/42
grieved	31.1	14/31	13/52
harmed	31.1	20/47	7/42
haunted	31.1	15/44	14/51
heartened	31.1	16/48	16/44
horrified	31.1	12/56	45/59
humbled	31.1	9/46	18/50
humiliated	31.1	8/40	14/47
hypnotized	31.1	11/49	16/49
impaired	31.1	16/39	26/64
impressed	31.1	18/43	41/49
incensed	31.1	17/47	27/53
inflamed	31.1	11/40	25/61
infuriated	31.1	11/48	30/49
inspired	31.1	9/32	45/56
interested	31.1	13/38	36/55
intimidated	31.1	10/44	31/45
intoxicated	31.1	14/43	27/54
intrigued	31.1	15/42	20/39
invigorated	31.1	11/47	23/40
irked	31.1	13/45	29/50
irritated	31.1	11/42	34/55
jaded	31.1	19/47	9/36
jollified	31.1	11/41	9/43
maddened	31.1	15/47	27/43

Verb	class	Result	Explanation
menaced	31.1	12/43	11/47
mesmerized	31.1	19/48	27/40
miffed	31.1	15/44	26/51
molested	31.1	14/36	10/46
mollified	31.1	9/35	17/52
mortified	31.1	15/40	29/44
mystified	31.1	15/43	39/47
nauseated	31.1	13/39	32/46
nettled	31.1	16/38	16/52
numbed	31.1	9/38	11/52
obsessed	31.1	18/40	21/47
occupied	31.1	10/37	15/51
offended	31.1	21/44	30/38
outraged	31.1	15/54	44/61
overawed	31.1	15/44	28/54
overwhelmed	31.1	15/47	38/60
pacified	31.1	14/47	13/53
pastered	31.1	11/35	12/46
peeved	31.1	12/56	32/53
perplexed	31.1	9/41	37/51
perturbed	31.1	9/42	36/52
piqued	31.1	15/44	20/47
placated	31.1	14/40	9/57
plagued	31.1	10/44	5/48
pleased	31.1	16/43	33/50
preoccupied	31.1	8/31	23/43
puzzled	31.1	5/25	28/50
rankled	31.1	7/37	15/44
ravished	31.1	12/35	10/43
reassured	31.1	10/33	12/47
recharged	31.1	12/32	9/50
refreshed	31.1	16/44	21/46
rejuvanated	31.1	13/41	18/60
relaxed	31.1	15/51	26/48
repelled	31.1	20/57	17/38
repulsed	31.1	22/55	28/47
revitalized	31.1	10/39	31/60
revolted	31.1	16/40	22/57
riled	31.1	11/50	19/45
ruffled	31.1	11/46	17/44
saddened	31.1	16/38	25/33
satiated	31.1	19/53	23/50

Verb	class	Result	Explanation
satisfied	31.1	11/39	31/51
scandalized	31.1	8/27	15/46
scared	31.1	14/39	30/46
solaced	31.1	16/47	5/47
spellbound	31.1	18/55	22/41
spooked	31.1	14/55	34/48
startled	31.1	6/50	34/51
stunned	31.1	13/42	28/47
stupefied	31.1	11/35	18/39
surprised	31.1	8/37	30/46
tantalized	31.1	10/38	27/51
taunted	31.1	8/42	5/58
terrorized	31.1	9/37	10/49
threatened	31.1	7/50	4/38
thrilled	31.1	7/41	34/49
titillated	31.1	3/33	22/44
tormented	31.1	6/33	5/51
tortured	31.1	11/41	7/56
transfixed	31.1	7/35	21/45
troubled	31.1	16/43	28/44
unnerved	31.1	18/51	32/47
unsettled	31.1	17/49	32/38
uplifted	31.1	17/47	13/43
upset	31.1	20/55	38/47
vexed	31.1	12/34	25/44
wearied	31.1	22/39	36/59
worried	31.1	12/47	26/43
wounded	31.1	16/44	12/49
wowed	31.1	15/40	38/56
abhorred	31.2	35/45	5/59
admired	31.2	32/43	1/39
adored	31.2	38/53	4/55
appreciated	31.2	34/60	2/44
bewailed	31.2	22/40	8/48
cherished	31.2	28/36	4/59
deified	31.2	19/32	9/48
deplored	31.2	23/36	6/48
despised	31.2	33/39	3/47
detested	31.2	41/46	3/52
disbelieved	31.2	22/33	4/42
disdained	31.2	32/43	5/51
disliked	31.2	34/41	1/48

