

The boundaries of language and thought: neural basis of inference making

Martin M. Monti^{* †} Lawrence M. Parsons[‡], and Daniel N. Osherson[†]

^{*}Cognition and Brain Sciences Unit, Medical Research Council, 15 Chaucer Rd, Cambridge CB2 7EF, UK, [†]Department of Psychology, Princeton University, Green Hall, Princeton NJ 08540, USA, and [‡]Department of Psychology, University of Sheffield, Sheffield S10 2TN, UK

Submitted to Proceedings of the National Academy of Sciences of the United States of America

Is human thought fully embedded in language, or do some forms of thought operate independently? To directly address this issue, we focus on inference making, a central feature of human cognition. In a 3T fMRI study we compare logical inferences relying on sentential connectives (e.g., not, or, if..then..) to linguistic inferences based on syntactic transformation of sentences involving ditransitive verbs (e.g., give, say, take). When contrasted with matched grammaticality judgments, logic inference alone recruited “core” regions of deduction (Brodmann Area 10p, 8m), while linguistic inference alone recruited perisylvian regions of linguistic competence, among others (BA 21, 22, 37, 39, 44, 45 and caudate). In addition, the two inferences commonly recruited a set of general “support” areas in fronto-parietal cortex (BA 6, 7, 8, 40, 47). The results indicate that logical inference is not embedded in natural language, while confirming the relative modularity of linguistic processes.

language | thought | inference making | fMRI

Abbreviations: fMRI, functional magnetic resonance imaging; LOT, language of thought

The interplay between language and thought is pivotal to the study of human cognition [1, 2, 3, 4]. A principal issue is the extent to which thinking is embedded in language [5, 6]. Within neuroscience, the matter has been explored in the domain of arithmetic [7], music [8], theory of mind [9], and deductive reasoning [10, 11]. In the case of deduction, however, the evidence is highly contradictory. Some studies report that logic inference recruits neural structures traditionally engaged by linguistic processing [10, 12]. These findings would suggest that thought is deeply rooted in language, a view consistent with certain psychological and philosophical perspectives on deduction [13, 14]. Other studies have failed to detect such activity [15, 16] which counts against the hypothesis that reasoning is tributary to language. These competing interpretations, however, are entirely based on indirect secondary observation qualitatively comparing language activations in one study to logic activations in another study across different methods, materials, paradigms and subjects. In the present study we address this issue directly by comparing, within the same group of participants, inference based on logic connectives to inferences based on the syntax and semantics of di-transitive verbs.

Consider, for example, arguments (1a) and (2a) in Table 1 (each with a single premise, above the line). The two arguments are equally valid in the sense that the truth of their respective premises guarantees the truth of their conclusions. But their validity rests on vocabulary from different lexical categories. The validity of argument (1a) is based on the sentential connectives “if then”, “or”, etc. whereas the validity of argument (2a) depends on the principal phrasal verb (and similarly for the invalidity of arguments (1b) and (2b), respectively).

	“Logic” Argument	“Linguistic” Argument
1a	$\frac{\text{If both X and Z then not Y.}}{\text{If Y then either not X or not Z.}}$	$2a \frac{\text{It was X that Y saw Z take.}}{\text{Z was seen by Y taking X.}}$
1b	$\frac{\text{If either Y or Z then not X.}}{\text{If X then not both Y and Z.}}$	$2b \frac{\text{What Z told X was Y.}}{\text{It was to Y that Z told X.}}$

The semantic contrast between the two cases is apparent although difficult to characterize precisely [17]. Roughly, connectives do not contribute to the topic of a sentence, hence appear “content-free.” In contrast, the verbs in (2a) refer to definite activities, and impart validity through the identification of semantic roles like agent, object, and patient (i.e., who did what to whom). In this report, arguments like (1a,b) will be termed “logic,” since their status depends on the sentential connectives studied in elementary logic. Arguments like (2a,b) will be termed “linguistic” inasmuch as their validity depends on comparing semantic roles across syntactic transformations.

In the study reported below, we asked healthy volunteers to evaluate a set of visually presented logic and linguistic arguments for validity (henceforth “inference task”) and, at a different time, for grammaticality (henceforth “grammar” or “baseline task”). To isolate inference making, activations for the inference and grammar tasks were contrasted over the whole brain and within specific regions of interest (ROIs).

Results

Behavioral Performance

During scanning, overall accuracy was 92.36% and 82.29% for grammar and inference trials, respectively. Judgment of validity was significantly more accurate for linguistic compared to logic arguments (89.58% and 75%, respectively; paired-t(14)=4.39, $p < 0.001$; all erroneous trials and their matched control trials were omitted from further behavioral and imag-

Reserved for Publication Footnotes

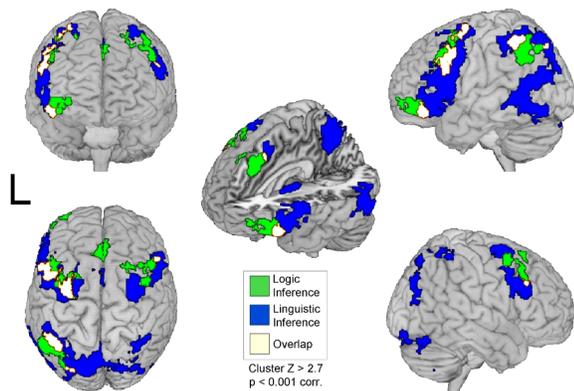


Fig. 1. Inference minus grammar contrast. Mean group activity for logic arguments (green/yellow) and linguistic arguments (blue/yellow).

