Published in final edited form as:

Cognition. 2014 October ; 133(1): 249–261. doi:10.1016/j.cognition.2014.05.021.

Syntax and intentionality: An automatic link between language and theory-of-mind

Brent Strickland, Matthew Fisher, Frank Keil, and Joshua Knobe

Yale University

Abstract

Three studies provided evidence that syntax influences intentionality judgments. In Experiment 1, participants made either speeded or unspeeded intentionality judgments about ambiguously intentional subjects or objects. Participants were more likely to judge grammatical subjects as acting intentionally in the speeded relative to the reflective condition (thus showing an intentionality bias), but grammatical objects revealed the opposite pattern of results (thus showing an unintentionality bias). In Experiment 2, participants made an intentionality judgment about one of the two actors in a partially symmetric sentence (e.g., “John exchanged products with Susan”). The results revealed a tendency to treat the grammatical subject as acting more intentionally than the grammatical object. In Experiment 3 participants were encouraged to think about the events that such sentences typically refer to, and the tendency was significantly reduced. These results suggest a privileged relationship between language and central theory-of-mind concepts. More specifically, there may be two ways of determining intentionality judgments: (1) an automatic verbal bias to treat grammatical subjects (but not objects) as intentional (2) a deeper, more careful consideration of the events typically described by a sentence.

Keywords

Theory-of-Mind; Language; Intentionality; Cognitive Biases; Heuristic Reasoning; Semantics; Syntax; Thematic Roles

Certain aspects of theory-of-mind appear to be automatic. For example, on the basis of visual cues alone, people spontaneously ascribe certain mental states to actors that may conflict with the judgments they make after careful reflection (e.g. Heider & Simmel, 1944; Gao, McCarthy & Scholl, 2010). Such automatic processes may arise from “core knowledge” structures often studied by infancy researchers (Spelke, 2000; Kinzler & Spelke, 2007) and are seen as continuing to operate automatically into adulthood (Cherries, Mitroff, Wynn & Scholl, 2009; Flombaum & Scholl, 2006). Due to their illusory nature, these attributions can be quite different from the judgments people make after taking the time to think more deeply about a situation.

Here we ask whether *linguistic* cues might also trigger the immediate impression of intentionality in a way similar to visual cues. In particular, the current studies suggest a connection between the grammatical subject position and a representation of intentional action. This link creates a bias for stronger intentionality attribution to grammatical subjects than to non-subjects (e.g. indirect objects and direct objects).

Automaticity in theory-of-mind

Consider the disposition to attribute intentions to creatures and people in the world around us. All available evidence suggests that this is an early emerging and fundamental part of cognition. At 10 months of age, infants can represent and evaluate goal states of simple props made to move around like animate objects (Hamlin, Wynn & Bloom, 2007). Twelve-month-old infants attribute intentions and rational principles to inanimate objects navigating simple spatial environments (Gergely & Csibra, 2003). Similarly, 11-month-old infants link intentionality with order creation but not destruction (Newman, Keil, Kuhlmeier & Wynn, 2011).

As is the case with many early emerging features of cognition, certain aspects of the representation of others' intentions appear to be automatic, arising in ways that people cannot control and that may even conflict with the judgments they make after careful reflection. One prominent example comes from the Heider and Simmel (1944) displays, in which simple animations create the illusion that basic geometric forms possess intentions even if the viewer is consciously aware that this is not the case. More recently it has been shown that the visual system automatically and irresistibly picks out "chasers" in simple animated displays (Gao, McCarthy & Scholl, 2010). In other words, the perceptual apparatus locks onto a certain type of intention (i.e. chasing), which in turn serves to structure the processes of visual attention.

Just as there is a tight link between theory-of-mind representations and high-level vision, there may analogously be a tight linkage between theory-of-mind representations and grammatical structure. The current studies ask whether links between the representation of intentional action in theory-of-mind and specific syntactic positions may introduce quick and reflexive biases on intentionality judgments.

Thematic roles

Linguists employ the notion of thematic roles to explain relationships between the syntactic structure of a sentence and its underlying semantics (Carlson & Tanenhaus, 1988; Dowty, 1991; Fillmore, 1968; Gruber 1965; Schein, 2002). To begin with a simple example, it seems that there is some sense in which the objects denoted by the underlined phrases in sentences (1)-(3) all carry out similar roles in the described event.

- (1) John rolled the ball.
- (2) George dropped the coin.
- (3) Mary moved the pencil.

Although the events described by these different sentences are in many ways quite different (the first involves rolling, the second involves dropping), there does seem to be an important respect in which the objects picked out by the underlined phrases are occupying the same role within each event. Linguists have captured this intuition by suggesting that in each of these cases the underlined phrase occupies the role of *THEME* (i.e. the object that undergoes movement). The key idea of this approach is that rules of the lexicon (or mental dictionary) stipulate that when verbs like the above are used and explicitly name two actors, the *THEME* occupies the syntactic position of direct object (Carlson & Tanenhaus, 1988; Gruber, 1965).

Work in this area has examined a number of different potential thematic roles (*THEME*, *PATIENT*, *INSTRUMENT*, *EXPERIENCER*, etc.), but the one most directly relevant to the present experiments is the role of *AGENT*. This is the role occupied by the underlined phrases in sentences (4)-(6).

(4) John rolled the ball.

(5) George dropped the coin.

(6) Mary moved the pencil.

Here, as in the above examples, the events described by these sentences are all different types, but the person picked out by the underlined phrase seems to perform a similar role in each case. The underlined phrase designates an actor who intentionally or volitionally brought about some state of affairs.

In the examples provided thus far, the *THEME* has appeared as the direct object, while the *AGENT* has appeared as the grammatical subject. But there is no simple relationship between the thematic role of an argument and the syntactic position in which it is expressed. For example, consider (7)-(9):

(7) The ball rolled.

(8) The coin dropped.

(9) The pencil moved.

The underlined phrases in these latter sentences appear as grammatical subjects (instead of as objects), but they nonetheless seem to occupy the same semantic role as the underlined phrases in (1)-(3), namely, *THEME*.

Nevertheless, languages do not just work in such a way that one can arbitrarily assign any syntactic position to any thematic role. For example, suppose we try to imagine a verb *shmite* that has the converse meaning of the English verb *bite* (for similar examples see Marantz, 1984). If such a word could exist, one could use the sentence (10) to express the idea that the man bit the dog.

(10) The dog shmit the man.

In such a hypothetical sentence, the subject would occupy the role *THEME*, while the direct object would occupy the role *AGENT*. While it would be logically possible for a verb like *shmite* to exist in English, we find no such examples. Indeed, within every language, the mappings between thematic roles and syntactic positions are highly predictable and do not

vary arbitrarily from verb to verb. Such regular mappings are the reason why just looking at the syntax of this sentence alone, one can make inferences about the likely assignment of thematic roles.

Indeed, the literature on syntactic bootstrapping has shown that young children are able to exploit such regular mappings from a very early age in order to make educated guesses at the meaning of verbs. For example, when 26 month-old children hear a sentence like “The duck is gorging the bunny” they are more likely to look at a picture representing a causal action (i.e. the duck doing something to the bunny) than to a synchronous action (i.e. the duck and bunny each performing an action with no causal interaction between the two). However, when children hear “The duck and bunny are gorging,” the children no longer preferentially direct their attention to the causal event (Naigles, 1990; see also Gertner, Fisher, & Eisengart, 2006). Similarly, when learning the meaning of a new word, children have a strong expectation for grammatical subjects (but not other grammatical positions) to be intentional, animate actors (Childers & Echols, 2004).

