

Neural Interaction between Logical Reasoning and Pragmatic Processing in Narrative Discourse

Jérôme Prado¹, Nicola Spotorno¹, Eric Koun¹, Emily Hewitt¹,
Jean-Baptiste Van der Henst¹, Dan Sperber^{2,3},
and Ira A. Noveck¹

Abstract

Logical connectives (e.g., *or*, *if*, and *not*) are central to everyday conversation, and the inferences they generate are made with little effort in pragmatically sound situations. In contrast, the neural substrates of logical inference-making have been studied exclusively in abstract tasks where pragmatic concerns are minimal. Here, we used fMRI in an innovative design that employed narratives to investigate the interaction between logical reasoning and pragmatic processing in natural discourse. Each narrative contained three premises followed by a statement. In Fully-deductive stories, the statement confirmed a conclusion that followed from two steps of disjunction–elimination (e.g., Xavier considers Thursday, Friday, or Saturday for inviting his girlfriend out; he removes Thursday before he rejects Saturday and declares “I will invite her out for Friday”). In Implicated-premise

stories, an otherwise identical narrative included three premises that twice removed a single option from consideration (i.e., Xavier rejects Thursday for two different reasons). The conclusion therefore necessarily prompts an implication (i.e., Xavier must have removed Saturday from consideration as well). We report two main findings. First, conclusions of Implicated-premise stories are associated with more activity than conclusions of Fully-deductive stories in a bilateral frontoparietal system, suggesting that these regions play a role in inferring an implicated premise. Second, brain connectivity between these regions increases with pragmatic abilities when reading conclusions in Implicated-premise stories. These findings suggest that pragmatic processing interacts with logical inference-making when understanding arguments in narrative discourse. ■

INTRODUCTION

Understanding discourse often requires people to make inferences based on logical connectives, such as *or*, *if*, and *not*. Consider for example the argument below:

- (1) (a) Xavier hesitates between Thursday or Friday for taking his girlfriend Claire out to dinner.
- (b) Claire cannot go on Thursday because she has a theater class.
- (c) Xavier says: “I’ll invite her for Friday.”

Xavier’s utterance in (1c) seems almost trivial because the negation in (1b) eliminates one option from the disjunction in (1a). The only possible conclusion is that Xavier will invite Claire out for Friday.

To date, our understanding of the neural bases of such propositional reasoning relies on studies in which participants are generally asked to evaluate the validity of arguments whose content is largely abstract (Prado, Chadha,

& Booth, 2011). For example, participants are typically presented with information about arbitrary relationships (e.g., *There is a Black Square or a Yellow Circle*) and are then provided further information in the form of a second premise (e.g., *There is not a Yellow Circle*). Overall, these neuroimaging studies indicate that propositional reasoning is associated with increased activity in a frontoparietal system that includes the inferior parietal lobule (IPL; Reverberi et al., 2007; Noveck, Goel, & Smith, 2004), the medial frontal gyrus (MFG; Reverberi et al., 2010; Monti, Parsons, & Osherson, 2009; Monti, Osherson, Martinez, & Parsons, 2007), the inferior frontal gyrus (Prado, Van Der Henst, & Noveck, 2010; Reverberi et al., 2007, 2010), and the dorsal and anterior middle frontal gyrus (dmFG and aMFG, respectively; Monti et al., 2007, 2009). Importantly, a recent study demonstrated that these frontoparietal regions are not only more activated but are also more functionally connected during propositional reasoning (Cocchi et al., 2014), in keeping with the view that interactions between brain regions make critical contributions to reasoning (Bazargani, Hillebrandt, Christoff, & Dumontheil, 2014; Mackey, Miller Singley, & Bunge, 2013). Taken together, these studies have been valuable insofar as they (a) help resolve some longstanding debates about reasoning representations (i.e., they address debates about the

¹Centre National de la Recherche Scientifique and Université de Lyon, Bron, France, ²Institut Jean Nicod, Ecole Normale Supérieure, Ecole des Hautes Etudes en Sciences Sociales and Centre National de la Recherche Scientifique, Paris, France, ³Central European University, Budapest, Hungary

format in which reasoning takes place; is it done through rules or models?; Prado et al., 2011) and (b) provide the basis for describing the brain structures involved in formal reasoning tasks.

However, laboratory tasks investigating logical reasoning do not necessarily resemble the logical inference-making activity found in everyday conversation. Unlike formal syllogisms, which are typically presented as an orderly series of premises with conclusions to be evaluated, utterances serve as the basis for constructing syllogisms on the fly. That is, premise information could come from nonexplicit sources and the premises could come in an unforeseeable order (Sperber & Wilson, 1986). In other words, pragmatic processing plays a central role in everyday conversations. To make this clear, consider, for example, the following exchange (from Noveck & Sperber, 2007):

(2) *Henry*: Do you want to go on working or shall we go to the cinema?

Jane: I am tired. Let's go to the cinema.

Although it might seem obvious from the exchange that Jane is too tired (to go on working), which is critical for making a logical deduction, she never actually *said* so explicitly. All she said was that she was tired as she provided the outcome (*Let's go to the cinema*) of a hidden process. That is, the explicit information in Jane's utterance carries with it implicit information. The explicit and implicit information together provides the following syllogism:

(3) Jane will continue working or go to the cinema.

Major premise, Henry's remark

Jane is too tired [to continue working].

Minor premise, Implicated by Jane

Therefore, Jane is going to the cinema.

Conclusion (stated)

Note how Jane actually provided the conclusion (through *Let's go to the cinema*) and, by doing so, *implied* the minor premise. Without the stated conclusion, her remark about being tired would have meant that (a) she is too tired to do *anything* or (b) she is too tired to go to the cinema but not too tired to continue working (perhaps on something that requires little effort). By providing the conclusion, Jane allows Henry to infer the implicated premise and to clarify what she meant by *tired* (she implicated that *she is too tired to keep on working* but *not too tired to go to the cinema*). Such natural exchanges allow for these kinds of enrichments and impromptu syllogisms. If explicit logical features of the exchange in (3) were transformed into an abstract problem, for example, as in (4), the conclusion would simply appear unjustified and her mentioning that she is tired would seem immaterial.

(4) W or C.

(Premise)

C.

(Conclusion)

Understanding conclusions of arguments in natural discourse prompts pragmatic inferences that are—by design—minimized in classical logical tasks. Moreover, success in inferring unstated premises from utterances is likely to hinge upon readers' ability to accurately decode the speaker or writer's intention, in other words their mindreading skills (Spotorno, Koun, Prado, Van Der Henst, & Noveck, 2012; Frith & Frith, 2006). Although some participants might not have the inclination to produce intermediary premises that make logical sense of Jane's stated conclusion in (2), others will. That is, tendencies to draw implicated premises are likely to vary across individuals and impact on processing (Nieuwland, Ditman, & Kuperberg, 2010).

The present fMRI study set out to investigate the interaction between logical reasoning and pragmatic processing in natural discourse. Seven-line long stories (hereafter Disjunctive stories) were designed to allow readers to integrate logical information in two different, but closely related, ways. In the Fully-deductive version of these stories, participants read about someone faced with three possibilities and, through disjunction elimination, reduce the three down to one. For example, the dinner invitation story was presented as follows:

- (5) (a) Xavier wants to surprise Claire by inviting her to a restaurant.
 (b) He hesitates between these three days of the week:
 (c) Thursday, Friday, or Saturday.
 (d) Claire cannot go on Thursday because she has a theater class.
 (e) Also, Saturday would not work because he is having dinner with his parents.
 (f) He said: "I'll invite her for Friday."
 (g) He hopes she will be available that day.