ing analysis, see Experimental Procedures). On average, grammar trials lasted 8.82 s (SD = 1.96), with logic arguments requiring 8.92 s (SD = 1.84) and linguistic ones requiring 8.71 s (SD = 2.13). No significant difference in either latency or accuracy was detected between the two types of stimuli when evaluated for grammaticality. Inference trials required 12.04 s (SD = 2.65), with 12.86 s for logic arguments (SD = 3.00) and 11.21 s for linguistic arguments (SD = 2.02); the difference is marginally significant (paired $t(14) = 2.02$, $p = 0.06$). While a marginal effect of argument type (i.e., logic, linguistic) on response time may suggest differential difficulty across the two tasks, the effect is driven by two subjects whose removal from the data renders the comparison non-significant (paired $t(12) = 1.24$, $p = 0.24$). Exclusion of the foregoing subjects from the functional brain analysis does not alter the results presented below, which include the full set of participants. (See Supporting Information (SI), Section 1).

Functional Brain Activity

Comparison of inference and grammar trials for logic arguments isolated a different network than the one uncovered by the same comparison performed on linguistic arguments. Inference on linguistic arguments activated regions typically reported for linguistic processing tasks [18, 19, 20, 21]. On the other hand, logic inference did not recruit the latter areas but rather a network of regions highly similar to that reported in previous studies of deduction with sentential connectives and quantifiers [16, 22].

Logic arguments. The results of the inference minus grammar contrast are reported in Figure 1 (regions in green and yellow; see Table S1 for the list of activations). Prefrontal cortex was activated in rostro-lateral sections of the left middle and superior frontal gyri (BA 10p), along with medial frontal (BA 8), left inferior (BA 47), right superior (BA 8) and bilateral middle (BA 9) frontal gyri. Frontal activations were detected in the left middle frontal gyrus (BA 6) and right frontal gyrus (BA 6). Finally, extensive activation was also detected in posterior parietal cortex, namely, in the left superior parietal lobule (BA 7) and inferior parietal lobule (BA 40), along with intra-parietal sulcus (BA 40) and supramarginal gyrus (BA 40). (Similar results are obtained when valid and invalid trials are analyzed separately; see Fig. S2.)

Linguistic arguments. The inference minus grammar contrast, when performed on linguistic arguments, revealed a very different set of activations (regions colored blue and yellow in Figure 1; see Table S2 for the list of activations). Notably, this comparison uncovered selective activity in perisylvian linguistic regions in bilateral inferior frontal gyri (in the vicinity of Broca's area, left BA 44/45, 45, and right BA 45), left posterior superior temporal gyrus (Wernicke's area, BA 22), and left middle and superior temporal gyri (BA 21 and 39, respectively). In addition, activity was detected in the left inferior temporal gyrus (BA 37), a region implicated in semantic aspects of linguistic processing [20, 21, 23], and the left caudate head, often implicated in grammatical rule processing [19, 24].

Beyond these language-related regions, activity was also uncovered in several foci in frontal cortex. Bilateral activation was detected in middle, superior, and medial gyri (in BA 6), along with cingulate cortex (BA 32). In the left hemisphere, activation included inferior frontal (BA 9, 46, 47) and precentral (BA 6) gyri. Right hemispheric activation was seen in the middle frontal gyrus (BA 46). Activations were also detected in posterior parietal cortex spanning bilateral inferior and superior parietal lobuli (BA 40 and 7, respectively) and precuneus (BA 7). Finally, extensive activation was seen in right lateralized crus I and II in posterior cerebellum. Recent evidence suggests that the latter regions support aspects of verbal working memory or phonological encoding [25]. (Similar results are obtained when valid and invalid trials are analyzed separately; see Fig. S3.) The foregoing results fully replicate a prior, unpublished, study involving similar linguistic inferences (see SI Section 4).

Overlap in Activity for Logic and Linguistic Inference.

The activations revealed by the two contrasts (namely, inference minus grammar for logic versus linguistic stimuli) overlap in bilateral inferior and middle frontal gyri (BA 6, 9), as well as in portions of the superior and inferior parietal lobules (BA 7, 40). See Figure 1, regions in yellow. Some overlap also appears in the left inferior frontal gyrus (BA 47) but the two peaks are distinct. The peak for linguistic inference closely matches previous studies of semantic processing (-42, 26, -2 compared to -46, 30, -6 for semantic comparisons in ref.[18]), whereas that for logic inference is approximately 2 cm more anterior and ventral (-46, 44, -14). The different activations between logic and linguistic inference are not tributary to differential activation in the baseline tasks. For, no activation appears when subtracting the grammar trials for one type of argument from those of the other. This is true even when the cluster size threshold is reduced to $Z > 2.3$ and a significance level of $p < 0.05$, corrected. Furthermore, the overlay of the regions uncovered by subtracting simple fixation from either baseline task revealed two highly similar patterns of activation (see Fig. S5).