From thematic roles to theory-of-mind

Different theories have been developed about the precise nature of this AGENT role (e.g. Dowty 1991; Jackendoff, 1987). Some theories have focused on the ways in which it might be connected to causation (Gruber, 1965; Jackendoff, 1972) and event initiation (Cruse, 1973). In addition to these, one other factor that a number of theorists have thought to be an important part of the notion of an AGENT is intentionality. It is thus widely agreed that there is some relationship between the thematic role AGENT and the notion of acting intentionally (Gruber, 1965; Holisky, 1987; Dowty, 1991; Hopper & Thompson, 1980).

Within the literature in theoretical linguistics, there have been a number of conflicting proposals about the precise relationship between thematic roles and their associated properties. The earliest theories posited a list of distinct thematic roles (AGENT, THEME, EXPERIENCER, etc.), with distinct properties associated with each role (e.g., Gruber, 1965). Some subsequent research aimed to explain the same phenomena not by directly stipulating a list of thematic roles but rather by developing decompositional theories from which claims about role assignment can be derived (e.g., Jackendoff, 1990; Pinker, 1989). Finally, according to a third view, there is no detailed list of distinct thematic roles but only two broad ‘proto-roles’ – proto-agent and proto-patient – each defined in terms of a prototype. The ‘proto-agent’ role is associated with a set of different features (intentionality, causation, existence, etc.), and the theory says that if one argument of the verb has more of these features than the other, this argument will tend to appear in the subject position (Dowty, 1991).

Although these theoretical differences are important, they will not be essential for present purposes. Our hypothesis is that the use of syntactic cues will show the trademark properties of an automatic link to theory-of-mind. In other words, when people hear about a behavior, they may have a deeper, more reflective way of determining whether or not it was performed intentionally, but at the same time, they might also be able to make use of a simple heuristic that draws on syntactic cues. In particular, because grammatical subjects are

typically assigned the AGENT role, and because AGENTs are typically intentional, people may tend to use the heuristic that the syntactic subject is acting intentionally. Our key claim is that this heuristic will continue to operate even in cases where a more reflective process would determine that it is not applicable.¹

This hypothesis finds some initial support in a series of recent studies based on proto-role theory. Kako (2006) examined the ways in which people's judgments about whether or not a person acted intentionally could be affected by thematic role assignment. One of these studies looked at ascriptions of AGENT-like properties to either grammatical subjects or direct objects using nonsense sentences like in (11) below.

(11) The grack mecked the zarg.

In sentences like this one, participants had a strong tendency to think that the grack acted more intentionally than the zarg, thus demonstrating that when the background knowledge about the meaning of the relevant sentence is entirely absent, participants display a tendency to consider the grammatical subject to have acted intentionally. However, our hypothesis goes a step further and posits a bias that is relatively automatic and may even contradict more detailed background knowledge of the event in question. Given that grammatical subjects are often intentional agents, a fast acting heuristic of this sort would be useful in terms of quickly guiding language comprehenders to an interpretation which is likely (though not guaranteed) to be correct.

If the hypothesis is correct, there should be cases in which the two processes lead to opposite judgments. For example, there should be cases in which the simple syntactic heuristic suggests that an individual acted intentionally, while the more complex reflective process suggests that this same individual acted unintentionally. In cases of this type, people's intentionality judgments should vary depending on which of the two processes they rely on more.

We test this novel hypothesis across three experiments. Experiment 1 shows that under time pressure there is a bias to treat grammatical subjects (but not other syntactic positions) as having acted intentionally. Experiment 2 establishes that there is a bias to view grammatical subjects as acting more intentionally than equivalent actors appearing in other syntactic positions. Experiment 3 shows that this bias can be overcome by encouraging people to think logically prior to making their judgments.

Experiment 1

Here, we directly tested a strong prediction from our two process view: there should be more of a bias to treat grammatical subjects (but not other syntactic positions) as intentional under time pressure than when participants are given an adequate opportunity to reflect on the

---

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

meaning of the sentence. Time pressure is predicted to block out the possibility of deep reflection, and thus any heuristic biases should become more prevalent.

In order to do this, we adapted a method from Rosset (2008) in which participants were asked to make intentionality judgments either under time pressure or in a reflective condition in which participants were encouraged to take their time and think deeply about the meaning of the sentence in question. Just as in Rosset's original study, participants were sequentially presented with sentences like (12) and (13) that were ambiguous with regard to intentionality.

(12) Paul bumped into Taylor. (participants asked about Paul)

(13) Philip eluded Chandler. (participants asked about Chandler)

While in Rosset's original study, participants only rated the intentionality of grammatical subjects, our participants rated intentionality for both grammatical subjects and objects (between items). So, participants in our study made a binary choice as to whether the grammatical subject or the direct/indirect object acted intentionally (depending on which name was underlined in the sentence).

We predicted that for grammatical subjects, participants would show an *intentionality* bias and therefore have a tendency to rate grammatical subjects as being more intentional in the speeded relative to the unspeeded condition on the ambiguous items. However for the grammatical objects, we predicted that this would not be the case. (We were neutral about whether they would show an *unintentionality* bias or simply no bias at all.)

Methods

Participants—Forty-three native English speakers completed a paid online study through Amazon's Mechanical Turk. Participants were randomly assigned to either the speeded or the reflective condition.

Materials and Procedure—The experiment was run using Qualtrics online software.

There was a fixed set of 17 "subject" sentences in which the grammatical subject was always underlined and 17 separate "object" sentences with the direct/indirect object was always underlined. All participants viewed all 34 sentences. A full list of experimental stimuli is available in Appendix 1. For both the subject and object sentences, there were 12 ambiguously intentional sentences that were designed to be ambiguous with regards to the intentionality of the person being asked about (again whose name in the actual experiment was underlined; e.g. "Paul bumped into Taylor"). Additionally, there were 5 control sentences in which the subject clearly acted intentionally (e.g. "Susan stole from Ron"), and 5 control sentences in which the object clearly did not act intentionally (e.g. "Andy pushed Mindy over").

Within the ambiguous object sentences, 6 were intuitively fully/partially reversible in the sense that if they are true with one actor as the subject of the sentence (and second as an object), they are also true/generally true with that actor as the object (and the second actor as the subject). One example of such a sentence is "Jack French kissed Paula."² The other 6

ambiguous object sentences were not intuitively reversible in this way (e.g. “Phillip eluded Chandler”). We included both in order to ensure a relatively wide range of sentence types.

Sentences appeared in two blocks, and the sentence orders were randomized within each block. The first block consisted of only the ambiguously intentional sentences, and the second block consisted only of the control sentences. The blocks always appeared in this order so that we could avoid the possibility that the control sentences would bias the responses in the ambiguous sentences. Each sentence appeared in the center of the display with either the grammatical subject or the object underlined (according to the sentence type). In the speeded condition, participants were instructed to indicate as quickly as possible whether or not the underlined actor in the sentence acted intentionally by keypress (“b” for intentional actions and “n” for actions that were not intentional). If they failed to respond in less than 2 seconds, the screen automatically advanced to the next item. In the reflective condition participants were told that “For each sentence you will have as much time as needed to respond. Do not necessarily respond with your first impression. How quickly you respond is not important, and instead we encourage you to think deeply about each sentence and just make sure that you respond as accurately as possible.”

Results

Each participant received four scores representing the percentage of “yes” responses for each of the following four categories: ambiguous subjects, ambiguous objects, control subjects and control objects.