Verb	class	Result	Explanation
distrusted	31.2	32/39	1/63
dreaded	31.2	48/54	3/62
envied	31.2	48/61	3/44
esteemed	31.2	32/39	8/55
exalted	31.2	13/35	8/44
execrated	31.2	16/48	6/52
favored	31.2	30/46	3/43
feared	31.2	55/63	4/43
grudged	31.2	17/26	5/44
hated	31.2	40/44	1/40
idolized	31.2	32/35	3/52
lamented	31.2	30/42	8/58
loathed	31.2	35/43	4/53
loved	31.2	38/42	2/56
missed	31.2	45/51	19/56
mourned	31.2	22/32	9/50
pitied	31.2	29/42	2/51
preferred	31.2	26/40	4/59
prized	31.2	31/46	9/49
regretted	31.2	37/44	12/50
relished	31.2	37/50	1/43
resented	31.2	22/29	4/54
respected	31.2	41/45	3/65
revered	31.2	34/44	3/47
rued	31.2	29/51	3/43
agonized over	31.3	24/34	5/44
angered over	31.3	35/44	4/51
anguished over	31.3	31/37	6/43
approved of	31.3	28/37	3/48
bled for	31.3	18/38	8/41
cared about	31.3	38/42	12/57
cared for	31.3	48/60	15/49
cheered at	31.3	7/46	1/54
cried for	31.3	22/36	12/48
delighted in	31.3	41/51	7/56
delighted over	31.3	12/29	2/47
despaired of	31.3	34/42	6/40
disapproved of	31.3	25/36	3/50
enthused at	31.3	10/52	5/44
enthused over	31.3	18/40	8/55
exulted at	31.3	9/32	5/45
exulted in	31.3	29/41	4/42

Verb	class	Result	Explanation
exulted over	31.3	26/44	14/51
feared for	31.3	43/43	4/44
felt for	31.3	39/44	9/47
fretted about	31.3	33/42	3/46
fretted over	31.3	25/43	7/52
fumed at	31.3	21/57	2/45
fumed over	31.3	24/44	2/45
gladdened at	31.3	12/46	11/59
gloated over	31.3	16/41	8/40
grieved for	31.3	39/46	4/33
grieved over	31.3	41/47	6/52
gushed over	31.3	18/51	4/54
hungered over	31.3	27/37	10/44
maddened at	31.3	23/45	5/43
marveled at	31.3	23/39	2/46
marveled over	31.3	19/36	2/40
meditated over	31.3	40/56	10/61
minded about	31.3	32/37	14/56
mooned about	31.3	15/27	12/51
mooned over	31.3	15/37	13/49
moped over	31.3	22/36	8/40
mourned for	31.3	37/43	6/49
mourned over	31.3	44/52	5/47
mused over	31.3	16/30	8/59
obsessed over	31.3	28/40	8/59
puzzled over	31.3	32/41	6/44
raged at	31.3	7/37	4/51
raged over	31.3	25/43	5/35
reacted to	31.3	29/48	5/45
reflected over	31.3	30/45	14/57
rejoiced about	31.3	27/46	7/47
rejoiced at	31.3	16/44	5/45
rejoiced in	31.3	27/40	6/42
rejoiced over	31.3	20/45	10/55
rhapsodied over	31.3	19/35	3/45
rhapsodized about	31.3	21/45	7/41
ruminated over	31.3	31/40	7/44
saddened at	31.3	24/43	6/57
sickened of	31.3	45/53	8/45
thrilled at	31.3	15/41	3/38
tired of	31.3	27/33	4/54
wearied of	31.3	32/39	3/43