Regions of Interest Analysis. To further assess the role of the regions reported above in logic and linguistic inference, we conducted an independent region of interest (ROI) analysis on the subject-wise regression coefficients of the four effects of interest (i.e., logic inference, logic baseline, linguistic inference, and linguistic baseline; see Methods). We focused on three sets of regions: "core" areas for logic inference (implicated in processes that lie at the heart of the logic calculus), "support" areas (implicated in domain-general cognitive functions), and linguistic areas that have been linked to linguistic competence (see ref. [22] and [26], respectively).

First, a 3-way analysis of variance (ANOVA) of tasks (inference v. grammar), material (logic v. linguistic arguments) and regions (ROIs) revealed a significant interaction amongst

the three factors ($F(12,182) = 4.82, p < 0.001$). This implies that the task by materials interaction varies across ROIs. The follow-up 2-way ANOVAs, one per ROI, revealed that “core” logic regions in left rostro-lateral prefrontal cortex (BA 10p) and medial superior frontal gyrus (BA 8) exhibited the predicted interaction effect ($F(1,14) = 4.86, p < 0.05$, and $F(1,14) = 6.35, p < 0.05$; see Figure 2). Specifically, in “core” regions, the inference versus grammar subtraction yielded an effect for logic materials only ($t(14) = 3.17, p < 0.05$; and $t(14) = 3.01, p < 0.05$). This pattern of results is consistent with our hypothesis that these regions perform operations that lie at the heart of logic reasoning.

Linguistic regions in the pars triangularis and pars opercularis of the inferior frontal gyrus, the left superior temporal gyrus and the left caudate, also exhibited the expected interaction effect ($F(1,14) = 28.26, p < 0.05$; $F(1,14) = 13.45, p < 0.05$; $F(1,14) = 7.162, p < 0.05$; and $F(1,14) = 8.81, p < 0.05$; see Figure 2). That is, in contrast to “core” logic regions, the inference minus grammar comparison for the language ROIs was significant for linguistic materials only ($t(14) = -3.11, p < 0.05$; $t(14) = -3.62, p < 0.05$; $t(14) = -2.58, p < 0.05$; and $t(14) = -3.08, p < 0.05$).

Finally, “support” regions exhibited no interaction effect,

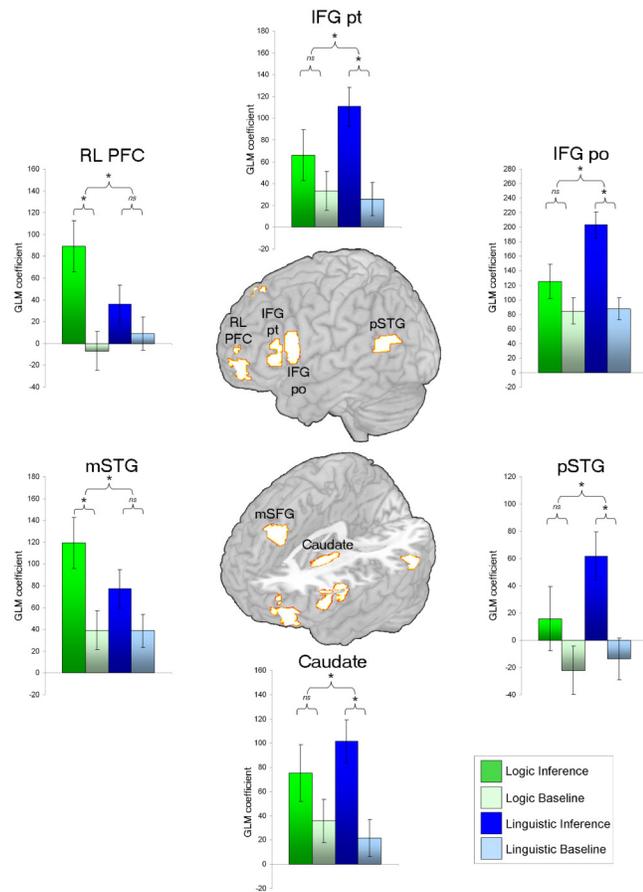


Fig. 2. ROI analysis of activity for “logic” and “linguistic” regions. Error bars represent inter-subject variability of the GLM regression coefficients (β). The ‘+’ sign indicates a significant task (inference v. grammar) by materials (logic arguments v. linguistic arguments) interaction, within the ROI. The ‘*’ sign highlights a significant simple effect of task. (RLPFC: Rostro-Lateral Prefrontal Cortex; IFG pt: Inferior Frontal Gyrus *pars triangularis*; IFG po: Inferior Frontal Gyrus *pars opercularis*; pSTG: posterior Superior Temporal Gyrus; and mSFG: medial Superior Frontal Gyrus)

consistent with the idea that these areas are implicated in general cognitive support, rather than specifically in deductive operations, either logic or linguistic (see Fig. S6).

Discussion

The interpretation of our findings depends on the status of a widely discussed thesis about the neural localization of language. The thesis (called “standard” in what follows) asserts that the principal elements of human linguistic capacity are embodied by structures proximal to the left Sylvian fissure (including inferior frontal and posterior temporal cortices). The matter is still debated inasmuch as recent experiments point to aspects of communicative competence that are supported by regions outside the traditional Wernicke-Broca circuit [19, 23, 24]. However, most neuroimaging studies aimed at localizing the central components of language comprehension do recruit areas consistent with the standard thesis. The latter studies include evaluating semantic equivalence of distinct sentences [18], morphological processing [27], detecting semantic roles [28], transforming sentence syntax [29], and comprehending discourse [30, 31].