We began by comparing the response times in the speeded vs. unspeeded condition as a verification that our design did lead to faster responses in the former condition. This was indeed the case. The average RT in the speeded condition was 1.26 sec, while the average RT in the unspeeded condition was 2.73. These means are the same regardless of whether the average RT’s are calculated by participant or by-item. Since the distribution of response times in the speeded condition failed to meet the requirement of homogeneity of variance (Shapiro-Wilk $p < .05$), we report non-parametric statistics below. A between participant analysis (independent samples Mann-U Whitney) revealed this difference to be significant ($p < .001$). Similarly, this difference was significant on a by-item analysis (Related-samples Wilcoxon signed rank test, $p < .001$).

After analyzing RTs as a manipulation check, we turned to an analysis of the pattern of participant responses. Given that these were aggregated binary responses whose distributions failed to meet the minimal requirements for a standard ANOVA, we report non-parametric statistics. We ran binomial mixed-models with speed and grammatical position as fixed effects, and a corresponding maximal structure for the random effects per item and per subject (cf. Barr et al., 2013). We did this separately for the ambiguous and control items (see also Figure 1 below for a graphical depiction of these results).

---

¹Here we are focusing on the potential relationship between intentionality and grammatical positions, but our account is open to the possibility that other such relationships may exist, such as one between subjecthood and causation.

For the ambiguously intentional items, there was a significant interaction between speed and grammatical position ($z=3.01$, $p<.01$). Post-hoc comparisons (with means and standard deviations calculated by participant) revealed that ambiguously intentional grammatical subjects were rated as being more intentional in the speeded ($M=.7$, $SD=.25$) relative to the unspeeded condition ($M=.55$, $SD=.21$), $z=2.17$, $p<.05$. Conversely, ambiguously intentional grammatical objects were rated as being significantly less intentional in the speeded ($M=.47$, $SD=.26$) relative to the unspeeded condition ($M=.64$, $SD=.25$), $z=-2.36$, $p<.05$.

For the control items however, none of the above effects obtained. Thus there was no significant interaction between speed and grammatical position ($z=-.79$, $p=.43$). Clearly intentional subjects were rated as being (nearly) equally intentional in the speeded condition ($M=.92$, $SD=.18$) relative to the unspeeded condition ($M=.96$, $SD=.10$), $z=.24$, $p=.81$. Clearly unintentional objects were also rated as being (nearly) equally intentional in the speeded ($M=.23$, $SD=.29$) and unspeeded ($M=.19$, $SD=.21$) conditions, $z=.44$, $p=.66$.

Finally, we calculated the average intentionality of the object sentences by-item in order to separately analyze the pattern of results for objects in partially symmetric (ambiguous) object sentences (e.g. “Harry dated Lola”) and those in clearly non-symmetric (ambiguous) object sentences (e.g. “Isaac dodged Ricardo”). We did this because the unspeeded levels of intentionality were considerably higher in the partially symmetric than non symmetric sentences ($M=.73$ vs. $M=.52$, $p<.05$), and thus the non symmetric items were more closely matched to the subject sentences for the unspeeded condition ($M=.55$).

Using a within-item Wilcoxon signed rank test, we found that in both sentence types, objects were rated as being less intentional in the speeded relative to the unspeeded condition. For (partially) symmetric sentences, speeded objects were thus rated as being less intentional ($M=.63$, $SD=.13$) than unspeeded objects ($M=.73$, $SD=.15$), $p<.05$. Furthermore, this difference held for every item tested regardless of its baseline level of intentionality. For clearly non-symmetric sentences, speeded objects were also rated as less intentional ($M=.31$, $SD=.11$) than unspeeded objects ($M=.52$, $SD=.13$), $p<.05$.

Discussion

Participants demonstrated an intentionality bias for grammatical subjects, and were thus more likely to say that a grammatical subject acted intentionally when under time pressure compared to a condition in which they thought deeply about the action. However, this was not the case for grammatical objects, which instead showed an unintentionality bias, with participants rating grammatical objects as being significantly less intentional in speeded relative to unspeeded conditions.

Experiment 2

In this experiment we began by looking at pairs of symmetric and partially symmetric sentences like (14) below (see also Gleitman et al, 1996).

(14) a. John married Susan.

b. Susan married John.

The striking thing about sentences like these is that if we learn that either of them is true, we can infer that the other must be true as well. For example, if we learn that John married Susan, we can infer that Susan married John.

The strongest form of symmetry would be one in which it was logically impossible for one of the sentences to be true while the other was false. In many cases, however, we find a weaker form of symmetry. Thus, consider (15).

- (15) a. John exchanged shoes with Susan.
- b. Susan exchanged shoes with John.

Intuitively, these sentences may appear to be symmetric, but on reflection, one might conclude that it actually is possible for one of the sentences to be true while the other was false. (Perhaps John darted in and swapped his own shoes for Susan's while no one was looking.) Sentences like these are therefore sometimes described as "partially symmetric" (Dowty, 1991; Fillmore, 1966).

For present purposes, the important thing to notice about partially symmetric sentences is that they license an inference that is similar in some ways to the one found for logically symmetric sentences. For example, it might not be possible to deduce (15b) from (15a) purely based on logic and facts about the meanings of words. Nonetheless, given our knowledge of how things usually work, it is still the case that if we learn that one of these sentences is true, we could reasonably infer that the other is almost certainly true as well. Thus, these sentences do not show a strictly logical symmetry, but they do show a certain kind of symmetry all the same.

Despite the apparent symmetry in sentence pairs like (14) and (15) above, we predicted that people would be influenced by syntactic cues and would therefore make different judgments in the different cases. In particular, given that we observed an intentionality bias for grammatical subjects and a bias to think that an action is not intentional for grammatical objects, we predicted that subjects would asymmetrically be rated as acting more intentionally than objects³.

Methods

Participants—236 native English speakers participated in a paid online study through Amazon's Mechanical Turk. Three participants were removed for failing a comprehension check meant to rule out participants who were simply guessing (described in the Methods and Procedure below). An additional 30 native English speakers participated in online norming study 1, and 8 native English speaking linguists (with a master's degree or PhD) participated in a second online norming study. One participant was removed from the expert linguist questionnaire for failing all control questions (see the description of Norming study 2 below).

²At first glance, these sentences may not appear to be ambiguous with respect to the intentionality of the object. For example, in "Jack French kissed Paula," it seems relatively clear that Paula did this intentionally. Nevertheless participants have a tendency to treat objects in these cases as being less intentional than subjects, thus rendering the objects' role somewhat ambiguous. This phenomenon is studied in more detail in Experiments 2 and 3 below.

Materials and Procedure

Norming study 1: Each participant viewed 45 sentences presented in a randomized order. The set of sentences contained 23 filler items, and 22 items of interest. These items of interest were sentences using verbs that we had identified as being potentially symmetric, which were chosen from Levin (1993) based on verb class.

For each verb, the participant was required to make a necessity rating as an estimation of its symmetry. So for example, for the verb “exchange” participants would have to answer the question: “If John exchanged shoes with Susan, is it necessarily the case that Susan exchanged shoes with John?”

For each verb of interest, we calculated the percentage of participants who agreed with the symmetry of the relevant sentence and used the 12 most symmetric items as the test items in Experiments 2a-2c below. Eight of these items were lexically symmetric (i.e. the verb form does not change when the syntactic positions of the subjects and objects are reversed), but 4 (loan/borrow; buy/sell) had verb forms that did change in the reversal. These verbs are referred to as “lexical doublets” in the relevant linguistics literature (Dowty, 1991), but we nevertheless treat them as partially (semantically) symmetric for our purposes since they meet the requirement that the event referred to by one description (e.g. “Susan bought shoes from John”) can often be referred to by the other description with the positions of the names reversed (e.g. “John sold shoes to Susan”). The final verb list is shown in Table 1.