One can notice that this story presents a brief introduction (through Lines 5a and 5b), three possibilities in the form of a disjunction (in 5c), and that the speaker then eliminates two of them (in 5d and 5e). By the time the conclusion is presented (in 5f), the participants were provided the means to draw it, making the conclusion a kind of verification.

A second condition, known as the Implicated-premise condition, was designed to prompt pragmatic inference-making at the time of the conclusion. To accomplish that, one minor adjustment was made to the stories. Instead of eliminating a second disjunct in line (5e), further justification is provided for eliminating the first one. In the story above, for example, line (5e) would replace "Saturday" with "Thursday" and read as:

- (5) (e') Also, Thursday would not work because he is having dinner with his parents.

With this version of the story, the elimination of the second disjunct (in 5e) is never explicitly stated but becomes

implied when Xavier states his conclusion in (5f). In other words, much like in the example in (2), Xavier's stated conclusion implies that he has eliminated Saturday from consideration, even if he never said so. Thus, the story still makes sense, but understanding the conclusion requires the reader to draw the premise implied by Xavier when the conclusion is presented.

We made three predictions. First, the frontoparietal regions that are typically involved in propositional reasoning (e.g., aMFG, dMFG, IPL) should show greater activity when processing the conclusion of the Disjunctive stories than for equivalent lines in Control stories in which logical arguments are not central (see

example in Table 1). Second, understanding conclusions of Implicated-premise stories—but not conclusions of Fully-deductive stories—requires participants to make further inferences, namely implicated premises. Given the call for an implied elimination, these frontoparietal regions should be more active and more functionally connected for conclusions of Implicated-premise stories than for conclusions of Fully-deductive stories. Third, given that the abilities related to drawing implicated premises reflect on pragmatic capacities, we anticipate that activity and/or connectivity between frontoparietal regions would be correlated with mindreading abilities (as measured through tests such as the Autism-spectrum Quotient

Table 1. Examples of Disjunctive (Fully-deductive and Implicated-premise) and Control Stories (Translated from French)

<i>Story Type</i>	<i>Example</i>
Disjunctive (Fully-deductive)	<p>Xavier wants to surprise Claire by inviting her out to a restaurant.</p> <p>He hesitates between these three days of the week: Thursday, Friday or Saturday.</p> <p>Claire cannot go on Thursday because she has a theater class.</p> <p>Also, Saturday would not work because he is having dinner with his parents.</p> <p>He said: "I'll invite her out for Friday."</p> <p>He hopes she will be available that day.</p> <p><i>Comprehension question:</i> In your opinion, does Xavier want to surprise Claire?</p> <p><i>Justification question:</i> In your opinion, is Xavier justified to think that he must invite Claire on Friday?</p>
Disjunctive (Implicated-premise)	<p>Xavier wants to surprise Claire by inviting her out to a restaurant.</p> <p>He hesitates between these three days of the week: Thursday, Friday or Saturday.</p> <p>Claire cannot go on Thursday because she has a theater class.</p> <p>Also, Thursday would not work because he is having dinner with his parents.</p> <p>He said: "I'll invite her out for Friday."</p> <p>He hopes she will be available that day.</p> <p><i>Comprehension question:</i> In your opinion, does Xavier want to surprise Claire?</p> <p><i>Justification question:</i> In your opinion, is Xavier justified to think that he must invite Claire on Friday?</p>
Control	<p>Alex comes back home after his soccer practice.</p> <p>He is very tired tonight and decides to relax.</p> <p>He is looking for something to do and turns on the TV.</p> <p>Unfortunately the TV is still broken.</p> <p>Instead, he chooses to read a comic book in his room.</p> <p>After a few minutes, he falls asleep.</p> <p>Alex needs a good rest to be ready for tomorrow.</p> <p><i>Comprehension question:</i> In your opinion, is Alex tired?</p>

All Control stories were followed by a comprehension question. Two thirds of the Disjunctive stories were followed by a comprehension question and one third by a justification question (see text).

[AQ]; Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001).

METHODS

Participants

Twenty right-handed adults (8 men, 12 women; mean age = 22 years) participated in the study. All participants were native French speakers, had normal or corrected-to-normal vision, and no history of neurological or psychiatric disorders. Each participant gave informed written consent before the experiment. All experimental procedures were approved by the local ethics committee (CPP Lyon Sud-Est IV).

Materials

The stimuli consisted of 60 short stories, 24 of which were Disjunctive stories, 12 were Control stories, and 24 were Filler stories. All stories were seven lines long, each with a maximum length of 91 characters (spaces included).

All Disjunctive stories contained at least one disjunction elimination argument, similar to the scenario about Xavier presented in the Introduction (see examples in Table 1). The first line of each such story introduced the background. The second line prepared the reader for the major premise (i.e., saying that there were three possibilities). This was followed by the third line, which made the three possibilities explicit (i.e., *It is A, B, or C*). The fourth line (constituting the first minor premise) eliminated one option from the major premise (e.g., *It is not A*). In half of the Disjunctive stories, the fifth line (constituting the second minor premise) eliminated another option from the major premise (e.g., *It is not B*). These stories included two disjunction–elimination inferences (i.e., *It is A, B, or C; It is not A; It is not B*) and the stated conclusion, *It is C*, which follows from the prior information. These were the Fully-deductive stories. In the other half, the option eliminated in the fourth line was eliminated again in the fifth line but with a different justification (the same one used to eliminate the second disjunct in the Fully-deductive version of each story). In effect, these stories contained just one disjunction–elimination inference (i.e., *It is A, B, or C; It is not A*). However, the conclusion in the story puts readers in a position to make sense of the speaker’s remark by assuming that the speaker made a further inference that represents a premise. Thus, these stories were members of the Implicated-premise condition. Twenty-four different story frameworks were created for the experiment (each framework involved a different situation and different characters). Each participant read either the Fully-deductive or the Implicated-premise version of each particular framework.

The final two lines of Disjunctive stories were compared with the final two lines of Control stories to define ROIs

(see below). Overall, Control stories presented simple events that occurred to a main character and were similar to Disjunctive stories in terms of number of sentences and number of propositions conveyed. However, unlike Disjunctive stories, they did not contain logical arguments (see example in Table 1). The final two lines of Control stories were similar to the final two lines of Disjunctive stories with respect to average word frequency (2654 vs. 2892, respectively, according to the LEXIQUE database; www.lexique.org; $t(34) = 0.46, p = .65$). However, the final two lines of Control stories had fewer words than Disjunctive stories (7.87 words vs. 8.44 words; $t(34) = 1.73, p = .09$) and slightly longer words than Disjunctive stories (4.99 characters vs. 4.14 characters; $t(34) = 4.38, p = .0001$). To ensure that differences in average number of words and word length did not drive activity in ROIs, we included both factors as covariates of no interest in an additional analysis of the contrast of Disjunctive versus Control stories (see Results).

Filler stories were part of a study on irony processing (Spotorno et al., 2012). These filler items were seven-line long stories that ended with utterances that were either banal or could be understood ironically or literally. This allowed us to present a variety of stories that helped disguise the focus of the present experiment.

Finally, a yes/no question was presented after each of the 60 stories, including the Control and Filler stories. For the most part, these were comprehension questions about a detail in the story (see examples in Table 1) to ensure that participants were paying attention to the stories throughout the experiment. Comprehension questions followed two thirds of the Disjunctive stories (eight in each condition). For example, the question for the story about Xavier was “In your opinion, does Xavier want to surprise Claire?” The remaining Disjunctive stories (four in each condition) were followed by a question concerning the speaker’s conclusion. For example, in the story about Xavier, the question was “In your opinion, is Xavier justified in thinking that he must invite Claire on Friday?” These justification questions were included to determine the extent to which participants noted a difference between the conclusions in the Fully-deductive and the Implicated-premise conditions.