Verb	class	Result	Explanation
wept for	31.3	24/35	11/62
wondered at	31.3	16/38	5/56
appealed to	31.4	19/47	22/49
grated on	31.4	13/43	25/48
jarred on	31.4	13/41	15/52
abused	33	11/39	8/49
acclaimed	33	13/44	8/51
accursed	33	12/35	3/49
assailed	33	16/46	6/52
assaulted	33	15/35	6/61
attacked	33	18/41	3/49
belittled	33	9/44	2/48
blamed	33	24/38	5/47
blasphemed	33	11/26	6/53
blessed	33	8/33	3/54
castigated	33	10/44	1/58
celebrated	33	23/39	5/56
censured	33	7/40	4/46
chastened	33	6/41	1/53
chastised	33	11/38	1/40
chided	33	13/43	5/39
commended	33	12/43	3/50
complimented	33	11/56	7/47
condemned	33	17/50	2/40
condoned	33	20/42	5/56
congratulated	33	7/43	3/43
criticized	33	14/51	4/60
cursed	33	9/41	6/42
damned	33	18/49	1/43
defamed	33	19/50	11/56
denigrated	33	10/38	6/43
denounced	33	7/33	3/53
deprecated	33	18/53	8/52
derided	33	11/47	4/43
disparaged	33	12/40	9/56
eulogized	33	19/30	5/39
excoriated	33	18/42	8/56
excused	33	13/35	4/45
extoled	33	17/40	4/48
faulted	33	16/39	3/56
felicitated	33	12/33	8/44
forgave	33	19/38	2/45

Verb	class	Result	Explanation
gibed	33	14/38	6/55
greeted	33	11/44	12/40
hailed	33	8/26	7/50
heralded	33	19/46	6/42
honored	33	20/47	6/53
impeached	33	12/32	2/50
incriminated	33	21/47	10/49
indicted	33	15/45	3/54
lambasted	33	10/37	3/44
lampooned	33	9/37	7/56
lauded	33	9/39	3/43
maligned	33	14/49	5/53
mocked	33	10/47	3/43
penalized	33	9/38	3/44
persecuted	33	22/54	1/36
praised	33	6/37	3/53
prosecuted	33	15/40	3/48
punished	33	7/27	2/43
rebuked	33	17/46	6/57
recompensed	33	19/46	14/47
remunerated	33	14/33	7/42
reprimanded	33	15/47	4/42
reproached	33	13/50	4/38
reproved	33	12/37	4/42
repudiated	33	15/39	6/57
reviled	33	23/37	5/58
ridiculed	33	13/44	1/45
scolded	33	10/50	2/46
scorned	33	18/36	2/54
slandered	33	15/34	5/39
snubbed	33	12/42	1/30
stigmatized	33	13/42	6/61
thanked	33	16/44	4/43
upbraided	33	18/44	0/50
victimied	33	16/39	5/60
vilified	33	18/45	4/36
abased	45.4	7/37	7/50
abated	45.4	22/52	11/52
abraded	45.4	15/42	6/46
activated	45.4	19/47	15/57
africanized	45.4	12/37	1/45
aged	45.4	13/34	16/48

Verb	class	Result	Explanation
americanized	45.4	12/40	9/50
anesthetized	45.4	14/45	5/55
anglicized	45.4	14/39	10/43
animated	45.4	13/34	16/46
apostatized	45.4	18/45	3/40
augmented	45.4	12/38	15/54
awaked	45.4	10/37	9/45
awakened	45.4	10/39	27/62
balanced	45.4	19/46	23/62
beautified	45.4	16/37	14/56
blackened	45.4	18/52	9/56
bleached	45.4	14/39	8/35
bloodied	45.4	17/34	7/47
capacitated	45.4	12/45	5/37
catholicized	45.4	21/45	4/54
cauterized	45.4	19/52	1/62
christianized	45.4	16/44	10/61
civilized	45.4	18/53	19/62
commercialized	45.4	13/40	10/58
constricted	45.4	4/41	12/39
contextualized	45.4	27/46	8/49
cooled	45.4	19/51	5/44
corrected	45.4	9/37	3/42
corrupted	45.4	16/50	18/51
cremated	45.4	40/56	6/56
deafened	45.4	11/30	24/48
decelerated	45.4	13/41	6/43
deflated	45.4	12/38	12/56
degraded	45.4	9/45	3/45
demobilized	45.4	12/46	8/55
devalued	45.4	15/39	5/49
dilated	45.4	22/59	10/44
disintegrated	45.4	29/47	4/47
domesticated	45.4	19/56	7/56
dried	45.4	11/42	3/39
effeminated	45.4	14/36	10/51
emaciated	45.4	21/53	9/40
embittered	45.4	18/45	28/47
embrocated	45.4	19/35	8/55
energized	45.4	15/40	15/40
europeanized	45.4	18/43	8/45
feminized	45.4	23/42	14/46