Norming study 2: In addition to testing the intuitions of linguistically naïve participants, we also examined the intuitions of linguists who have received extensive training in thinking about subtle differences between linguistic items. Participants were recruited by sending out a recruitment email to postdoctoral researchers and graduate students working in linguistics. Our sample included 5 people who have obtained a PhD in linguistics, and 2 who have obtained a master’s degree. The sub-disciplines included 3 semanticists, 1 syntactician, 1 expert in pragmatics/psycholinguistics, 1 neuro-linguist, and 1 sociolinguist.

Each participant viewed a series of sentence pairs, each presented on the screen at one time and arranged vertically (e.g. “John married Susan”; “Susan married John”). For each sentence pair participants were instructed to indicate whether that sentence pair was partially symmetric or fully symmetric by indicating on a scale of 0 – 100 what percentage of the situations typically described by one of the sentences in the pair could also be described by the other sentence in the pair. So for example if participants judged that “John married Susan” and “Susan married John” refer to exactly the same possible set of scenarios, they should choose 100 as their answer. However, if they judged that the set of situations which the sentences typically refer to were mostly (but not fully) overlapping, then they might choose a number like 80, or if they judged that they were only very slightly overlapping, then they might choose a number like 20.

Participants judged all the sentence pairs listed above, with the exception that the relevant “buy/sell” and “loan/borrow” pairs were only listed once. In addition to these ten word pairs, participants also saw the partially symmetric pairs from Study 3 below, as well as 4 control sentences (e.g. “John slapped Susan”/“Susan slapped John”) for which the answer should be

“0”. These served as control questions that allowed us to assess whether the participants were being cooperative in the task. Table 2 below lists all the verbs tested along with the percentage of linguists saying that the verb is fully symmetric and the average rating.

Experiment 2a: Each participant was randomly assigned to view only one sentence and to be tested about a single grammatical position. They thus rated a single actor in a single sentence involving a symmetric verb taken from the list above. All sentences contained “John” and “Susan” as the proper nouns designating the actors in the event. The syntactic position of the names was randomized between participants.

Each participant rated the intentionality of the action of either the grammatical subject (e.g. “John”) or the direct/indirect object (e.g. “Susan”) on a scale of 1 (not at all intentionally) to 7 (very intentionally). An example of what a participant would see on the screen (when being probed about a grammatical subject) is: “John exchanged shoes with Susan. How intentionally did John act?”

After making their response to the primary question, each participant also saw a comprehension check question in which they read the sentence “John beat up Susan.” They were then asked to indicate which of the two characters acted more intentionally. Participants who failed this question were excluded from the analyses.

Results

We analyzed the results using both by-subject and by-item non-parametric analyses (since the results again did not meet the basic assumptions required for ANOVA analyses). In an independent samples Mann-Whitney U (by-subject) analysis, grammatical subjects were judged to have acted more intentionally ($M=6.4$, $SD=.93$) than objects ($M=4.85$, $SD=1.85$), $p<.001$. For the by-item analysis, we calculated the mean rating for the grammatical subjects and objects of each verb and then conducted a related samples Wilcoxon signed rank test. This analysis also revealed that participants judged the grammatical subject as having acted more intentionally ($M=6.37$, $SD=.21$) than the grammatical object ($M=4.81$, $SD=.8$), $p<.01$. This pattern of results (with higher intentionality ratings for grammatical subjects) held for all of the verbs tested. Note that in this and all subsequent tests, the subject of “buy” was paired with the object of “sell”. This was also the case for sell/buy, loan/borrow, and borrow/loan.

To determine whether the degree of reversibility for each item correlates with the difference in judged intentionality between subjects and objects, we ran a series of Spearman’s correlational analyses. We started by asking whether the difference in intentionality between subjects and objects correlated with naïve intuitions about reversibility as assessed in Norming study 1 on a by item-basis. This correlation did not approach significance: $r_s=-.07$, p (two-tailed)=.83. We then ran a similar analysis correlating difference scores on individual items with the average (expert) symmetry judgments made by linguists in Norming study 2. This analysis also revealed no significant correlation ($r_s=-.04$, $p=.91$). We also ran a third analysis asking whether the percentage of experts agreeing with the reversibility of each item correlated with the difference in intentionality between subjects and object. This analysis also revealed no significant difference ($r_s=-.05$, $p=.88$).

Moreover, as one can appreciate in Table 3 below, the effect held for all verbs and is prominent even for purely or highly symmetric verbs as assessed by experts. For example, with the verb “marry,” the grammatical subject was rated as being .7 points higher in intentionality despite the fact that all experts agreed that this verb is completely symmetric. Thus it is unlikely that intentionality ratings can be explained by differences in perceived symmetry alone.

We also ran a final correlational analysis asking whether the percentage of naïve participants agreeing with symmetry correlated with the percentage of experts agreeing with symmetry, and found that there was indeed a marginal correlation ($r_s = -.54$, $p = .07$). This suggests that despite the fact that naïve ratings of symmetry do not predict differences in intentionality between subjects and objects, they do predict expert ratings of symmetry.

Experiment 2b

Our hypothesis is that the effect found in Experiment 2a is due to syntactic position, but one might worry that it is simply due to the order of mention, with subjects being mentioned first and thereby being viewed as more intentional. In order to test for possible effects of word order we ran control Experiment 2b in which participants saw sentences like (16).

(16) John and Susan exchanged books.

In these sentences, while John may be mentioned before Susan as in sentence (16) above, both “John” and “Susan” are grammatical subjects. If order of mention alone is sufficient for creating an intentionality bias, we would expect John to be judged as more intentional. However, if the intentionality bias is due to grammatical position, then we should expect no difference in intentionality ratings between the two nouns.

Methods

Participants—229 native English speaking paid online participants from Amazon’s Mechanical Turk completed a short survey. Seven participants were removed for failing the comprehension check.

Materials and Procedure—Experiment 2b was identical to Experiment 2a except that both names appeared as grammatical subjects, with the order of presentation of each name randomized. Thus for each verb, there were two possible orderings of “John” and “Susan”. An example sentence is: “John and Susan exchanged books.” For buy/sell and loan/borrow, the sentence incorporated “each other” in order to make it as similar to the other sentences as possible (e.g. “John and Susan bought books from each other”).

Results

Calculated by participant (in an independent samples Mann-Whitney U analysis), there was no significant difference between the first noun ($M = 5.99$, $SD = 1.29$) and the second noun ($M = 6.06$, $SD = 1.23$), $p = .76$. There was also no difference between the first noun ($M = 6.03$, $SD = .54$) and the second noun ($M = 6.03$, $SD = .48$) when analyzed by item in a related samples Wilcoxon signed rank test, $p = .72$. This suggests that syntactic position, not word order, leads to the intentionality bias in Experiment 2a.

Experiment 2c

To rule out any other possible effects of word order or word position we created sentences like (16) and (17) below.

(16) It was John that exchanged books with Susan.

(17) It was John that Susan exchanged books with.

In these sentences, we clefted either the grammatical subject (as in sentence (16)) or the grammatical object as in sentence (17). If any other low level confound is the sole producer of an intentionality bias, one would expect to see no difference between (16) and (17). However, if there is an intentionality bias for grammatical subjects over and above any low level cues, then John should be judged to be more intentional in (16) than in (17).

Methods

Participants—238 native English speaking paid online participants from Amazon's Mechanical Turk completed a short survey. Eight participants were removed for failing the comprehension check.

Materials and Procedure—Experiment 2c was identical to 2a with the following exceptions. Each verb appeared in each of the following conditions: a clefted subject ("It was John that exchanged books with Susan") or a clefted object ("It was John that Susan exchanged books with"). Thus, participants rated the intentionality of either the clefted grammatical subject or object. Again, each participant rated only one actor for one specific verb.