Experimental Procedure

Participants performed the experiment in four runs of 15 stories each. Each trial started with the presentation of a visual fixation mark (i.e., a central cross) in the center of the screen. The cross was red for 7 sec, orange for 1 sec, and green for 1 sec. The first line of the story immediately followed the disappearance of the green central cross. Each line was displayed in a left-justified manner at the center of the screen. Participants read the stories line by line in a self-paced manner (i.e., each sentence remained on the screen until the participant pressed a key). There was a 0.5-sec interval between the

disappearance of a line and the presentation of the next line. After the disappearance of the last (seventh) line, a fixation mark (white central cross) reappeared for 0.5 sec. The question was then presented, prompting the participant to respond *yes* or *no*. The trial ended with the participant's response. Each trial was directly followed by a period of visual fixation (ranging from 2 to 4 sec). The presentation order of the stories was pseudorandomized, such that each run contained six Disjunctive stories (three Fully-deductive stories and three Implicated-premise stories), three control stories, and six filler stories. Among the six Disjunctive stories presented in each run, four stories (two Implicated-premise stories and two Fully-deductive stories) were followed by a standard comprehension question and two stories (one Implicated-premise story and one Fully-deductive story) were followed by a justification question. Participants were instructed to read at a normal rate. The experimental session began with three practice trials. Behavioral responses were recorded using an MR-compatible keypad placed below the right hand. Visual stimuli were generated using Presentation software (Neurobehavioral Systems, Albany, CA) and projected onto a translucent screen that was viewed by the participants through a mirror attached to the head coil.

After the fMRI session, we evaluated each participant's mindreading skills using the AQ (Baron-Cohen et al., 2001), which has been shown to be inversely correlated with the ability to correctly identify intentions from actions of others (Marsh et al., 2010). In other words, a relatively high score on the AQ questionnaire is associated with relatively low mindreading skills and a relatively low score is associated with relatively high mindreading skills. AQ scores ranged from 9 to 21 (with an average of 15), indicating that none of our participants were considered to be on the autism spectrum (Baron-Cohen et al., 2001).

Behavioral Analyses

First, we analyzed responses to comprehension and justification questions. Nonparametric testing was used because these responses were not normally distributed across conditions (Shapiro–Wilk test, $p < .05$). Second, we analyzed mean reading times of Disjunctive stories. Because these were normally distributed (Shapiro–Wilk test, $p > .05$), we performed an ANOVA with the within-subject factors Line (1 through 7) and Story version (Implicated-premise, Fully-deductive). All F tests are reported with Greenhouse–Geisser correction to correct for violations of sphericity. For all analyses, significant effects were further explored by post hoc Bonferroni tests. In such cases, Bonferroni-adjusted p values are reported (i.e., p values obtained by multiplying uncorrected p values by the number of tests performed). Bonferroni-adjusted p values of less than .05 are considered to be significant.

Imaging Acquisition and Preprocessing

Images were collected using the 1.5-T MRI system (Siemens Sonata Maestro Class; Siemens, Erlangen, Germany) of the *CERMEP Imagerie du vivant* in Lyon. The fMRI BOLD signal was measured using a T2*-weighted echo-planar sequence (repetition time = 2500 msec, flip angle = 90°, echo time = 60 msec). Twenty-six axial slices (4.40-mm thickness, field of view = 23 cm, 64 × 64 matrix) were acquired per volume. Following functional image acquisition, a high-resolution T1-weighted anatomical image (repetition time = 1880 msec, echo time = 3.93 msec, field of view = 256 mm, flip angle = 158, 176 × 256 × 256 matrix, slice thickness = 1 mm) was collected for each participant.

Data processing was performed using SPM8 (www.fil.ion.ucl.ac.uk/spm). The first four images of each run were discarded to allow for T1 equilibration effects. The remaining functional images were corrected for slice acquisition delays and spatially realigned to the first image of the first run to correct for head movements. The realigned functional images were then normalized to the SPM EPI template based in Montreal Neurological Institute (MNI) space using third-order B-spline interpolation. The quality of the normalization was verified in each participant by visually checking the registration and ensuring an adequate correspondence between each individual's brain and the MNI template. Finally, the normalized functional images were spatially smoothed with an isotropic Gaussian filter (8-mm FWHM). All coordinates are reported in MNI space.

Brain Activity Analyses

Statistical analysis was performed according to the general linear model (GLM). The last two sentences of stories (i.e., conclusion and wrap-up sentence) constituted the period of interest and were modeled as epochs with onsets time-locked to the presentation of the sixth line and with offsets time-locked to the presentation of the question. Regressors of no interest encoded activity from Lines 1 through 5. Those were modeled as epochs with onsets time-locked to the presentation of the first line and with offsets time-locked to the presentation of the sixth line. Because the task was self-paced, different regressors were constructed for each participant based on their own self-paced timings. This was done separately for the four types of stories (Fully-deductive stories, Implicated-premise stories, Control stories, and Filler stories) and for each run, yielding eight separate regressors per run. Finally, additional regressors of no interest reflecting head motion and brain activity associated with reading questions were also included in the model. All epochs were convolved with a canonical hemodynamic response function (HRF). The time series data were high-pass filtered (1/128 Hz), and serial correlations were corrected using an autoregressive AR(1) model.

Data were analyzed using an ROI approach. Specifically, brain responses associated with the last two lines of Implicated-premise and Fully-deductive stories were compared within ROIs that were involved in reading the final two lines of all Disjunctive stories (as compared with Control stories). This involved the following three steps. First, for each subject, Lines 6 and 7 of Disjunctive (Fully-deductive and Implicated-premise) stories were contrasted to Lines 6 and 7 of Control stories using a linear combination of regression coefficients associated with each type of story ([1 1 -2]). Second, individual contrasts were submitted to one-sample t tests across all participants. The resulting statistical maps were thresholded for significance using an individual voxel threshold of $p < .05$, corrected for family wise error (FWE) rate at the voxel level using the random field theory (Nichols & Hayasaka, 2003). This multiple comparison correction method was chosen to delineate ROIs because it gives much higher anatomical specificity than cluster-based thresholding (Woo, Krishnan, & Wager, 2014). Third, ROIs were defined as 6-mm spheres around the local maximum of each cluster. The main ROIs were located in the aMFG, dMFG, and dorsal IPL (dIPL) in each hemisphere (see Results).

For each participant, we calculated the average activity for Implicated-premise and Fully-deductive stories (each compared with Control stories) within an ROI by averaging the fMRI signal (beta weight) across the voxels within that ROI. Activity was entered in a GLM analysis with the within-subject factors Story version (Implicated-premise, Fully-deductive), Hemisphere (Left, Right), and ROI (aMFG, dMFG, dIPL), as well as the continuous predictor AQ scores. Because reading time of the conclusion (i.e., Line 6) differed as a function of Story version (see Results), this difference was included as a nuisance covariate in the GLM analysis to ensure that none of the effects were because of differences in performance. Our hypotheses only concerned variations of activity related to differences in Story version. Therefore, we only report main effects and interaction involving this factor. Significant effects were also explored by more restricted GLM analyses carried out in each hemisphere and by post hoc Bonferroni t tests. In such cases, Bonferroni-adjusted p values are reported (i.e., p values obtained by multiplying uncorrected p values by the number of tests performed). Bonferroni-adjusted p values less than .05 are considered to be significant.