Verb	class	Result	Explanation
fertilized	45.4	12/44	9/54
halted	45.4	16/40	3/43
healed	45.4	10/35	12/46
hellenized	45.4	14/39	5/44
improved	45.4	24/58	32/62
incinerated	45.4	25/52	6/53
incubated	45.4	16/37	6/54
levitated	45.4	24/46	5/44
liquefied	45.4	21/44	5/48
mellowed	45.4	8/30	18/54
moderated	45.4	16/38	7/54
modernized	45.4	16/41	14/56
muddied	45.4	14/40	11/51
neutralized	45.4	17/50	7/42
objectified	45.4	13/40	5/40
obscured	45.4	24/52	14/44
paralyzed	45.4	24/44	20/59
perfected	45.4	33/47	18/49
popularized	45.4	18/47	13/55
publicized	45.4	6/40	3/45
quieted	45.4	14/55	4/46
quietened	45.4	16/32	8/38
reanimated	45.4	21/52	17/49
reddened	45.4	17/44	13/51
resuscitated	45.4	8/33	4/45
reversed	45.4	11/43	13/54
revived	45.4	15/42	10/49
sank	45.4	17/40	5/48
secularized	45.4	18/42	8/46
shushed	45.4	7/41	0/52
slowed	45.4	15/52	32/61
softened	45.4	11/39	16/43
sovietized	45.4	17/45	9/64
stabilized	45.4	15/53	10/42
steadied	45.4	14/41	5/48
sterilized	45.4	5/33	6/44
strengthened	45.4	14/37	28/50
submerged	45.4	16/50	4/48
tamed	45.4	18/53	20/48
toppled	45.4	15/37	9/29
toughened	45.4	7/32	10/52
tranquilized	45.4	10/48	3/60

Verb	class	Result	Explanation
wakened	45.4	11/45	10/43
warmed	45.4	12/44	12/47
weakened	45.4	9/42	33/59
westernized	45.4	19/49	12/50
woke	45.4	6/24	16/47
allured	59	14/43	34/63
arm-twisted	59	8/49	8/49
bamboozled	59	16/50	16/41
blackmailed	59	17/57	12/60
bluffed	59	15/43	11/51
bribed	59	18/46	14/48
bullocked	59	13/35	4/60
cajoled	59	12/40	3/48
coaxed	59	11/48	13/53
coerced	59	13/52	8/36
commissioned	59	25/44	6/46
compelled	59	11/40	20/55
dared	59	10/42	12/50
deceived	59	17/45	23/59
deluded	59	10/38	22/56
duped	59	16/46	11/47
ensnared	59	16/49	12/38
entrapped	59	21/49	8/51
fooled	59	17/48	28/42
forced	59	6/39	8/63
harried	59	8/43	14/64
hijacked	59	18/41	7/44
hoodwinked	59	9/30	18/55
hustled	59	19/51	11/45
impelled	59	14/42	10/41
induced	59	9/34	5/47
influenced	59	18/48	41/53
inveigled	59	22/46	10/48
lured	59	15/55	11/44
manipulated	59	13/34	8/43
misled	59	19/37	24/46
obligated	59	14/45	22/55
obliged	59	11/39	9/44
panicked	59	8/40	30/42
pressured	59	8/51	10/60
prompted	59	7/41	7/51
roused	59	9/39	17/52

Verb	class	Result	Explanation
seduced	59	13/33	19/57
spurred	59	10/50	7/56
sweet-talked	59	8/50	12/45
tricked	59	13/40	17/57