Results

We again analyzed the results using both by-subject and by-item analyses. In the by-subject analysis (an independent samples Mann-Whitney U analysis), grammatical subjects were judged to have acted more intentionally ($M=6.14$, $SD=1.22$) than prepositional objects ($M=4.76$, $SD=1.77$), $p<.001$. The by-item analysis also revealed that participants judged the grammatical subject as having acted more intentionally ($M=6.13$, $SD=.41$) than the prepositional object ($M=4.69$, $SD=.79$), $p<.01$. This pattern of results (with higher intentionality ratings for grammatical subjects) held for all of the verbs tested.

Furthermore, we replicated the results of the Spearman's correlational analysis from Experiment 2a. The difference in intentionality between subjects and objects (for each item) again failed to correlate with the percentage of participants agreeing with the reversibility of each item: $r_s=.14$, $p=.67$. Intentionality differences also failed to correlate with average expert reversibility ratings ($r_s=-.004$, $p=.99$), and failed to correlate with the percentage of experts agreeing with reversibility for each item ($r_s=-.09$, $p=.77$). Finally, the effect held for all items but one (see Table 3). Thus it is again unlikely that intentionality ratings can be explained by differences in perceived symmetry alone.

Discussion—Experiments 2a-2c show that participants use grammatical position as a heuristic for intentionality judgments. Participants judged the grammatical subject of

symmetric and partially symmetric sentences as acting more intentionally than the prepositional object. Thus the same actor from two equivalent (or similar) events will be treated differently depending on how that event is described. Experiments 2b and 2c show that these effects are not due to simple word order or other low level factors.

The effects reported here build on those discovered by Kako (2006) in which participants systematically judged grammatical subjects of transitive sentences involving non-words (e.g. The grack mecked the zarg.) as being more volitional than grammatical objects in similar sentences. Experiments 2a-2c extend these findings in two ways. First, they carefully rule out the possibility that Kako's original results can be explained by simple word order. Secondly, the current findings suggest that the intentionality bias arises even for verbs that people regard, as entirely symmetric. This result potentially supports the conclusion from Experiment 1 that participants may have two ways of thinking about such sentences, especially given the fact that while expert and naïve judgments about symmetry are (marginally) correlated with each other, neither type of judgment correlates with the difference in intentionality ratings between subjects and objects. In Experiment 3, we further investigate the possibility that there are two ways of assessing intentionality for sentences like these.

Experiment 3

In the previous experiment we demonstrated a bias to treat grammatical subjects as more intentional than other syntactic positions. In Experiment 3, we asked if this bias would persist when people were encouraged to think more about the symmetry of the verbs.

Here, we employed a slightly more complicated design than in the previous experiments. We first probed participants on their intentionality judgments for both grammatical subjects and objects in partially symmetric sentences as well as non-symmetric sentences. We then encouraged them to engage in deeper reflection by asking if they agreed that a different set of partially symmetric (and non-symmetric filler sentences) were in fact symmetric. Finally we re-tested their intentionality judgments on a new set of stimuli after the intervention.

We predicted that people would again rate grammatical subjects as being more intentional than objects. However, we also predicted that this propensity to judge grammatical subjects as acting more intentionally would be reduced or removed entirely after thinking deeply about the meaning of the sentences. The reason for this latter prediction is that in partially symmetric sentence pairs, the typical set of events that one description refers to (e.g. "John cuddled with Susan") can often also be referred to by the second sentence in the pair (e.g. "Susan cuddled with John"). While these overlapping reference sets may not immediately come to mind (thus leaving more room for asymmetries in intentionality), upon entering a reflective mentality, these are likely to become more salient and thus reduce differences in the attribution of intentionality to grammatical subjects and objects.

Moreover, based on our previous results we can actually make specific predictions about this process with respect to subjects and objects. Given that the types of events we describe in the current experiment are all volitional social activities (e.g. cuddling, making love, French kissing), we would expect baseline levels of intentionality to be high for both

subjects and objects. Therefore grammatical subjects would be expected to be rated close to ceiling in intentionality ratings even under deep reflection. We should thus not expect to see any difference for grammatical subjects in the pre- vs. post-intervention phase (since the heuristic bias for intentionality pushes in the same direction as the conclusion of deeper reflection). However, given that there is also an unintentionality bias for grammatical objects (Exp. 1), we should expect that upon reflection, intentionality ratings for objects should increase (since the heuristic bias in this case pushes in an opposite direction from the conclusion derived via deeper reflection).

Methods

Participants—111 native English-speaking participants from Amazon’s Mechanical Turk completed a short survey for monetary compensation. An additional 50 native English speakers participated in an online norming study.

Materials and Procedure

Norming Study and Stimulus Selection: We used participant ratings in a norming study to identify the symmetric verbs with the greatest difference between intentionality ratings for the subject and the object. These verbs were selected because they allow for the most appropriate test of the hypothesis that a logical intervention should reduce the difference in intentionality ratings between grammatical subjects and objects for symmetric verbs. For example, if we were to choose verbs in which participants naturally gave very high ratings for both the subject and object, it would be more difficult to observe a reduction in the discrepancy between these after an intervention (even if such an effect were to exist) because there would be less room for movement on intentionality ratings.

Each participant thus rated the intentionality of both subjects and objects (on a scale of 1-7) for three verbs taken randomly from a list of twelve verbs containing those with a higher than 50% necessity judgments from the Experiment 1 norming study. Participants also rated the intentionality of both subjects and objects for 6 control sentences which were clearly asymmetric. The six verbs (displayed in Table 4) with the greatest difference were used in the initial phase or the post-intervention phase of Experiment 3.

For our partially symmetric intervention items, we chose three items which were relatively high in symmetry ratings in Experiment 2, but not at ceiling: “date,” “change with,” and “trade with” (see Table 3 above).

Main Study: In the initial phase, participants saw nine sentences, three of which were partially symmetric and six of which were not (e.g. “Nina tripped Kevin”). They rated both the subject and object on how intentionally that person acted (on a scale of 1-7).

During the intervention phase, participants were then shown nine different sentences, three of which were partially symmetric and six of which were clearly asymmetric. For each sentence, participants were asked whether the sentence entails the same sentence with the actors reversed. For example, participants might be asked “If Bill swapped books with Susan, is it necessarily the case that Susan swapped books with Bill?” Participants responded “Yes” or “No.”

In the post-intervention phase, participants received another set of nine sentences, three of which were partially symmetric and six of which were not. None of the verbs from the initial phase appeared in this post-intervention phase. Participants again rated the intentionality of both the grammatical subject and grammatical object on a seven-point scale.

The relevant partially symmetric and non-symmetric sentences were grouped into two blocks of nine, and the order of these blocks was counterbalanced across participants. All orders within blocks and during the intervention phase were randomized.

Results

Eight participants were removed from the analyses due to their failure to respond correctly to more than half of the control questions in the intervention phase of the experiment.

We again performed non-parametric statistical analyses because the data failed to meet the basic assumptions required for an ANOVA. For the by-participant analysis, each participant received four scores, representing the mean of all their responses for the subjects and objects in the pre-intervention phase and the post-intervention phase. Planned comparisons using a related samples Wilcoxon signed rank test showed that while the grammatical subjects failed to differ significantly from one another in the pre- and post-intervention stages ($M=6.33$, $SD=.73$ vs. $M=6.3$, $SD=.82$, $p=.97$), grammatical objects did differ significantly ($M=4.86$, $SD=1.19$ vs. $M=5.12$, $SD=1.33$, $p<.05$). We further analyzed the difference scores (i.e. intentionality of subjects – intentionality of objects) for the pre- and post-intervention phases (again using a Wilcoxon signed rank test), and found a significant difference ($p<.05$). This analysis is important because it demonstrates that the amount of difference in intentionality between subjects and objects differed as a function of the experiment phase.