To investigate whether any differences between Fully-deductive and Implicit-premise stories could be observed outside ROIs, we also report results of whole-brain analyses. Effects are reported at an uncorrected voxel-level threshold of $p < .001$ and a spatial extent threshold of $p < .05$, FWE-corrected for multiple comparisons.

Brain Connectivity Analyses

Hypotheses about functional connectivity were tested using an extension of the psychophysiological interaction

(PPI) approach (Gitelman, Penny, Ashburner, & Friston, 2003; Friston et al., 1997). A standard PPI analysis assesses whether certain brain areas (target regions) display activity that can be explained in terms of an interaction between the influence of a distal area (source region) and an experimental parameter. In other words, a PPI analysis tests whether activity in a source area contributes to activity in a target area to a greater (or lesser) extent in one condition versus another. In this study, we aimed to determine whether each pair of ROIs that showed differences in brain activity between story versions (i.e., Implicated-premise vs. Fully-deductive stories) also showed differences in functional connectivity. To make this possible, we used a generalization of the PPI approach called multiregional PPI (Cocchi et al., 2014). Multiregional PPI allows one to assess the connectivity between each pair of a priori defined brain regions.

Each ROI included in the multiregional PPI analysis served as a seed region in GLM analyses in which activity of each of the other ROIs was the dependent variable. Each of these GLMs included three regressors: (i) the average time series of the seed region (i.e., the “physiological” part of the PPI), (ii) the story version (Implicated-premise vs. Fully-deductive, coded as 1 and -1) after it had been convolved with a standard HRF (the “psychological” parts of the PPI), and (iii) the interaction between the physiological and psychological factors (i.e., the “interaction” part of the PPI). To compute this interaction regressor, the BOLD signal in the seed region was deconvolved by using a Bayesian estimation algorithm (Gitelman et al., 2003). The regressor coding the story version was then multiplied to the deconvolved seed activity regressor to produce the interaction term. This interaction term was then convolved with a standard HRF.

This procedure yielded a $n \times n$ connectivity matrix (where n is the number of ROIs) for each participant. Each cell (i, j) of the connectivity matrix contained the parameter estimate corresponding to the interaction term. This interaction term quantified the degree to which activity of region j (the target) was explained by the interaction between activity of region i (the seed) and the story version (Implicated-premise vs. Fully-deductive). Each set of parameter estimates (n seeds $\times n - 1$ efferent connections) in the connectivity matrix was (1) submitted to a one-sample t test and (2) correlated with AQ scores across participants. This allowed us to identify (1) the pairs of regions showing significant changes in connectivity as a function of the story version across participants as well as (2) the pairs of regions whose changes in connectivity (as a function of the number of disjunction-elimination inferences) covaried with the AQ score across participants. Note that, although the direction of the flow of information cannot be inferred from PPI analyses, a PPI analysis is not symmetric. In other words, a pair of ROI can show an increase of connectivity between a source and a target whereas the reverse analysis (using the target as a source and the source as a target) may not show a significant effect

(Cocchi et al., 2014). In the description of the results, the term “from” refers to source regions whereas the term “to” refers to target regions. p Values were corrected for multiple comparisons using the false discovery rate procedure.

RESULTS

Behavior

First, we analyzed the responses to the comprehension questions. Accuracy for these questions was 99% (range = 83–100%, $SD = 4\%$) when they followed Control stories and 97% (range = 88–100%, $SD = 5\%$) when they followed Disjunctive stories, respectively. These rates did not vary as a function of AQ (Control stories: Spearman $r = .007$, $p = .98$; Disjunctive stories: Spearman $r = -.06$, $p = .80$). However, the difference between accuracy for Control and Disjunctive stories was significant (Wilcoxon matched-pairs signed-ranks test, $z = 5.51$, $p < .001$). Therefore, remembering details of the story was harder in Disjunctive than Control stories. Nonetheless, the very high levels of accuracy observed for both Control and Disjunctive stories demonstrate that participants paid attention to both types of stories throughout the experiment.

Second, we analyzed the responses to the justification questions that followed some of the Disjunctive stories. Participants agreed with the conclusion in 90% (range = 50–100%, $SD = 17\%$) of the trials when the question followed a Fully-deductive story. Agreement rates dropped to 40% (range = 0–100%; $SD = 36\%$) when the question followed an Implicated-premise story. The difference between Implicated-premise and Fully-deductive stories was significant (Wilcoxon matched-pairs signed-ranks test, $z = 3.62$, $p < .001$), indicating that participants did note that the stated conclusions in the Fully-deductive stories were better supported by the explicit information. Nevertheless, 40% acceptance rates in the Implicated-premise condition indicates that, in their effort to make sense of the conclusion, a substantial number of participants do not distinguish between conclusions drawn from Implicated-premise and those that follow from the presented information. Interestingly, the rate of acceptance of the conclusion in Implicated-premise stories tended to be negatively correlated with AQ scores (Spearman $r = -.42$, $p = .06$). In other words, the participants with lower scores (i.e., indicating relatively strong mindreading skills) were more likely than those with higher scores (indicating that they were closer to the autistic spectrum) to accept the conclusion in the Implicated-premise condition.

Third, reading times of Disjunctive stories were analyzed using an ANOVA with the within-subject factors Line (1 through 7) and Story version (Implicated-premise, Fully-deductive). This analysis revealed a main effect of Line, $F(2.6, 49) = 54.77$, $p < .00001$, indicating that overall reading times varied between lines (see Figure 1). Most importantly, however, there was an interaction of Line \times




Figure 1. Average reading time of Disjunctive stories as a function of line and number of premises. Only Line 6 (i.e., the conclusion) was associated with longer reading time for Implicated-premise when compared with Fully-deductive stories.

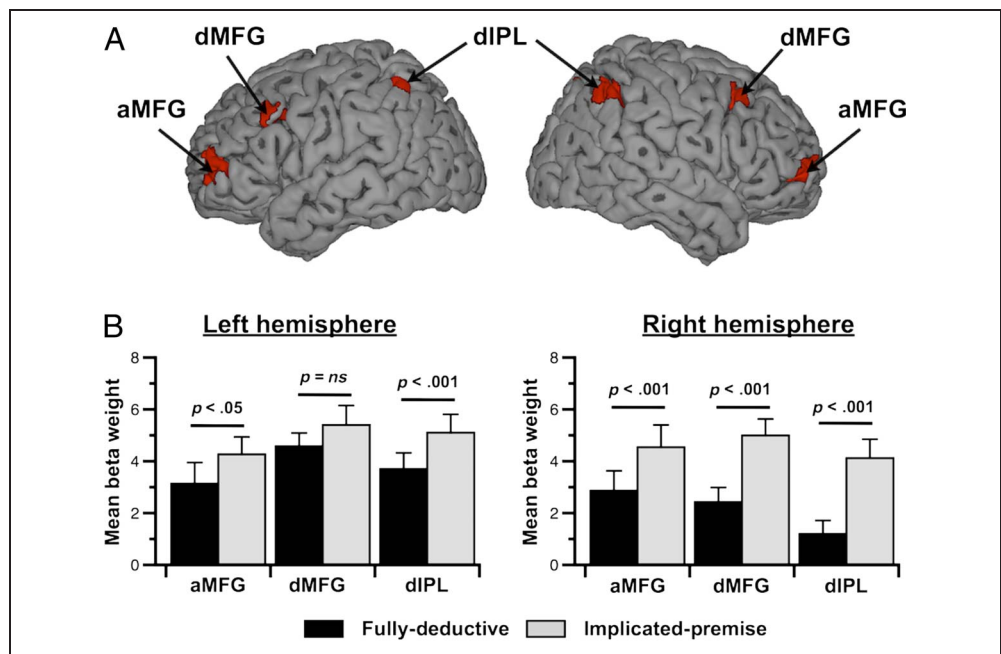
Premise, $F(4.6, 86.9) = 4.71$, $p = .001$. Post hoc Bonferroni tests revealed that reading Implicated-premise stories was longer than reading Fully-deductive stories at Line 6 (i.e., the conclusion; $p = .0005$) but not at any other line (all $ps > .39$; see Figure 1). Therefore, only the conclusion in the Implicated-premise condition prompted longer reading times than the one in the Fully-deductive condition.