There is, of course, a definitional component to evaluating the standard thesis. Thus, if the pragmatics of message selection [32] is counted as a core linguistic capacity then virtually any neural area implicated in cognition must be considered a “language structure.” The standard thesis, however, is generally taken to bear on a more cohesive set of abilities, namely, those allowing humans to correlate sound and meaning according to an algorithm that can be socially shared and communicated to infants. One factor in the decision to consider a given cognitive operation as linguistic in this narrower sense is whether it can be shown to be corollary to elementary grammatical principles. Inference based on ditransitive verbs, for example, corresponds to transformation of sentences *via* familiar syntactic processes like passivization and clefting. The activations elicited by such inferences are indeed largely confined to perisylvian areas, in conformity with the standard thesis. It is less clear whether deduction involving sentential connectives can be related to grammar in the same way. Quantifiers have been fruitfully analyzed as determiners [33], which has allowed grammatical elucidation of many forms of quantified inference [34]. But validity based on connectives like “if-then” has yet to be linked to syntactic mechanisms in similar fashion. Thus, classification of such deductive ability as “linguistic” versus “extra-linguistic” might be decided by comparing, as done here, its neural basis with the brain regions identified by the standard thesis as embodying the foundations of language.

We offer one more general observation before summarizing our findings about the neural representation of linguistic versus logic inference. There may well be a “language of thought” (LOT) that underlies much of human cognition [35] without LOT being structured like English or other natural languages. Even if tokens of LOT provide the semantic interpretations of English sentences, such tokens might also arise in the minds of aphasic individuals and even in other species, and may not resemble the expressions found in natural language. Hence, qualifying logical deduction as an “extra-linguistic” mental capacity is not to deny that some sort of structured representation is engaged when humans perform such reasoning. On the other hand, it is possible that LOT (in humans) coincides with the “logical form (LF)” of natural language sentences, as studied by linguists [36]. Indeed, LF (serving as the language of thought) might be pervasive in the cortex, functioning well beyond the language circuit specified by the standard thesis. In this case, the “non-linguistic” character of logical deduc-

tion would be limited to the operators that manipulate tokens of LF in view of detecting validity. As indicated by our findings, such operators would be implemented in prefrontal sites exterior to the left perisylvian circuit.

Proceeding on the basis of the standard thesis, our experimental results establish one sense in which thought does not rely on language. Obviously, when an argument is presented in verbal form, sentences must be linguistically encoded before reasoning can be initiated. However, our findings indicate that subsequently, inference involving sentential connectives relies on a circuit that is largely independent of areas recruited by semantic and syntactic processes specific to natural language. Conversely, inference based on verbal argument structure appears to make generous use of inferior frontal and posterior perisylvian regions that have previously been reported for language tasks.

Specifically, the linguistic inferences performed by our participants activated regions that are claimed to be responsible for building phrase-structure hierarchies and integrating them syntactically (namely, BA 44/45 and posterior STG, respectively [37]). Furthermore, repetitive transcranial magnetic stimulation applied to a focus closely matching our peak activation in the inferior frontal gyrus implicates this area in syntactic judgment (-50, 18, 18 here, versus -48, 16, 20 in ref.[38]). Consistent with the standard view, the linguistic system thus appears to be sufficient to draw inferences that hinge on the semantics of phrasal verbs; it does not rely on core regions for connective-based reasoning.

In striking contrast, logic inference (compared to its grammar baseline) produced no activation in the language areas described above, while recruiting regions previously proposed as “core” for inference with sentential connectives, including left rostro-lateral (BA 10p) and medial superior (BA 8) prefrontal cortices (see Fig. S17) [22]. The rostro-lateral section of polar cortex has been implicated in inferences that require integrating relational information [39, 40] and multiple sub-operations [41, 42]. Further, activity in this region has been found to peak in a time-locked fashion to the integration of multiple pieces of information [43]. The mesial superior frontal cortex has been associated with selection and coordination of multiple sub-goals [44] as well as tasks requiring choice among multiple rules to transform an initial state into a final one [45]. Moreover, these same areas have been reported for deduction with quantifiers [16], as well as for deontic reasoning [15], in experiments with very different materials and experimental design. These two regions may thus be central to constructing the derivational path that allows elementary operations to convert premises into conclusion.

Comparison of the inference minus baseline subtraction for each type of material reveals overlap in the two networks for deduction (logic v. linguistic). As described below, most of the overlap included regions that have been associated with increased cognitive load, working memory, and executive processes, as observed in our prior study. Confirming this general support function, the ROI analysis revealed no differential recruitment of these areas for logic versus linguistic inferences. Bilateral portions of the inferior and middle frontal gyri (in BA 6 and 9) have been reported to co-vary with working memory load in the absence of further mental operations [49]. In addition, left lateral BA 6 has been linked with serially updating verbal information [50] and supporting the manipulation of representations in short-term memory across numerical, spatial and verbal domains [51]. Activation in left BA 47 has been associated with working memory and executive aspects of semantic processing [52]. Activation in parietal cortex has been reported in studies of reasoning [53], and has

also been linked with maintenance of compound rules across delay periods [54]. More specifically, the inferior parietal lobule (BA 40) has been reported to support manipulation of information in working memory [51], as well as representation of numerical distance and spatial information [55]. Superior parietal regions have instead been associated with executive functions related to updating of information, maintenance of order relations and allocation of spatial attention [50]. The latter parietal regions may thus be implicated in representing the structure of arguments. Overall, in light of the foregoing independent reports, our findings lend credence to the idea that these regions serve a “support” role in inference making.