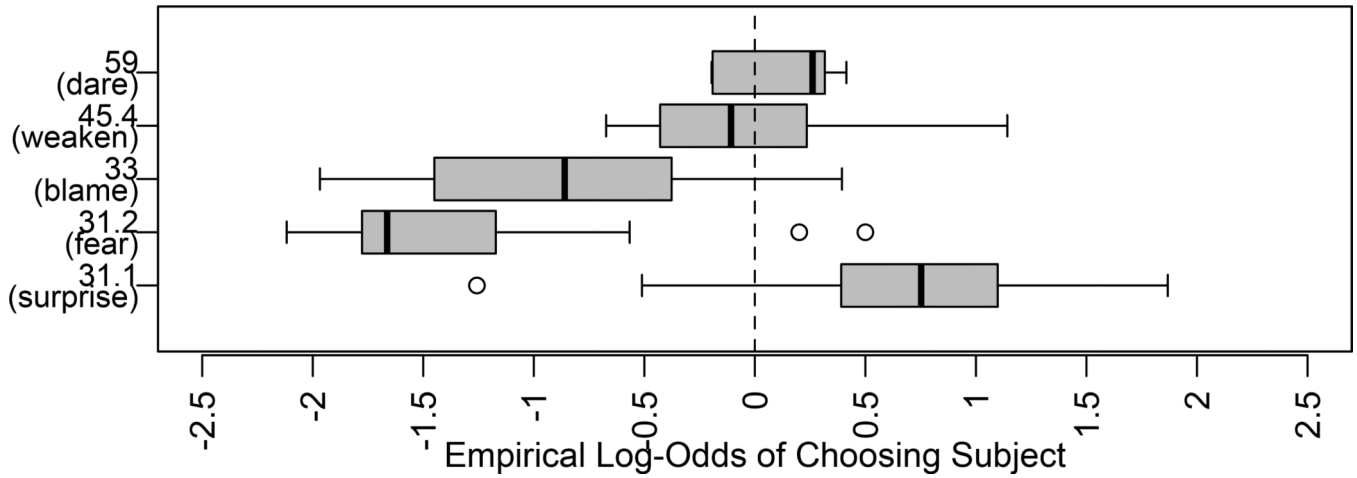

Figure 1.

Pronoun interpretation biases in explanatory ("implicit causality") sentences (e.g., 2–4) for the five verb classes (identified by VerbNet class number[16]) investigated in Hartshorne and Snedeker (2013), plotted by verb. For verbs that were tested in both Exp. 1 and Exp. 2, the results of the two experiments have been combined.

[16]Because most classes do not lend themselves to intuitive names, these classes are numbered. Throughout, we use the VerbNet numbers and give examples.

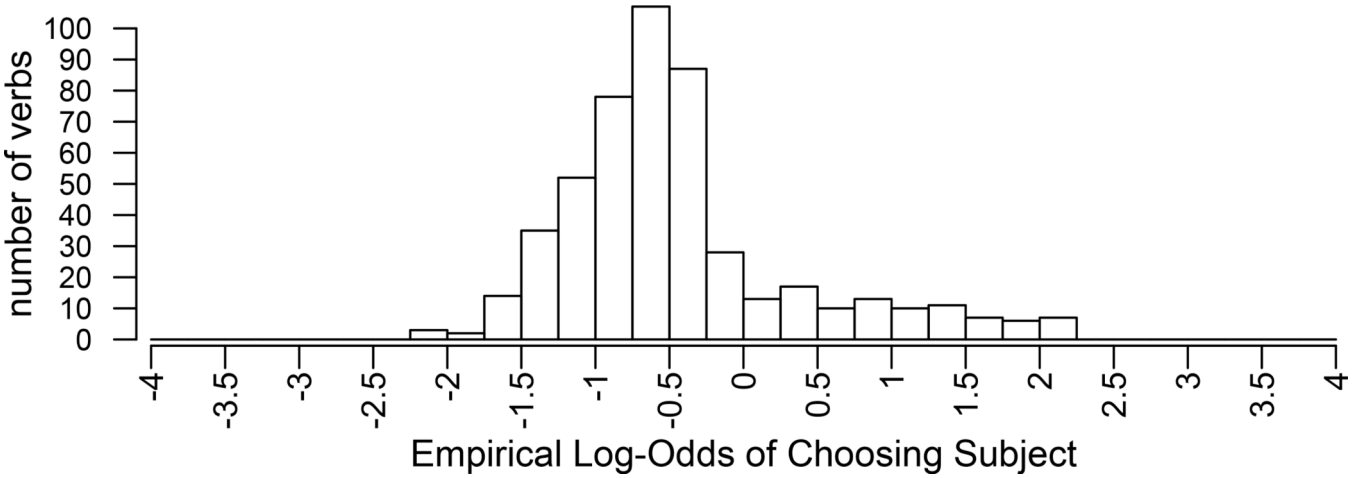

Figure 2.