ROI Definition

As detailed in the Methods, all ROIs were defined using the contrast of Disjunctive versus Control stories (using the final two lines). This contrast showed significant activation in the bilateral aMFG, bilateral dMFG, and bilateral dIPL (see Figures 2A and 3A). These regions served as ROIs in the main analyses (see Table 2 for coordinates). Disjunctive versus Control stories also elicited activation in the left Cuneus ($x = -9$, $y = -79$, $z = -26$, $Z = 6.14$) and right Precuneus ($x = 6$, $y = -67$, $z = 49$, $Z = 5.05$). These regions served as additional ROIs in separate analyses.

The final two lines of Control stories tended to contain fewer words than the final two lines of Disjunctive stories. The words used in the final two lines of Control stories were also slightly longer than those in the final two lines of Disjunctive stories (see Methods). To ensure that the definition of ROIs in the contrast of Disjunctive versus Control stories was not affected by these factors, we included number of words and word length as regressors of no interest in additional analyses of the contrasts of Disjunctive versus Control stories. This did not change any of the results above (i.e., enhanced activity for Disjunctive over Control stories was still found in bilateral aMFG, bilateral dMFG, and bilateral dIPL). Therefore, enhanced activity in these ROIs is more likely explained

Figure 2. Brain activity analysis. (A) Brain regions associated with more activity in Disjunctive than Control stories. For each hemisphere, this contrast was used to define ROIs in the aMFG, dMFG, and dIPL. (B) Brain activity in Disjunctive stories as a function of ROI, hemisphere, and story version. Greater activity for Implicated-premise than Fully-deductive stories was observed in all ROIs except the left dMFG. Mean beta weight is relative to Control stories.



by the presence of a logical conclusion than by a difference in number of words or word length.

ROI Analyses

Brain activity was analyzed using a GLM with the within-subject factors Story version (Implicated-premise, Fully-deductive), Hemisphere (Left, Right), and ROI (aMFG, dMFG, dIPL), as well as the continuous predictor AQ scores

(see Figure 2B). Difference in reading time of the conclusion (i.e., Line 6) between Implicated-premise and Fully-deductive stories was also included as a nuisance covariate. We found a main effect of Story version, $F(1, 17) = 11.38$, $p = .004$, indicating greater activity for Implicated-premise than Fully-deductive stories across ROIs. Hemisphere also interacted with Story version, $F(1, 17) = 11.56$, $p = .003$, such that the difference in activity between the Implicated-premise and Fully-deductive stories was greater in the right than in the left hemisphere. Follow-up GLM analyses were then carried out separately on left and right ROIs. In the left hemisphere, we found a main effect only for Story version, $F(1, 17) = 4.75$, $p = .044$. Post hoc Bonferroni tests indicated greater activity in Implicated-premise than Fully-deductive stories in the left aMFG ($p = .01$) and left dIPL ($p = .0008$), but the effect did not reach significance in the left dMFG ($p = .16$). In the right hemisphere, we found both a main effect of Story version, $F(1, 17) = 16.43$,




Figure 3. Brain connectivity analysis. On brain renderings are depicted the changes of functional connectivity between the regions that were more activated in Implicated-premise than Fully-deductive stories, as a function of story version and mindreading skill (1 = left aMFG; 2 = left dIPL; 3 = right aMFG; 4 = right dMFG; 5 = right dIPL). Directional arrows indicate changes of connectivity from a source to a target region (see text). (A) Across all participants, connectivity increased in Implicated-premise stories (as compared with Fully-deductive stories) from the right dMFG to the left aMFG and left dIPL, as well as from the left RLFFC to the left dIPL. (B) Connectivity in Implicated-premise stories (as compared with Fully-deductive stories) increased as a function of mindreading skill from the right dIPL to all other brain regions.

Table 2. MNI Coordinates of the ROIs Defined in the Contrast of Disjunctive versus Control Stories

Anatomical Location	~BA	MNI Coordinates			Z Score
		x	y	z	
L. aMFG	10	-36	56	7	5.64
L. dMFG	9	-48	23	37	5.36
L. dIPL	40	-42	-58	52	5.42
R. aMFG	10	33	59	4	5.59
R. dMFG	9	42	17	46	5.15
R. dIPL	40	48	-55	46	5.38

L. = left; R. = right; ~BA = approximate Brodmann's area.

$p = .008$, and an interaction of ROI \times Story version, $F(2, 34) = 4.36$, $p = .02$. Post hoc Bonferroni tests indicated greater activity in Implicated-premise than Fully-deductive stories in the right aMFG ($p < .0001$), right dMFG ($p < .00001$), and right dIPL ($p < .00001$).

Paired t tests further revealed no difference between Story versions in the two additional ROIs activated (left cuneus and right precuneus; both $t_s < 1.83$, both $p_s > .08$).

Whole-brain Analyses

Data were complementarily analyzed with whole-brain analyses. We found greater activity for Implicated-premise than Fully-deductive stories in the right dIPL ($x = 48$, $y = -64$, $z = 49$, $Z = 4.15$), right dMFG ($x = 45$, $y = 20$, $z = 49$, $Z = 3.94$), and MFG ($x = 6$, $y = 32$, $z = 43$, $Z = 4.21$).

Multiregional PPI Analyses

As detailed above, we found 5 ROIs (left aMFG, left dIPL, right aMFG, right dMFG, right dIPL) in which activity was greater for Implicated-premise stories than Fully-deductive stories. To test whether connectivity between these 5 ROIs differed between conditions, we measured the patterns of connectivity across ROIs using multiregional PPI analyses (Cocchi et al., 2014). Across participants, time series of these 5 ROIs were positively correlated with each other (average $r = .58$, $SD = 0.06$), and there was no significant difference among correlation coefficients ($\chi^2 = 15.41$, $p = .08$). Correlation coefficients were as follows: Left aMFG–Left dIPL = .58, Left aMFG–Right aMFG = .66, Left aMFG–Right dMFG = .50, Left aMFG–Right dIPL = .53, Left dIPL–Right aMFG = .51, Left dIPL–Right dMFG = .54, Left dIPL–Right dIPL = .65, Right aMFG–Right dMFG = .55, Right aMFG–Right dIPL = .58, Right dMFG–Right dIPL = .66. More importantly, we determined whether connectivity between each pair of ROIs varied as a function of Story version (Implicated-premise, Fully-deductive; see Methods). After correction for multiple comparisons, we found three pairs of ROIs in which connectivity was greater in Implicated-premise stories than Fully-deductive ones (see Figure 3A). Two of these connections were from the right dMFG (source region) to target regions of the left hemisphere: left aMFG ($p = .039$) and left dIPL ($p = .042$). In other words, right dMFG activity contributed to activity in both left aMFG and left dIPL to a greater extent in Implicated-premise stories than Fully-deductive stories. Left aMFG activity also contributed to left dIPL activity to a greater extent in Implicated-premise than Fully-deductive stories ($p = .042$). All three connections were associated with a larger positive relationship between ROIs during Implicated-premise stories than Fully-deductive ones, as indicated by positive betas for each of the physiological terms. None of the other connections differed as a function of the number of premises in Disjunctive stories (all $p_s > .129$).