Some studies of deduction fail to uncover activity in the “core” regions noted above (e.g., [10, 46, 47]). One explanation for the discrepancy may be reliance by others on extremely simple arguments, for example, modus ponens (‘if-p-then-q’ and ‘p’, ‘therefore q’; [47]). The more challenging deductions figuring in the present experiment provoke extended and vigorous reasoning, not to be expected from elementary schemata like modus ponens. Furthermore, when such simple arguments are repeated numerous times throughout the experiment (including training sessions with feedback until 90% accuracy is achieved; e.g. [47]), pattern matching may well suffice to produce accurate responses. Indeed, studies that employ more demanding arguments report activity like ours (see above). Further methodological issues may also explain discrepant findings (see [22, 47]). Certain studies [46], for example, allow accurate responding to be based on heuristics that do not involve deduction (see [48]). Use of heuristics might be expected when participants are “instructed to respond as quickly as possible” and “arguments are drawn from the easy side of the spectrum of difficulty” ([46] p 505-6). When combined with the decision to delay the sampling, or modeling, of brain activity until the presentation of the argument’s conclusion (e.g., [10, 46]), such experimental design (i.e. simple arguments and time pressure) incurs the risk of missing much of the inferential process, which may proceed as soon as premises are read.

The dissociation in neural circuits observed in the present study suggests that inference making encompasses a heterogeneous collection of cognitive processes sensitive to the specific vocabulary upon which the inference relies (for example, sentential connectives compared to verbs). For, while arguments (1a) and (2a) above are equally valid in the sense that the truth of their premises guarantees the truth of their respective conclusions, the detection of validity in the two cases appears to rest on different neural structures, presumably embodying distinct cognitive algorithms. Our results may therefore clarify the apparent lack of replicability that characterizes earlier studies of logical inference: disparate functional activity across studies may be due to the use of deductive arguments based on different kinds of vocabularies (e.g., connectives, quantifiers and comparative adjectives).

It seems hard to reconcile the dissociation observed here with the claim that language and logic are a unitary phenomenon [56, 57] or with the belief that linguistic processes play a central role in logical deduction [13]. Our results, rather, favor theories in which the role of language in logic is limited to decoding arguments into a format suitable for language, and then encoding the result of inference back into language. In contrast, when the inference relies on the semantics of phrasal verbs, linguistic competence appears to be sufficient for drawing inferences. Indeed, neither of our two “core” regions for deduction appears in studies of discourse comprehension, even when implicit causal reasoning is involved [30, 31]. Moreover, the linguistic reasoning in the present study was explicit and semantic (like the logic reason-

ing), yet still failed to elicit activity in rostro-lateral prefrontal and medial superior frontal areas of cortex.

With respect to cognitive theories of deduction, it is relevant that none of our experiments (either here or in ref. [22]) show any sign of right parietal engagement with logic arguments. The absence of such activity is inconsistent with the Mental Models Theory of deductive reasoning [58] given the prediction that “reasoning itself, as the model theory predicts, also implicates parietal regions thought to mediate spatial representations, including regions in the right hemisphere” [53]. Rather, in each of our recent studies (as independently confirmed by others), logic inference is associated with the activation of regions that have been linked to goal/sub-goal processing in tasks requiring multiple steps and embedded operations [41] as well as coordination and selection of rules for path-finding from an initial to a final state [44, 45]. Such cognitive processes are intriguingly consistent with a “rule-based” account of deduction [59, 60, 61]. The latter theory conceives inference as the construction of a chain of structured representations within LOT leading from premises to conclusions *via* intermediate transformations licensed by elementary rules of inference.

Overall, our brain data contribute direct evidence that logic thought does not rely on natural language. These findings integrate and extend previous behavioral studies of numerical cognition [2, 3] and theory of mind [9] by providing converging evidence from a completely different, and central, domain of human cognition, that much of thought is not embedded in language.

Materials and Methods

Participants. Fifteen (7 female) Princeton University undergraduate students with no history of neurological disorder participated in the experiment for monetary compensation after giving informed consent. All subjects were right-handed, native English speakers and reported no prior training in logic.

Stimuli. Participants evaluated 32 grammatically correct arguments (half logic, half linguistic; see Table S4) for deductive validity and for syntactic well-formedness. All arguments concerned three elements: ‘X’, ‘Y’ and ‘Z’. In the logic trials, elements represented the phrasal constituents of each argument, while in the linguistic trials they represented the agent, object and patient of each statement. Assignment of X, Y, Z to constituents or roles was randomized for each trial and each individual. For each type of inference, half of the arguments were valid. Grammar trials involved all 32 well-formed arguments plus another 16 ungrammatical arguments obtained by misordering words in either the premise or the conclusion. Half of the ungrammatical trials resembled the logic arguments, and half the linguistic ones. (Grammatically incorrect trials were never analyzed and only served to make the grammar task credible.) To illustrate, a non-grammatical logic argument may include the following sentence: “If Z then either not and Y or not X.” Similarly, a non-grammatical linguistic argument may include the statement “It was by Y that to X was reported to Z.”