A histogram of implicit consequentiality biases by verb, in result sentences. Note that at this scale one extremely subject-biased verb (*feared for*, class 31.3) is not visible.

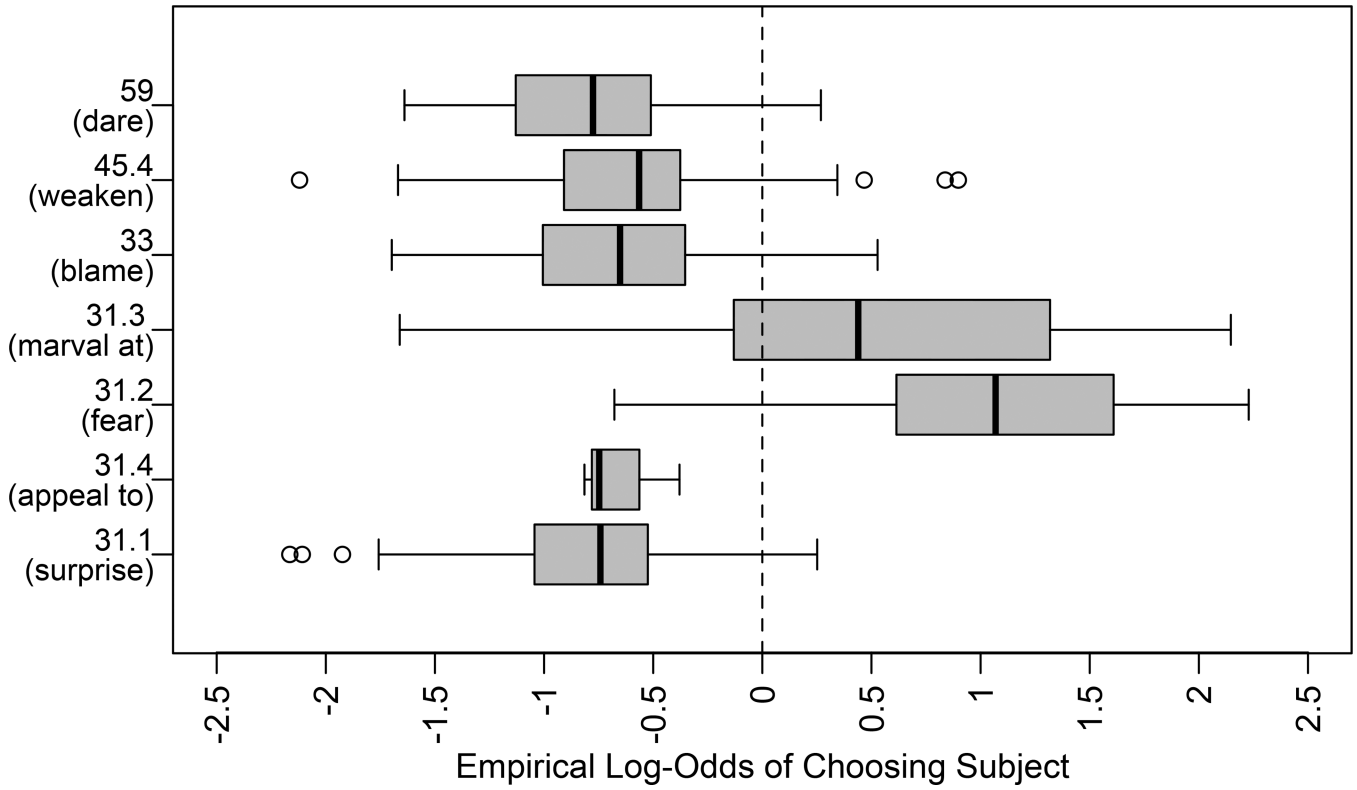

Figure 3.
Implicit consequentiality biases by verb, analyzed separately by verb class. Note that at this scale one extremely subject-biased verb (*feared for*, class 31.3) is not visible.