Second, we assessed whether the AQ score was predictive of differences in connectivity between Implicated-premise and Fully-deductive stories among these ROIs. We found six pairs of ROIs for which differences in connectivity between Implicated-premise and Fully-deductive stories were negatively correlated with AQ score across participants (see Figure 3B). In other words, there was greater connectivity among these pairs of ROIs for participants with relatively high mindreading skills (i.e., further from the autistic spectrum) than for participants with relatively low mindreading skills in Implicated-premise stories (relative to Fully-deductive stories). Four of these connections were from the right dIPL (source region) to target regions of the left and right hemisphere: right aMFG ($p = .009$), right dMFG ($p = .009$), left dIPL ($p = .009$), and left aMFG ($p = .025$). Connections from the right dMFG (source region) to the left dIPL ($p = .028$) and right aMFG ($p = .015$) were also greater in Implicated-premise stories compared with the Fully-deductive ones. All connections were associated with a larger positive relationship between ROIs during Implicated-premise stories than Fully-deductive ones as a function of mindreading skill, as indicated by positive betas for each physiological term. None of the other connections were significantly modulated by AQ score (all $p_s > .11$).

Control Analyses

As described above, we found both activity and connectivity differences between conclusions of Implicated-premise and Fully-deductive stories. However, it is important to note that conclusions of Implicated-premise were also associated with longer reading times than conclusions of Fully-deductive stories. This is unlikely to have affected our results because we included differences in conclusion reading time as nuisance covariate in the second-level GLM analysis. However, to further ensure that this difference did not have an effect on our results, we ran another set of first-level analyses adding the trial-by-trial reading time of the conclusion as regressor of no interest for each subject. The results obtained with these additional analyses did not alter any of the results obtained with our initial model. All of the ROIs that were significantly more active in Implicated-premise versus Fully-deductive stories and all the connections that were significant (and varied with AQ) in the main analysis remained so when reading times were included as a covariate. This indicates that none of our results were due to differences in reading times.

DISCUSSION

Propositional arguments are most often understood in a context of discourse comprehension where pragmatic processing plays a central role. However, their neural substrates have always been studied in abstract tasks where

pragmatic concerns are minimal. Here we used fMRI to explore the neural interaction between logical reasoning and pragmatic processing in narrative discourse. Participants read short stories in which we embedded disjunction–elimination arguments. In the Fully-deductive version of the stories, three disjuncts were reduced to one through two disjunction–elimination steps, making the conclusion a confirmation. The remaining Disjunctive stories only included one disjunction–elimination, such that the same conclusion allowed for an Implicated-premise.

Reading a Conclusion that Is Not Entirely Grounded by Explicit Information Elicits Enhanced Activity in and Connectivity between Frontoparietal Regions

Consistent with prior work (e.g., Prado et al., 2011; Monti et al., 2009), we found more activity for Disjunctive than for Control stories in a frontoparietal network including the aMFG, dMFG, and dIPL. Although there was more activity for conclusions of Implicated-premise than Fully-deductive stories in most of these regions, our results point to some functional heterogeneity within this brain system. First, the difference in activity between Implicated-premise stories and Fully-deductive ones was larger in right-lateralized than in left-lateralized regions, suggesting an important role for the right hemisphere in reconciling the stated conclusion with the prior information. Second, in Implicated-premise stories (as compared with Fully-deductive stories), connectivity analyses revealed that activity in both left-lateralized regions was influenced by activity in two separate right-lateralized source regions: the dMFG and the dIPL. Because mindreading skill did not affect these connections to the same extent, dMFG and dIPL are likely to have different functions.

The Roles of the Right dMFG and Left aMFG in Processing Nonconfirming Conclusions

Activity in the right dMFG contributed to activity in the left aMFG to a greater extent in Implicated-premise than Fully-deductive stories. The size of this contribution was not correlated with mindreading skill and was thus similar across all participants. Although both right dMFG and left aMFG are often reported in neuroimaging studies of logical reasoning (Goel, 2007), studies indicate that these regions might have different functions. On the one hand, it has been suggested that the left aMFG (i.e., the lateral part of BA 10/46 and 10/47) is involved in logical inference-making (Monti et al., 2007, 2009; Rodriguez-Moreno & Hirsch, 2009), especially when arguments require several logical steps to convert premises into conclusions (Prado, Mutreja, & Booth, 2013; Monti et al., 2007, 2009). On the other hand, the right lateral pFC (including the dMFG) is typically involved in situations in which the conclusion of an argument conflicts with what might be expected, either because of prior beliefs (Stollstorff, Vartanian, & Goel, 2012;

Goel & Dolan, 2003; Goel, Buchel, Frith, & Dolan, 2000) or because the conclusion is not entirely supported by the premises (Goel, Stollstorff, Nakic, Knutson, & Grafman, 2009; Parsons & Osherson, 2001). For example, the right lateral pFC is activated when the conclusion of a syllogism is not consistent with one's beliefs about the world (Goel et al., 2000). It is also engaged when the validity of a conclusion cannot be entirely determined by preceding premises, whether these are relational (Goel et al., 2009) or propositional (Parsons & Osherson, 2001). Overall, these findings are consistent with growing evidence that the right lateral pFC is sensitive to conflict and uncertainty in information processing (Bach, Seymour, & Dolan, 2009; Shackman, McMenamin, Maxwell, Greischar, & Davidson, 2009; Huettel, Stowe, Gordon, Warner, & Platt, 2006). Here, we found increased activity in and connectivity between the right dMFG and the left aMFG when the stated conclusion was not based on a fully developed argument but would be sensible with further pragmatic processing (i.e., in Implicated-premise stories). Thus, we propose that the right dMFG might detect a conflict between background information and a stated conclusion. It is possible that this region might signal the left aMFG to resolve this conflict by inferring the premise that was implied by the speaker or writer. This is consistent with the fact that there was more activity in the aMFG in Implicated-premise than in Fully-deductive stories. It also suggests that the left aMFG might support the cognitive operations that allow readers to infer a conclusion from explicit premises (Prado et al., 2013; Monti et al., 2009; Rodriguez-Moreno & Hirsch, 2009) as well as to infer implicated premises from nonconfirming conclusions.

How Mindreading Interacts with the Reasoning Brain System

In abstract logical tasks, participants might resolve conflicts by simply considering an argument invalid and rejecting it (especially when the task explicitly requires such an evaluation). Unlike in abstract tasks, however, producing arguments in spoken or written discourse is an act of communication, and every act of communication involves intentions from speakers or writers (Sperber & Wilson, 1986). Therefore, an utterance expressing a speaker's conclusion that is not predictable by context (as is the case in Implicated-premise stories) is not abandoned or rejected on purely logical grounds. Rather, it prompts listeners to make an implicit premise that the speaker intends them to make. In keeping with this idea, we found that the reader's general ability to mindread (i.e., to make further inferences based on the speaker or writer's intention) predicted connectivity between the brain regions involved in the conclusion stage of Implicated-premise stories (relative to Fully-deductive stories, whose conclusion followed from the prior information). Specifically, the better the reader was at mindreading (as indexed by a relatively low AQ score), the more several regions

(including the left aMFG and left dIPL) were connected to the right dIPL.