Task. In each trial subjects were presented with a single argument and asked to perform one of two tasks. In the inference trials, participants were required to assess the validity of the 32 arguments. In particular, subjects were told to “assume the first statement to be true, and judge whether the second statement followed out of necessity.” In the baseline trials, participants were presented with the same arguments used in the inference task, plus 16 ungrammatical arguments, and instructed to “assess whether there was any grammatical defect in either sentence of the argument.” In the latter task, participants were explicitly instructed to disregard validity and to consider each argument’s premise and conclusion separately.

Design and Procedure. In an event-related design, each trial began with a 2 s task instruction cue. Subsequently, the argument’s premise was presented alone for three seconds. The argument was then completed by adding the conclusion to the display. The entire argument was then visible for up to 15 s. Upon the subject’s response (**via** button box key-press) the trial was terminated and a 12 s fixation rest allowed the hemodynamic signal to return to baseline.

The 32 grammatically correct arguments were administered in four scans. Within each, 8 arguments of the same type (i.e., logic or linguistic) were presented twice: once for inferential evaluation, once for grammatical evaluation. In addition, in every scan four ungrammatical arguments were also presented, interspersed among the grammar trials. In half of the scans participants first assessed arguments for inferential

validity and then for grammaticality; in the other half arguments were assessed in the reverse order (counterbalanced for each subject). Seven participants assessed logic arguments in the first two scans, followed by the linguistic ones. The remaining eight underwent tasks in the reverse order. The remaining eight participants underwent tasks in the reverse order. Scans lasted an average of approximately 8.5 minutes.

Stimuli were presented using E-PrimeTM software. To assure synchrony between stimulus onset and data acquisition, the procedure was triggered by scanner pulse at the beginning of each run, and re-synchronized at the end of each trial.

fMRI Acquisition. Image data were acquired with a 3T Siemens Allegra scanner. T2* sensitive images were acquired with a gradient echo sequence (TR = 2,000 ms, TE = 30 ms, FA = 90 deg, FOV = 192 X 192 mm) in 32 ascending interleaved slices, AC-PC aligned, with a .33 distance factor and a 3 mm3 resolution. Structural images were acquired with a standard T1 sensitive MP-RAGE sequence in 176 axial slices, with a 1mm isovoxel resolution.

fMRI Data Analysis. Analysis methods were performed using FSL 4.1 (FMRIB Software Library, Oxford University [62]). Prior to functional analyses, each individual EPI time-series was brain extracted, motion corrected to the middle time point (6 parameter rigid-body transformation), and smoothed with a 5 mm FWHM kernel. Autocorrelation was corrected using a pre-whitening technique. Statistical analyses were performed using an event-related general linear model. Each individual data set was coregistered to the MNI152 standard template brain using 7 and 12 parameter optimization methods. Group mean statistics for each contrast were generated with a mixed-effects model resulting from the use of within-session variance (i.e., fixed-effects) at the single subject level and between-session variance (i.e., random-effects) at the group level. Group statistical parametric maps were thresholded using clusters determined by $Z > 2.7$ and a (corrected) cluster significance of $p < 0.001$.

For each participant two contrasts were performed: inference minus grammar for logic arguments and inference minus grammar for linguistic arguments. All trials incorrectly evaluated were excluded from the analysis. To preserve comparability between inference and grammar trials, if an argument was incorrectly evaluated in one task (logic v. grammar) it was also excluded from the other. In addition, the number of volumes analyzed across the two tasks was equalized for each argument. For grammar trials, we included the second volume through the response volume (i.e., the volume that includes the subject’s response); the first volume was excluded because only the premise appears. For the inference trials, we included the same number of volumes, counting back from the response volume.

To assess interaction effects between task (inference v. grammar) and material (logic arguments v. linguistic arguments) we adopted an ROI approach. Specific ROIs were chosen using a conjunction of functional and anatomical criteria. First, we selected all voxels uncovered by the main effect of inference (logic inference + linguistic inference minus logic grammar + linguistic grammar). The resulting regions of activity were then intersected with anatomical masks (chosen from the AAL atlas) which included “core” and “support” regions of logic inference [22] as well as regions typically reported for linguistic processing [26]. The “core” regions were in middle and superior frontal gyrus (rostro-lateral prefrontal cortex; BA 10p) and medial superior frontal gyrus (BA 8m). The “support” regions fell in bilateral superior frontal gyrus (BA 6), left inferior (BA 47) and middle (BA 8,9) frontal gyri, medial superior frontal cortex (BA 6) and left superior and inferior parietal lobuli (BA 7, 40). Language regions included the pars opercularis (BA 44) and pars triangularis (BA 45) of the left inferior frontal gyrus, left superior temporal gyrus (BA 22), and left caudate.

We stress that use of these two selection criteria assures independence of the ROI analysis from the full-brain analysis. Thus, selecting (i) anatomical regions on the basis of previous research on reasoning and language processing, and (ii) voxels uncovered by the overall effect of inference (i.e., collapsing across logic and linguistic arguments) assures that our ROIs are not biased towards either argument type.

For each region, the input to the ROI analysis was the subject-wise regression coefficient of the four effects of interest (i.e., logic inference, logic baseline, linguistic inference, and linguistic baseline). The statistical analysis was conducted in three successive steps. First, an overall 3-way ANOVA was conducted with task (inference v. grammar) and material (logic arguments v. linguistic arguments) as within factors, and ROI as the across factor. A significant interaction of the three factors indicated a different task by material interaction pattern across ROIs. This analysis was then followed by a series of 2-way ANOVAs, one per ROI, including the two within factors. Finally, significant task by material interactions were followed-up by t-test, to assess the simple effect of inference across materials.