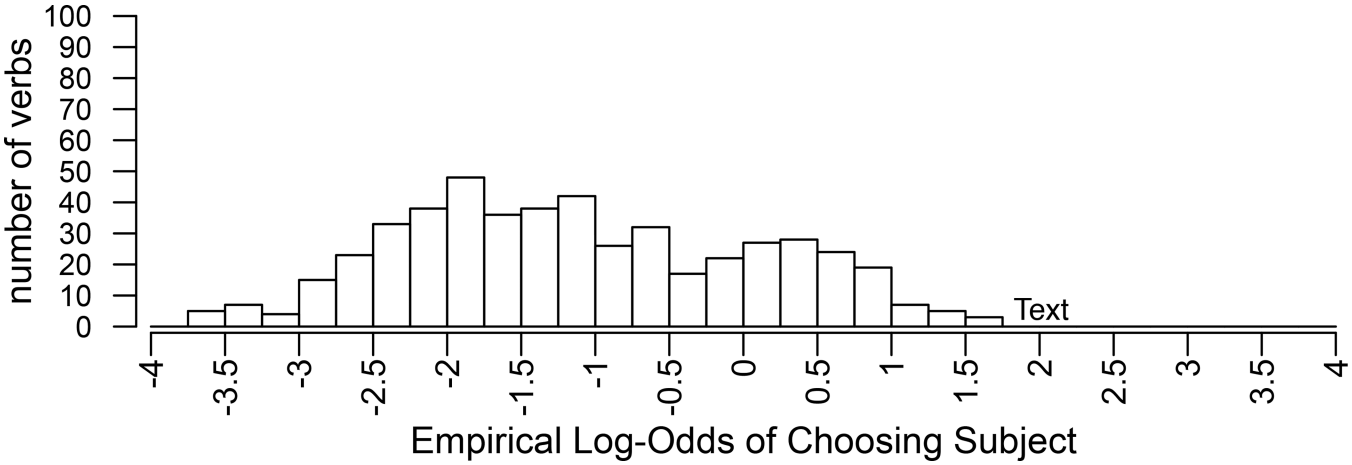

Figure 4.

A histogram of implicit causality biases by verb. Note that at this scale two extremely object-biased verbs (*shushed*, class 45.4; *upbraided*, class 33) are not visible.