Regions around the dIPL and intraparietal sulcus have been implicated in a wide range of tasks in fMRI studies. However, meta-analyses suggest that these regions play a key role in the maintenance and manipulation of information in working-memory (Wager & Smith, 2003) and might be involved in the short-term storage of premises during logical reasoning. Our task clearly required participants to store the premises in working memory because these disappeared after they were read. In Implicated-premise stories, retrieving these premises from working memory (e.g., *There is an A, a B, or a C* and *There is not an A*) and confronting them when presented a conclusion (*There is a C*) was necessary to infer the implicit premise that made the argument logically valid (*There is not a B*). We speculate that the increased activity observed in the bilateral dIPL for the conclusion of Implicated-premise stories, when compared with the Fully-deductive ones, as well as the greater connectivity between right and left dIPL reflects the retrieval of the earlier premises. The enhanced connectivity between the right dIPL and the left aMFG might then reflect the integration of these premises with the presented conclusion. Critically, however, this increase in connectivity was less apparent in participants with lower mindreading skill than in participants with higher mindreading skill. This is consistent with a growing body of research associating low mindreading skill (e.g., in patients with ASD; Happé, 1993) with impaired communication between brain regions (Just, Cherkassky, Keller, & Minshew, 2004). It is also consistent with other similar effects in the neuroimaging literature viewed more largely. Nieuwland et al. (2010) found that those who scored higher on the AQ were also less likely to enrich sentence meanings of weak quantifiers pragmatically as measured by ERPs (i.e., they were less likely to transform *Some* into *Some but not all*). Here, like in Nieuwland et al. (2010), all of our participants had what can be considered normal mindreading skills. But it is possible that impaired connectivity between dIPL and left aMFG in individuals with abnormally low mindreading skills (e.g., patients with ASD) might prevent them from inferring implicit premises from conclusions.

An intriguing result is that only the right dIPL increased its connectivity with all other brain regions during conclusions of Implicated-premise stories (compared with conclusions of Fully-deductive stories), suggesting that it might play a more important role than the left dIPL in inferring implicit premises. Interestingly, the neural mechanisms supporting mindreading have consistently been found to be right-lateralized across neuroimaging studies (Van Overwalle & Baetens, 2009). For example, foci of activation typically include regions around the right TPJ (Van Overwalle & Baetens, 2009). Although the right dIPL cluster is more dorsal than the right TPJ, mindreading-related activity around the coordinates reported here has also been demonstrated in several studies

(Vistoli, Brunet-Gouet, Baup-Bobin, Hardy-Bayle, & Passerieux, 2011; Ortigue, Thompson, Parasuraman, & Grafton, 2009). Therefore, the superiority of the right over the left dIPL for understanding a speaker's conclusion in the Implicated-premise stories might result from a role for this region in mindreading or perhaps from the proximity of this region to mindreading mechanisms in the right TPJ.

On the Naturalness of Implicated-premise Stories

Our aim was to construct Implicated-premise stories that ultimately resemble the rather everyday exchange illustrated in (2) and that can be, at the same time, maximally similar to the Fully-deductive ones. Although these stories could appear intricate, we do not think that they are unusual nor do we think they raise any particular difficulties for participants. Nevertheless, we directly address three concerns about the Implicated-premise condition. First, it could be argued that working from three options renders the Implicated-premise condition unusual. However, this would have to be extended to the Fully-deductive condition and yet this condition arguably remains unremarkable. Second, the Implicated-premise condition seems complex because the speaker is making a declaration when there are two options left. But that is what is occurring in the example in (2), too, and that exchange appears rather everyday. Finally, although this situation arguably occurs in everyday life, it might be considered infelicitous to present participants two reasons to eliminate one option (which occurs in the fourth and fifth lines in the Implicated-premise stories to maintain comparability to the Fully-deductive ones). One could then wonder whether there might be some extra processing occurring as early as with the fifth line of Implicit-premise stories (as compared with Fully-deductive ones). This could be problematic because, although all the lines before the conclusion were modeled out in our analyses, the timing or our design and sluggishness of the HRF make it difficult to disentangle activity elicited by the fifth line (second premise) from that of the sixth line (conclusion) of stories. However, we think that the effects observed here are unlikely to be driven by greater activity for the fifth line of Implicit-premise stories. Indeed, if the fifth line of Implicated-premise stories were associated with greater processing than the fifth line of Fully-deductive stories, one would expect this line to be read more slowly in Implicated-premise stories than Fully-deductive ones. However, only the sixth line (i.e., the conclusion) differed across conditions in terms of reading times. Therefore, extra-processing in Implicit-premise stories (as compared with Fully-deductive ones) is more likely to occur upon the presentation of the sixth line rather than the fifth line. Overall, we are confident that our design is capturing a pragmatic inference that occurs in nature when processing conclusions that are not entirely supported by explicit premises, while starting off with three options.

Relationship with Prior Neuroimaging Work on Bridging Inference

In this study, we investigated the way an implicated premise is integrated specifically with logical processing. We found that this kind of implicit activity made itself apparent in neural areas that are typically associated with logical reasoning and, given that this is a pragmatic process, that this activity varied with mindreading skills. We do not assume that this implicit activity is a generalized procedure that can be applied to all cases that call for implicated premises. In this respect we distinguish the present work from neuroimaging papers on *bridging*, which focused on one sort of causal inference (e.g., Ferstl, Neumann, Bogler, & von Cramon, 2008; Virtue, Parrish, & Jung-Beeman, 2008). This bridging literature has shown that when reading, for example, that a toy was “scattered in many different pieces,” a participant is required to do more inferential work when she had read earlier that the toy had fallen ambiguously “from the kitchen” as opposed to “from the third floor” (Ferstl et al., 2008; Virtue et al., 2008). These studies have found that extra inferential effort in such cases is linked to activity in both inferior frontal gyrus and lateral temporal lobes. In this study, although we did find difference in activity between Implicated-premise and Fully-deductive stories in the frontal cortices, we did not find differences in the lateral temporal lobes during the conclusion stage. Thus, it remains unclear whether bridging inferences during discourse processing are general pragmatic processes that are similar to the type of pragmatic inferences reported here.

Conclusion

Over the past two decades or so, neuroimaging studies have identified a set of frontal and parietal regions involved in logical reasoning (Prado et al., 2011). However, because all of these studies have used abstract logical tasks where pragmatic concerns are minimal, the interaction between logical reasoning and pragmatic processing in natural discourse was unknown. The present findings significantly extend this literature by providing a system level account of the brain mechanisms that enable readers to understand logical conclusions that rely on identifiable pragmatic inferencing in a task that closely conforms to natural discourse.

Reprint requests should be sent to Jérôme Prado, Laboratoire Langage, Cerveau et Cognition, Centre National de la Recherche Scientifique, 67 Boulevard Pinel, 69675 Bron cedex, France, or via e-mail: jprado@isc.cnrs.fr.