We also performed similar by-item analyses which yielded a similar pattern of results. Thus pre-intervention subjects were rated as having acted (virtually) equally intentionally ($M=6.28$, $SD=.37$) as post-intervention subjects ($M=6.27$, $SD=.34$), $p=.92$. However, pre-intervention objects were rated as having acted (marginally) less intentionally ($M=4.85$, $SD=.56$) than post-intervention objects ($M=5.05$, $SD=.52$), $p=.075$. The results of the difference score analysis also replicate ($p<.05$), with the difference between intentionality ratings for subjects and objects being greater in the pre-compared to post-intervention phase. It is also worth noting that the post-intervention difference in intentionality between subjects and objects was less than the pre-intervention difference for all the verbs we tested.

In contrast, the asymmetric control items showed no such difference between the pre- and post-intervention phases (as assessed by participant), and in fact went in the opposite direction from the partially symmetric verbs. Grammatical subjects failed to differ from one another ($M=4.35$, $SD=.82$ vs. $M=4.3$, $SD=.74$, $p=.61$). Grammatical objects did differ significantly ($M=1.99$, $SD=.93$ vs. $M=1.78$, $SD=.91$, $p<.001$); however they went in the opposite direction from the partially symmetric items, with object ratings being lower after the intervention. An analysis of the difference scores did not reveal a significant change in the differences in intentionality ratings for grammatical subjects compared to objects in the pre- vs. post-intervention phases ($p=.19$).

Discussion

The main finding from Experiment 3 was that after the intervention that encouraged participants to reflect about what would actually happen in the events described by certain sentences, they adjusted their intentionality ratings for symmetric subjects and objects by making them more similar. This was accomplished by increasing the intentionality of grammatical objects after the intervention.

Of course, our main focus here is on empirical questions about which processes lead to which judgments as opposed to normative questions about which judgments are actually correct. Nonetheless, one natural way of thinking about this result is that when in a non-reflective state of mind, grammatical subjects were heuristically assigned (correct) high intentionality ratings while objects were subjected to an unintentionality bias like that observed in Experiment 1. By entering a more reflective mode of thought, participants were able to access their deeper knowledge of the event in question and thereby reduce the asymmetry between subjects and objects by mitigating the effects of the unintentionality bias for objects.

General Discussion

Three experiments investigated the relationship between syntax and intentionality judgments. Experiment 1 showed that people are more prone to overascribe intentionality to ambiguously intentional grammatical subjects when under time pressure compared to when they think deeply about the meaning of that sentence. Conversely they are prone to under-ascribe intentionality to ambiguously intentional grammatical objects when under time pressure. Experiment 2 showed that people have a bias to treat the grammatical subject of a partially symmetric verb (like “exchange”) as being more intentional than an equivalent indirect object. Experiment 3 showed that this bias can be overcome by encouraging people to think in a more reflective mindset.

Together, these three experiments tell a coherent story. There are at least two ways of generating an intentionality judgment from verbal reports. One way involves using a heuristic judgment that is generated from certain associations between the syntactic positions and intentionality (with subjects associated with more intentionality while objects are associated with less). The second way involves considering the event described by the sentence in a deeper more reflective manner that is less susceptible to syntactic biases. The use of such a quick heuristic under time pressure may be a useful device in on-line language processing for deriving an inference that is often accurate, although it may lead people astray in exceptional cases like the ones studied here.

Possible mechanisms

What is the heuristic system picking up on that leads to the strong intentionality bias for grammatical subjects and an unintentionality bias for grammatical objects? One potential answer comes from the literature on symmetric and partially symmetric predicates. It has been noted that differences in discourse structure may introduce asymmetries into interpretation (Gleitman et al., 1996). In the context of our experiments, one might suggest that manipulating syntactic position can affect which phrase is seen as the topic of discourse

and which is seen as the focus. Perhaps it is this fact about discourse structure – rather than anything involving thematic roles – that explains the original effect.

While this argument may have some initial plausibility, Experiment 2c helps rule this out as an explanation for our effects. Clefting a name is a well-known device for focusing that name in discourse (see for e.g. Almor, 1999). Given that we found greater intentionality ratings for clefted syntactic subjects than for clefted syntactic objects in symmetric predicates, focus alone is unlikely to account for increased intentionality ratings in grammatical subjects.

A second plausible mechanism that could yield these biases is an on-line influence of thematic roles. In the case of grammatical subjects for example, perhaps the thematic role of an AGENT actively induces an intentionality bias during on-line processing. Similarly for grammatical objects, perhaps an underlying connection to a thematic role like PATIENT or THEME could account for the observed unintentionality bias in grammatical objects.

However, another possibility is that the effects discovered here are due to a form of statistical learning. Perhaps the AGENT thematic role does not actively induce the intentionality bias, but only serves to structure the lexical entries of verbs in long-term memory. People might then pick up on a statistical pattern whereby phrases that occupy a particular grammatical position tend to refer to AGENTS and therefore to entities acting intentionally, while phrases that occupy a different position tend to refer to PATIENTS or THEMES and therefore to entities acting unintentionally. (It would be somewhat difficult to say precisely which grammatical position people associate with intentional agency, but one possibility would be that people see a connection between the AGENT thematic role and ‘deep structure’ subjects; Baker, 1988.) People may gradually pick up on these statistical associations, and rely on them to some degree under time pressure.

These two views could be explicitly tested by examining verbs in which grammatical position and typical thematic role come apart. For example, in experiencer verbs like “love,” the grammatical subject is not an AGENT but is instead an EXPERIENCER (i.e. the one experiencing the emotion associated with “love”). If participants display a bias to overestimate the intentionality of EXPERIENCER subjects in speeded relative to unspeeded situations, this would be evidence for the statistical association view. If participants show less of an intentionality bias for EXPERIENCER subjects under time pressure than for AGENT subjects, then this would be evidence for the view that the effect is driven by true thematic role ascriptions.

The Intentionality Bias Effect: A Theory-of-Mind Reflex?

Rosset’s (2008) recent findings on the intentionality bias bear directly on related issues concerning the nature of theory-of-mind. Her third experiment was nearly identical to our Experiment 1 in that it also tested people’s judgments for ambiguously intentional actions in speeded and reflective conditions. In that experiment, participants displayed a bias to interpret the described actions as being intentional under time pressure. However, only the grammatical subject position was tested. In our Experiment 1, when examining grammatical objects we found the opposite of an intentionality bias.

Thus, while the original Rosset results were interpreted as showing that there is a general bias to interpret all human action as being intentional, those findings might be more simply explained by an intentionality bias that applies to syntactic subjects but is not a trait of theory-of-mind more generally.

Even if Rosset's results could potentially be re-explained in this way however, a general intentionality bias might also still exist. For example, a general intentionality bias might be present, but is suppressed for grammatical objects (as in our Experiment 1). Further studies could more closely examine this issue by using non-verbal stimuli involving videos or pictures.

Two Systems

Dual process theory has proposed that there are two systems for reasoning, with one system acting quickly and unreflectively and the other acting slowly and deliberately (Sloman, 1996; Kahneman, 2011). Some researchers have suggested that this general approach can also be applied in the study of theory-of-mind (Cohen & German, 2009; Apperly & Butterfill, 2009). They suggest that there is a quick heuristic system for automatically calculating others' mental states, alongside a slower more effortful system. It is possible that our present results could be explained within a theory of this sort. On this view, a fast-acting heuristic system would act on linguistic information but not deeper real world knowledge while a second, slower system would be less influenced by grammatical position and instead rely on a deeper understanding of the real world events that particular sentences refer to.