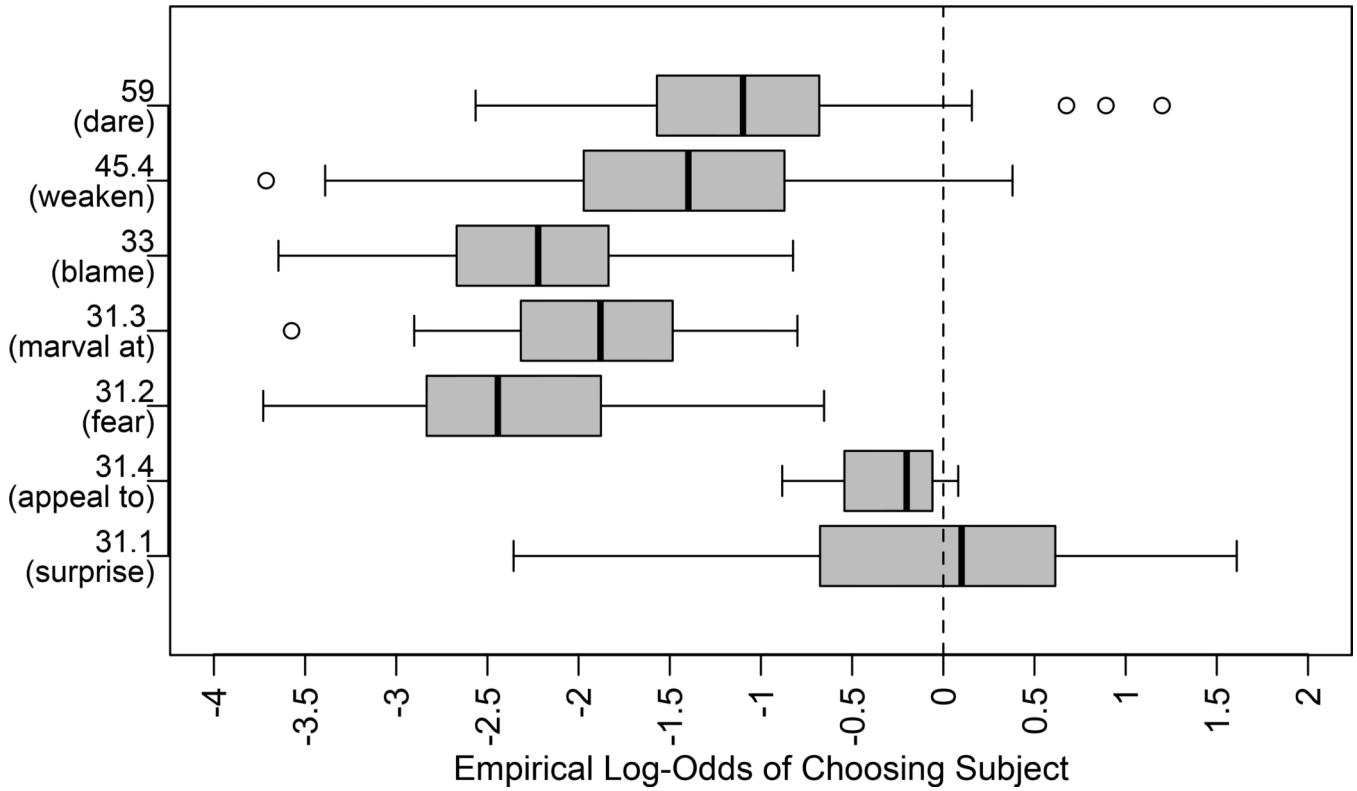

Figure 5.

Boxplots of explanation biases, analyzed separately by verb class. Note that at this scale two extremely object-biased verbs (*shushed*, class 45.4; *upbraided*, class 33) are not visible.

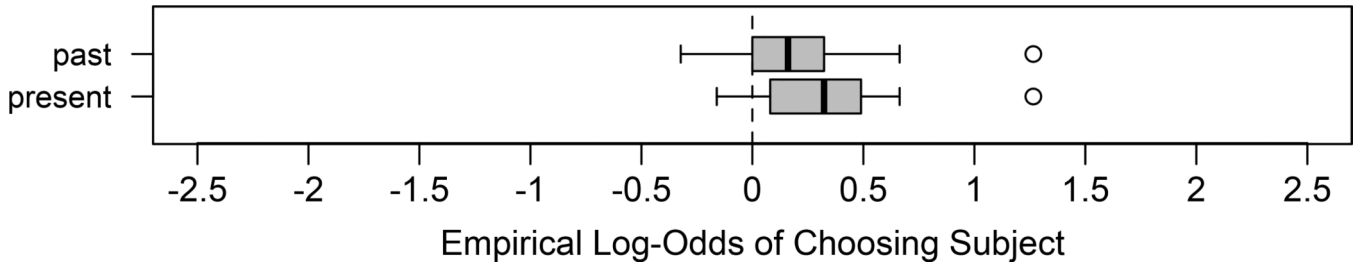

Figure 6.
Results of the follow-up experiment. Boxplots of implicit causality biases for class 59, in present and past tense.

Table 1

The verb classes investigated in Experiment 1.

Class	N	Examples
31.1	191	amuse, baffle, disappoint, surprise, worry
31.2	38	admire, dislike, envy, fear, resent, respect
31.3	61	agonize over, care about, fear for
31.4	3	appeal to, grate on, jar on
33	73	acclaim, blame, forgive, praise, scorn
45.4	95	age, animate, awaken, mellow, popularize
59	41	coax, dare, deceive, influence, trick

Note that the verb class numbers are those used by VerbNet (Kipper et al., 2006), an extension of the original work by Levin (1993).