REFERENCES

Bach, D., Seymour, B., & Dolan, R. (2009). Neural activity associated with the passive prediction of ambiguity and risk for aversive events. *The Journal of Neuroscience*:

- The Official Journal of the Society for Neuroscience*, 29, 1648–1656.
- Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The autism-spectrum quotient (AQ): Evidence from asperger syndrome/high-functioning autism, males and females, scientists and mathematicians. *Journal of Autism and Developmental Disorders*, 31, 5–17.
- Bazargani, N., Hillebrandt, H., Christoff, K., & Dumontheil, I. (2014). Developmental changes in effective connectivity associated with relational reasoning. *Human Brain Mapping*, 35, 3262–3276.
- Cocchi, L., Halford, G. S., Zalesky, A., Harding, I. H., Ramm, B. J., Cutmore, T., et al. (2014). Complexity in relational processing predicts changes in functional brain network dynamics. *Cerebral Cortex*, 24, 2283–2296.
- Ferstl, E., Neumann, J., Bogler, C., & von Cramon, D. (2008). The extended language network: A meta-analysis of neuroimaging studies on text comprehension. *Human Brain Mapping*, 29, 581–593.
- Friston, K. J., Buechel, C., Fink, G. R., Morris, J., Rolls, E., & Dolan, R. J. (1997). Psychophysiological and modulatory interactions in neuroimaging. *Neuroimage*, 6, 218–229.
- Frith, C. D., & Frith, U. (2006). The neural basis of mentalizing. *Neuron*, 50, 531–534.
- Gitelman, D. R., Penny, W. D., Ashburner, J., & Friston, K. J. (2003). Modeling regional and psychophysiological interactions in fMRI: The importance of hemodynamic deconvolution. *Neuroimage*, 19, 200–207.
- Goel, V. (2007). Anatomy of deductive reasoning. *Trends in Cognitive Science*, 11, 435–441.
- Goel, V., Buchel, C., Frith, C., & Dolan, R. J. (2000). Dissociation of mechanisms underlying syllogistic reasoning. *Neuroimage*, 12, 504–514.
- Goel, V., & Dolan, R. J. (2003). Explaining modulation of reasoning by belief. *Cognition*, 87, B11–B22.
- Goel, V., Stollstorff, M., Nacic, M., Knutson, K., & Grafman, J. (2009). A role for right ventrolateral prefrontal cortex in reasoning about indeterminate relations. *Neuropsychologia*, 47, 2790–2797.
- Happé, F. (1993). Communicative competence and theory of mind in autism: A test of relevance theory. *Cognition*, 48, 101–119.
- Huettel, S., Stowe, C., Gordon, E., Warner, B., & Platt, M. (2006). Neural signatures of economic preferences for risk and ambiguity. *Neuron*, 49, 765–775.
- Just, M., Cherkassky, V., Keller, T., & Minshew, N. (2004). Cortical activation and synchronization during sentence comprehension in high-functioning autism: Evidence of underconnectivity. *Brain: A Journal of Neurology*, 127, 1811–1821.
- Mackey, A. P., Miller Singley, A. T., & Bunge, S. A. (2013). Intensive reasoning training alters patterns of brain connectivity at rest. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33, 4796–4803.
- Marsh, A. A., Kozak, M. N., Wegner, D. M., Reid, M. E., Yu, H. H., & Blair, R. J. (2010). The neural substrates of action identification. *Social Cognitive and Affective Neuroscience*, 5, 392–403.
- Monti, M. M., Osherson, D. N., Martinez, M. J., & Parsons, L. M. (2007). Functional neuroanatomy of deductive inference: A language-independent distributed network. *Neuroimage*, 37, 1005–1016.
- Monti, M. M., Parsons, L. M., & Osherson, D. N. (2009). The boundaries of language and thought in deductive inference. *Proceedings of the National Academy of Sciences, U.S.A.*, 106, 12554–12559.
- Nichols, T., & Hayasaka, S. (2003). Controlling the familywise error rate in functional neuroimaging: A comparative review. *Statistical Methods in Medical Research*, 12, 419–446.

- Nieuwland, M. S., Ditman, T., & Kuperberg, G. R. (2010). On the incrementality of pragmatic processing: An ERP investigation of informativeness and pragmatic abilities. *Journal of Memory and Language*, *63*, 324–346.
- Noveck, I. A., Goel, V., & Smith, K. W. (2004). The neural basis of conditional reasoning with arbitrary content. *Cortex*, *40*, 613–622.
- Noveck, I. A., & Sperber, D. (2007). The why and how of experimental pragmatics: The case of “scalar inferences.” In N. Burton-Roberts (Ed.), *Advances in pragmatics* (pp. 184–212). Basingstoke, UK: Palgrave.
- Ortigue, S., Thompson, J., Parasuraman, R., & Grafton, S. (2009). Spatio-temporal dynamics of human intention understanding in temporo-parietal cortex: A combined EEG/fMRI repetition suppression paradigm. *PLoS One*, *4*, e6962.
- Parsons, L., & Osherson, D. (2001). New evidence for distinct right and left brain systems for deductive versus probabilistic reasoning. *Cerebral Cortex (New York, N.Y.: 1991)*, *11*, 954–965.
- Prado, J., Chadha, A., & Booth, J. (2011). The brain network for deductive reasoning: A quantitative meta-analysis of 28 neuroimaging studies. *Journal of Cognitive Neuroscience*, *23*, 3483–3497.
- Prado, J., Mutreja, R., & Booth, J. R. (2013). Fractionating the neural substrates of transitive reasoning: Task-dependent contributions of spatial and verbal representations. *Cerebral Cortex*, *23*, 499–507.
- Prado, J., Van Der Henst, J.-B., & Noveck, I. A. (2010). Recomposing a fragmented literature: How conditional and relational arguments engage different neural systems for deductive reasoning. *Neuroimage*, *51*, 1213–1221.
- Reverberi, C., Cherubini, P., Frackowiak, R. S. J., Caltagirone, C., Paulesu, E., & Macaluso, E. (2010). Conditional and syllogistic deductive tasks dissociate functionally during premise integration. *Human Brain Mapping*, *31*, 1430–1445.
- Reverberi, C., Cherubini, P., Rapisarda, A., Rigamonti, E., Caltagirone, C., Frackowiak, R. S. J., et al. (2007). Neural basis of generation of conclusions in elementary deduction. *Neuroimage*, *38*, 752–762.
- Rodriguez-Moreno, D., & Hirsch, J. (2009). The dynamics of deductive reasoning: An fMRI investigation. *Neuropsychologia*, *47*, 949–961.
- Shackman, A., McMenamin, B., Maxwell, J., Greischar, L., & Davidson, R. (2009). Right dorsolateral prefrontal cortical activity and behavioral inhibition. *Psychological Science*, *20*, 1500–1506.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
- Spotorno, N., Koun, E., Prado, J., Van Der Henst, J. B., & Noveck, I. A. (2012). Neural evidence that utterance-processing entails mentalizing: The case of irony. *Neuroimage*, *63*, 25–39.
- Stollstorff, M., Vartanian, O., & Goel, G. (2012). Levels of conflict in reasoning modulate right lateral prefrontal cortex. *Brain Research*, *1428*, 24–32.
- Van Overwalle, F., & Baetens, K. (2009). Understanding others’ actions and goals by mirror and mentalizing systems: A meta-analysis. *Neuroimage*, *48*, 564–584.
- Virtue, S., Parrish, T., & Jung-Beeman, M. (2008). Inferences during story comprehension: Cortical recruitment affected by predictability of events and working memory capacity. *Journal of Cognitive Neuroscience*, *20*, 2274–2284.
- Vistoli, D., Brunet-Gouet, E., Baup-Bobin, E., Hardy-Bayle, M.-C., & Passerieux, C. (2011). Anatomical and temporal architecture of theory of mind: A MEG insight into the early stages. *Neuroimage*, *54*, 1406–1414.
- Wager, T. D., & Smith, E. E. (2003). Neuroimaging studies of working memory: A meta-analysis. *Cognitive, Affective & Behavioral Neuroscience*, *3*, 255–274.
- Woo, C. W., Krishnan, A., & Wager, T. D. (2014). Cluster-extent based thresholding in fMRI analyses: Pitfalls and recommendations. *Neuroimage*, *91*, 412–419.