However it is also possible that the present results could be explained without a dual-processing account. Instead of multiple systems, there could be a single system that prioritizes various intentionality cues differently. Under time pressure, only those cues highest on the list of priorities may be employed. If verbal cues were ranked more highly on this list than cues from a deeper knowledge base, then that could potentially explain these results without appeal to two systems.

Existing research in other domains has employed a variety of methods in establishing the two systems account, including examining competition between systems (Cushman, 2008), cognitive load (Greene et al., 2008), and neuroimaging (Lieberman, 2003). Future work could use these strategies in deciding between the two-systems and one-system account of these particular effects and of theory-of-mind more generally.

Acknowledgments

We would like to thank Emmanuel Chemla, Bernard Victorri, Frederic Landgrain, Amit Almor, Robert Frank, Zoltán Szabó, Larry Horn, Beth Levin, Ray Jackendoff, Philippe Schlenker, David Pietraszewski, Phil Wolff and Amy Skerry for helpful comments on this work. We would also like to thank Nina Wexelblatt, Ike Silver, Florian Pellet and Alison Goldblatt for help collecting these data. This research was supported in part by the National Institute of Health's Grant R37-HD023922 to Frank C. Keil.

References

- Almor A. Noun-phrase anaphora and focus: The informational load hypothesis. *Psychological Review*. 1999; 106(4):748–765. [PubMed: 10560327]

- Apperly IA, Butterfill SA. Do humans have two systems to track beliefs and belief-like states? *Psychological Review*. 2009; 116(4):953–970. [PubMed: 19839692]
- Baker, MC. *Incorporation: A Theory of Grammatical Function Changing*. The University of Chicago Press; Chicago, Illinois: 1988.
- Barr DJ, Levy R, Scheepers C, Tily H. Random-effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*. 2013; 68(3):255–278.
- >Brainard D. The psychophysics toolbox. *Spatial Vision*. 1997; 10:433–436. [PubMed: 9176952]
- Carlson, G.; Tanenhaus, M. Thematic roles and language comprehension. In: Wilkins, W., editor. *Syntax and Semantics, Vol. 21: Thematic Relations*. Academic Press; 1988. p. 263-89.
- Cherries, EW.; Mitroff, SR.; Wynn, K.; Scholl, BJ. Do the same principles constrain persisting object representations in infant cognition and adult perception?: The cases of continuity and cohesion. In: Hood, B.; Santos, L., editors. *The origins of object knowledge*. Oxford University Press; 2009. p. 107-134.
- Childers JB, Echols CH. Two ½ year- old children use animacy and syntax to learn a new noun. *Infancy*. 2004; 5:109–125.
- Cohen AS, German TC. Encoding of others' beliefs without overt instruction. *Cognition*. 2009; 111:356–363. [PubMed: 19376507]
- Cruse DA. Some thoughts on agentivity. *Journal of Linguistics*. 1973; 9:1–204.
- Cushman FA. Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*. 2008; 108:353–380. [PubMed: 18439575]
- Dowty DR. Thematic proto-roles and argument selection. *Language*. 1991; 67:547–619.
- Fillmore, C. The Ohio State University project on linguistic analysis, report (13). Ohio State University; Columbus: 1966. *Toward a modern theory of case*; p. 1-24.
- Fillmore, C. The case for case. In: Bach, E.; Harms, R., editors. *Universals in Linguistic Theory*. 1968. p. 1-90.
- Fillmore, C. The case for case reopened. In: Cole, P.; Sadock, J., editors. *Syntax and Semantics 8: Grammatical Relations*. Academic Press; New York: 1977. p. 59-82.
- Flombaum JI, Scholl BJ. A temporal same-object advantage in the tunnel effect: Facilitated change detection for persisting objects. *Journal of Experimental Psychology: Human Perception & Performance*. 2006; 32(4):840–853. [PubMed: 16846283]
- Gao T, McCarthy G, Scholl BJ. The wolfpack effect: Perception of animacy irresistibly influences interactive behavior. *Psychological Science*. 2010; 21(12):1845–1853. [PubMed: 21078895]
- Gergely G, Csibra G. Teleological reasoning in infancy: The naive theory of rational action. *Trends in Cognitive Sciences*. 2003; 7:287–292. [PubMed: 12860186]
- Gertner Y, Fisher C, Eisengart J. Learning words and rules: Abstract knowledge of word order in early sentence comprehension. *Psychological Science*. 2006; 17:684–691. [PubMed: 16913951]
- Gleitman L, Gleitman H, Miller C, Ostrin R. 'Similar' and similar concepts. *Cognition*. 1996; 58:321–376. [PubMed: 8871343]
- Greene JD, Morelli SA, Lowenberg K, Nystrom LE, Cohen JD. Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*. 2008; 107:1144–1154. [PubMed: 18158145]
- Grice, HP. *Logic and conversation*. In: Cole, P.; Morgan, J., editors. *Syntax and Semantics. Vol. 3*. Academic Press; 1975. 1989. p. 22-40. Reprinted as ch.2 of Grice
- Gruber, JS. *Studies in lexical relations*. 1965. PhD Dissertation, MIT, Mass
- Hamlin JK, Wynn K, Bloom P. Social evaluation by preverbal infants. *Nature*. 2007; 450:557–559. [PubMed: 18033298]
- Heider F, Simmel M. An experimental study of apparent behavior. *American Journal of Psychology*. 1944; 57:243–259.
- Holisky DA. The case of the intransitive subject in Tsova-Tush (Batsbi). *Lingua*. 1987; 71:103–132.
- Hopper PJ, Thompson SA. Transitivity in grammar and discourse. *Language*. 1980; 56:251–299.
- Jackendoff, R. *Semantic interpretation in generative grammar*. MIT Press; Cambridge: 1972.
- Jackendoff R. The status of thematic relations in linguistic theory. *Linguistic Inquiry*. 1987; 18:369–412.