Following our previous hypothesis [22], we expected differential activations for “core” versus “support” regions. The latter areas were expected not to be differentially recruited for inference across materials, consistently with their characterization as subserving general cognitive support. In contrast, “core” regions were expected to show a significant task by material interaction, with a significant simple effect of task for logic inferences, but not linguistic ones. Similarly, extrapolating from the literature referenced earlier, we expected linguistic regions to also show a significant task by materials interaction in the reverse direction.

1. Gordon P (2004) Numerical cognition without words: evidence from Amazonia. *Science* 306:496-9.
2. Frank MC, Everett DL, Fedorenko E, Gibson E (2008) Number as a cognitive technology: evidence from Piraha language and cognition. *Cognition* 108:819-24.
3. Butterworth B, Reeve R, Reynolds F, Lloyd D (2008) Numerical thought with and without words: Evidence from indigenous Australian children. *Proc Natl Acad Sci U S A* 105:13179-84.
4. Li P, Gleitman L (2002) Turning the tables: language and spatial reasoning. *Cognition* 83:265-94.
5. Gleitman L, Papafragou A (2005) in *Cambridge Handbook of Thinking and Reasoning*, eds Holyoak K, Morrison R (Cambridge Univ Press, Cambridge), pp 633-661.
6. Culicover PW, Jackendoff R (2006) The simpler syntax hypothesis. *Trends Cogn Sci* 10:413-8.
7. Varley RA, Klessinger NJ, Romanowski CA, Siegal M (2005) Agrammatic but numerate. *Proc Natl Acad Sci U S A* 102:3519-24.
8. Brown S, Martinez MJ, Parsons LM (2006) Music and language side by side in the brain: a PET study of the generation of melodies and sentences. *Eur J Neurosci* 23:2791-803.
9. Varley R, Siegal M (2000) Evidence for cognition without grammar from causal reasoning and 'theory of mind' in an agrammatic aphasic patient. *Curr Biol* 10:723-6.
10. Goel V, Gold B, Kapur S, Houle S (1998) Neuroanatomical correlates of human reasoning. *J Cogn Neurosci* 10:293-302.
11. Parsons LM, Osherson D (2001) New Evidence for Distinct Right and Left Brain Systems for Deductive versus Probabilistic Reasoning. *Cereb Cortex* 11:954-65.
12. Goel V, Dolan RJ (2004) Differential involvement of left prefrontal cortex in inductive and deductive reasoning. *Cognition* 93:B109-21.
13. Polk TA, Newell A (1995) Deduction as a verbal reasoning. *Psychological Review* 102:533-566.
14. Quine WV (1970) *Philosophy of logic* (Prentice Hall, Englewood Cliffs).
15. Fiddick L, Spampinato MV, Grafman J (2005) Social contracts and precautions activate different neurological systems: an fMRI investigation of deontic reasoning. *Neuroimage* 28:778-86.
16. Rodriguez-Moreno D, Hirsch J (2009) The dynamics of deductive reasoning: An fMRI investigation. *Neuropsychologia* 47:949-961.
17. Sainsbury M (1991) *Logical Forms: An Introduction to Philosophical Logic* (Blackwell, Oxford).
18. Dapretto M, Bookheimer SY (1999) Form and content: dissociating syntax and semantics in sentence comprehension. *Neuron* 24:427-32.
19. Moro A et al. (2001) Syntax and the brain: disentangling grammar by selective anomalies. *Neuroimage* 13:110-8.
20. Price CJ (2000) The anatomy of language: contributions from functional neuroimaging. *J Anat* 197:335-59.
21. Stowe L, Haverkort M, Zwarts F (2005) Rethinking the neurobiological basis of language. *Lingua* 115:997-1042.
22. Monti MM, Osherson DN, Martinez MJ, Parsons LM (2007) Functional neuroanatomy of deductive inference: a language-independent distributed network. *Neuroimage* 37:1005-16.
23. Rodd JM, Davis MH, Johnsrude IS (2005) The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. *Cereb Cortex* 15:1261-9.
24. Ullman MT (2001) A neurocognitive perspective on language: the declarative/procedural model. *Nat Rev Neurosci* 2:717-26.
25. Petacchi A, Laird AR, Fox PT, Bower JM (2005) Cerebellum and auditory function: an ALE meta-analysis of functional neuroimaging studies. *Hum Brain Mapp* 25:118-28.
26. Bookheimer S (2002) Functional MRI of language: new approaches to understanding the cortical organization of semantic processing. *Annu Rev Neurosci* 25:151-88.
27. Sahin NT, Pinker S, Halgren E (2006) Abstract grammatical processing of nouns and verbs in Broca's area: evidence from fMRI. *Cortex* 42:540-62.
28. Bornkessel I, Zysset S, Friederici AD, von Cramon DY, Schlesewsky M (2005) Who did what to whom? The neural basis of argument hierarchies during language comprehension. *Neuroimage* 26:221-33.