- Jackendoff, R. Semantic structures. MIT Press; Cambridge, MA: 1990.
- Kahneman, D. Thinking, fast and slow. Allen Lane; London: 2011.
- Kako E. Thematic role properties of subjects and objects. *Cognition*. 2006; 101:1–42. [PubMed: 16289069]
- Levin, B.; Rappaport-Hovav, M. Unaccusativity: At the syntax-lexical semantics interface. MIT Press; Cambridge, MA: 1995. *Linguistic Inquiry Monograph* 26
- Levin, B.; Rappaport Hovav, M. Argument realization. Cambridge University Press; Cambridge, UK: 2005.
- Lieberman, MD. Reflective and reflexive judgment processes: A social cognitive neuroscience approach. In: Forgas, JP.; Williams, KR.; von Hippel, W., editors. *Social judgments: Implicit and explicit processes*. Cambridge University Press; New York: 2003. p. 44–67.
- Marantz, A. On the nature of grammatical relations. MIT Press; Cambridge, MA: 1984.
- Naigles L. Children use syntax to learn verb meanings. *Journal of Child Language*. 1990; 17:357–314. [PubMed: 2380274]
- Newman GE, Keil FC, Kuhlmeier V, Wynn K. Sensitivity to design: Early understandings of the link between agents and order. *Proceedings of the National Academy of Sciences*. 2010; 107:17140–17145.
- Noveck IA, Reboul A. Experimental pragmatics: A Gricean turn in the study of language. *Trends in Cognitive Sciences*. 2008; 12(11):425–431. [PubMed: 18805727]
- Parsons T. Thematic relations and arguments. *Linguistic Inquiry* 26.4. 1995:635–662.
- Pelli D. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*. 1997; 10:437–442. [PubMed: 9176953]
- Pelphrey KA, Morris JP, McCarthy G. Grasping the intentions of others: The context of a perceived action influences activity in the superior temporal sulcus during social perception. *Journal of Cognitive Neuroscience*. 2004; 16(10):1706–1716. [PubMed: 15701223]
- Phillips W, Barnes JL, Mahajan N, Yamaguchi M, Santos LR. ‘Unwilling’ versus unable’: Capuchin monkeys’ (*Cebus apella*) understanding of human intentional action. *Developmental Science*. 2009; 12(6):938–945. [PubMed: 19840049]
- Pinker, S. Learnability and cognition: The acquisition of argument structure. The MIT Press; Cambridge: 1989.
- Rosset E. It’s no accident: Our bias for intentional explanations. *Cognition*. 2008; 108:771–780. [PubMed: 18692779]
- Samson D, Apperly IA, Braithwaite J, Andrews B. Seeing it their way: Evidence for rapid and involuntary computation of what other people see. *Journal of Experimental Psychology: Human Perception and Performance*. 2010; 36(5):1255–1266. [PubMed: 20731512]
- Saxe R, Kanwisher N. People thinking about thinking people: The role of the temporo-parietal junction in “theory of mind”. *NeuroImage*. 2003; 19:1835–1842. [PubMed: 12948738]
- Schein, B. Events and the semantic content of thematic relations. In: Preyer, G.; Peter, G., editors. *Logical Form and Language*. Oxford University Press; Oxford: 2002. p. 263–344.
- Sloman SA. The empirical case for two systems of reasoning. *Psychological Bulletin*. 1996; 119:3–22.
- Spelke ES. Core knowledge. *American Psychologist*. 2000; 55:1233–1243. [PubMed: 11280937]
- Spelke ES, Kinzler KD. Core knowledge. *Developmental Science*. 2007; 10:89–96. [PubMed: 17181705]
- Wynn K. Do infants have numerical expectations or just perceptual preferences? *Developmental Science*. 2002; 2:207–209.

Highlights

- We ask if syntactic position influences the nature of the intentionality bias.
- Under time pressure, there is an intentionality bias for grammatical subjects but an unintentionality bias for grammatical objects (study 1).
- For logically symmetric sentences (study 2), subjects are viewed as more intentional than objects. This is true even when controlling for word order.
- This bias can be mitigated when participants are encouraged to think reflectively (study 3).

Experiment 1: Intentionality in subjects vs. objects

Figure 1. Proportions of “intentional” answers for each condition in Experiment 1. Error bars show SE mean.

Table 1

The twelve most symmetric verbs from the norming study.

Verb	% Agreeing with symmetry
Exchange	94
Trade	94
Date	94
Sell – Buy	94
Marry	94
Buy – Sell	91
Borrow – Loan	84
Partner	84
Change	78
Swap	78
Loan – Borrow	72
Barter	71
Average	86

Table 2

Expert judgements on symmetry

Verb (verb pair)	% of experts agreeing 100% with symmetry	Average symmetry rating
Marry	100	100
Partner	71	98
Sell/buy	71	97
Swap	57	92
Exchange	57	92
Date	57	91
Trade	43	92
Change	29	84
Correspond	29	80
Loan/borrow	14	79
Barter	14	72
Cuddle	0	80
Make love	0	68
French kiss	0	67
Kill (control)	0	0
Slap	0	0
Lie to	0	0
Harm	0	0

Table 3

Columns 2 and 3 depict differences in rated intentionality between grammatical subjects and objects for Experiments 2a and 2c respectively. A positive score indicates a higher intentionality rating for grammatical subjects.

Verb(pair)	Exp. 2a Sub.- Obj	Exp. 2c Sub.- Obj.	%Experts agreeing fully with symmetry	Avg. expert symmetry rating	% Naïve agreeing with symmetry
Marry	.7	.6	100	100	94
Partner	3.1	2.57	71	98	84
Buy/sell	.2	1.56	71	97	91
Sell/buy	1.89	1.28	71	97	94
Swap	1.35	.8	57	92	78
Exchange	2.67	.71	57	92	94
Date	1.27	1.74	57	91	94
Trade	1.53	1.66	43	92	94
Change	2.07	0	29	84	78
Borrow/loan	.57	3.79	14	79	84
Loan/borrow	1.67	1.99	14	79	72
Barter	1.68	.52	14	72	71

Table 4

The six verbs with the greatest difference between subject and object intentionality rating. We have also included the symmetry ratings collected for these items in Norming studies 1 and 2 (during Experiment 2 above).

Verb	Subject Intentionality	Object Intentionality	Expert avg. symmetry	%Naive participants agreeing with symmetry
Cuddle	6.73	4.53	80	53
Partner	5.63	3.50	98	84
French kiss	5.88	3.88	67	55
Borrow	6.79	4.86	79	84
Correspond	6.29	4.64	80	53
Make love to	6.23	4.85	68	63

Appendix 1

Experimental Stimuli with mean proportion of intentionality responses

Ambiguous Subject Sentences	Speeded	Unspeeded
<u>Zach</u> preceded Dara	M=.93	M=.45
<u>Tristan</u> impressed Alexis	M=.65	M=.75
<u>Paul</u> bumped into Taylor	M=.57	M=.5
<u>Hunter</u> inspired Ruby	M=.75	M=.7
<u>Garrett</u> hindered Keaton	M=.67	M=.65
<u>John</u> passed by Sue	M=.74	M=.65
<u>Sam</u> ran into Tricia	M=.8	M=.45
<u>Isaac</u> scraped Abby	M=.59	M=.55
<u>Edward</u> crashed into Megan	M=.53	M=.45
<u>Charles</u> came upon Johanna	M=.74	M=.45
<u>Gary</u> attracted Sandra	M=.68	M=.5
<u>Dave</u> astonished Jenna	M=.74	M=.45
Ambiguous Object Sentences		
<i>Non-symmetric</i>		
Isaac dodged <u>Ricardo</u>	M=.17	M=.45
Seth avoided <u>Casey</u>	M=.41	M=.4
Phillip eluded <u>Chandler</u>	M=.29	M=.55
Jacob hid from <u>Nicholas</u>	M=.35	M=.5
Logan fled from <u>Rodney</u>	M=.41	M=.75
Cole consulted <u>Brittney</u>	M=.2	M=.45
<i>Partially symmetric</i>		
Harry dated <u>Lola</u>	M=.61	M=.8
Jack French kissed <u>Paula</u>	M=.67	M=.7
John married <u>Sarah</u>	M=.78	M=.95
Laura exchanged stories with <u>Danielle</u>	M=.67	M=.9
Daniel teamed up with <u>Betsy</u>	M=.65	M=.65
Gary connected with <u>Belle</u>	M=.4	M=.55
Control Subject Sentences		
<u>Susan</u> stole from Ron	M=.91	M=.95
<u>Nick</u> murdered Rhonda	M=.1	M=.1
<u>Casey</u> lied to Mandy	M=.95	M=.9
<u>Cole</u> bullied Abby	M=.91	M=.1
<u>Logan</u> warned Virginia	M=.85	M=.95
Control Object Sentences		
Sandra poisoned <u>Gary</u>	M=.24	M=.15
Erin insulted <u>Chandler</u>	M=.14	M=.2
Robert cut <u>Annie</u>	M=.2	M=.25
Seth offended <u>Randy</u>	M=.22	M=.2
Andy pushed <u>Mindy</u> over	M=.22	M=.1