29. Ben-Shachar M, Hendler T, Kahn I, Ben-Bashat D, Grodzinsky Y (2003) The neural reality of syntactic transformations: evidence from functional magnetic resonance imaging. *Psychol Sci* 14:433-40.
30. Kuperberg G, Lakshmanan B, Caplan D, Holcomb P (2006) Making sense of discourse: an fMRI study of causal inferencing across sentences. *Neuroimage* 33:343-61.
31. Virtue S, Haberman J, Clancy Z, Parrish T, Jung Beeman M (2006) Neural activity of inferences during story comprehension. *Brain Res* 1084:104-14.
32. Grice P (1989) *Studies in the Way of Words* (Harvard Univ Press, Cambridge).
33. Chierchia G, McConnell-Ginet S (2000) *Meaning and Grammar: An Introduction to Semantics* (MIT Press, Cambridge) pp 501-528.
34. Ludlow P (2002) in *Logical Form, Language and Ontology: On Contemporary Developments in the Philosophy of Language and Linguistics*, ed Preyer G (Oxford Univ Press, Oxford), pp 132 - 168.
35. Fodor J (2008) *LOT 2: The Language of Thought Revisited* (Oxford Univ Press, Oxford).
36. Preyer G, Peter G (2002) *Logical Form and Language* (Oxford Univ Press, Oxford).
37. Grodzinsky Y, Friederici AD (2006) Neuroimaging of syntax and syntactic processing. *Curr Opin Neurobiol* 16:240-6.
38. Udden J et al. (2008) The inferior frontal cortex in artificial syntax processing: an rTMS study. *Brain Res* 1224:69-78.
39. Christoff K et al. (2001) Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *Neuroimage* 14:1136-49.
40. Kroger JK et al. (2002) Recruitment of anterior dorsolateral prefrontal cortex in human reasoning: a parametric study of relational complexity. *Cereb Cortex* 12:477-85.
41. Ramnani N, Owen AM (2004) Anterior prefrontal cortex: insights into function from anatomy and neuroimaging. *Nat Rev Neurosci* 5:184-94.
42. van den Heuvel OA et al. (2003) Frontostriatal system in planning complexity: a parametric functional magnetic resonance version of Tower of London task. *Neuroimage* 18:367-74.
43. De Pisapia N, Braver TS (2008) Preparation for integration: the role of anterior prefrontal cortex in working memory. *Neuroreport* 19:15-9.
44. Koechlin E, Corrado G, Pietrini P, Grafman J (2000) Dissociating the role of the medial and lateral anterior prefrontal cortex in human planning. *Proc Natl Acad Sci U S A* 97:7651-6.
45. Volz KG, Schubotz RI, von Cramon DY (2005) Variants of uncertainty in decision-making and their neural correlates. *Brain Res Bull* 67:403-12.
46. Goel V, Buchel C, Frith C, Dolan RJ (2000). Dissociation of mechanisms underlying syllogistic reasoning. *Neuroimage* 12, 504-14.
47. Reverberi C et al. (2007). Neural basis of generation of conclusions in elementary deduction. *Neuroimage* 38, 752-62.
48. Reverberi C, Rusconi P, Paulesu E, Cherubini P. (2009). Response demands and the recruitment of heuristic strategies in syllogistic reasoning. *Q J Exp Psychol* 62, 513-30.
49. Rypma B, Prabhakaran V, Desmond JE, Glover GH, Gabrieli JD (1999) Load-dependent roles of frontal brain regions in the maintenance of working memory. *Neuroimage* 9:216-26.
50. Tanaka S, Honda M, Sadato N (2005) Modality-specific cognitive function of medial and lateral human Brodmann area 6. *J Neurosci* 25:496-501.
51. Hanakawa T et al. (2002) The role of rostral Brodmann area 6 in mental-operation tasks: an integrative neuroimaging approach. *Cereb Cortex* 12:1157-70.
52. Poldrack RA et al. (1999) Functional specialization for semantic and phonological processing in the left inferior prefrontal cortex. *Neuroimage* 10:15-35.
53. Knauff M, Fangmeier T, Ruff CC, Johnson-Laird PN (2003) Reasoning, models, and images: behavioral measures and cortical activity. *J Cogn Neurosci* 15:559-73.
54. Bunge SA, Kahn I, Wallis JD, Miller EK, Wagner, AD (2003) Neural circuits subserving the retrieval and maintenance of abstract rules. *J Neurophysiol* 90:3419-28.
55. Pinel P, Dehaene S, Riviere D, LeBihan D (2001) Modulation of parietal activation by semantic distance in a number comparison task. *Neuroimage* 14:1013-26.
56. Montague R (1970) *Formal Philosophy Selected Papers of Richard Montague* (Yale Univ Press, New Haven).
57. Partee B, Hendricks HL (1997) in *Handbook of Logic and Language*, eds van Benthem J, ter Meulen A (MIT Press, Cambridge), pp 5-91.
58. Johnson-Laird PN, Byrne R (1991) *Deduction* (Erlbaum, Hillsdale).
59. Osherson D (1975) *Reasoning: Representation and Process, Logic and models of logical thinking* (Erlbaum, Hillsdale).
60. Braine MD, O'Brien DP (1991) A theory of if: A lexical entry, reasoning program, and pragmatic principles. *Psychol Rev* 98:182-203
61. Rips L (1994) *The Psychology of Proof* (MIT Press, Cambridge).
62. Smith SM et al. (2004) Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23:S208-19.