Trait-level paranoia shapes inter-subject synchrony in brain activity during an ambiguous social narrative

Emily S. Finn*[1], Philip R. Corlett[2], Gang Chen[3], Peter A. Bandettini[1], R. Todd Constable[4]

[1]Section on Functional Imaging Methods, Laboratory of Brain and Cognition, National Institute of Mental Health; Bethesda, Md.

[2]Department of Psychiatry, Yale School of Medicine; New Haven, Conn.

[3]Scientific and Statistical Computing Core, National Institute of Mental Health; Bethesda, Md.

[4]Department of Radiology and Biomedical Imaging, Yale School of Medicine, New Haven, Conn.

*Corresponding author: emily.finn@nih.gov

ABSTRACT

Individuals often interpret the same event in different ways. How do personality traits modulate patterns of brain activity evoked by a complex stimulus? Participants listened to a narrative during functional MRI describing a deliberately ambiguous social scenario, designed such that some individuals would find it highly suspicious, while others less so. Using inter-subject correlation analysis, we identified several brain areas that were differentially synchronized during listening between participants with high- and low trait-level paranoia, including theory-of-mind regions. Event-related analysis indicated that while superior temporal cortex responded reliably to mentalizing events in all participants, anterior temporal and medial prefrontal cortex responded to such events only in high-paranoia individuals. Analyzing participants' speech as they freely recalled the narrative revealed semantic and syntactic features that also scaled with paranoia. Results indicate that trait paranoia acts as an intrinsic 'prime', producing different neural and behavioral responses to the same stimulus across individuals.

That two individuals may see the same event in different ways is a truism of human nature. Examples are found at many scales, from low-level perceptual judgments to interpretations of complex, extended scenarios. This latter phenomenon is known as the "Rashomon effect"[1] after a 1950 Japanese film in which four eyewitnesses give contradictory accounts of a crime and its aftermath, raising the point that for multifaceted, emotionally charged events, there may be no single version of the truth.

What accounts for these individual differences in interpretation? Assuming everyone has access to the same perceptual information, personality traits may bias different individuals toward one interpretation or another. Paranoia is one such trait, in that individuals with strong paranoid tendencies may be more likely to assign a nefarious interpretation to otherwise neutral events. While paranoia in its extreme is a hallmark symptom of schizophrenia and other psychoses, trait-level paranoia, like many experiences and behaviors associated with psychiatric illness, exists as a continuum rather than a dichotomy[2,3]. On a behavioral level, up to 30 percent of people report experiencing certain types of paranoid thoughts (e.g., 'I need to be on my guard against others') on a regular basis[4] and trait-level paranoia in the population follows an exponential, rather than bimodal, distribution[5].

Few neuroimaging studies have investigated paranoia as a continuum between normality and pathology; the majority simply contrast healthy controls and patients suffering from clinical delusions. However, a handful of reports from subclinical populations describe patterns of brain activity that scale parametrically with tendency toward paranoid or delusional ideation. For example, it has been reported that higher-paranoia individuals show less activity in the medial temporal lobe during memory retrieval and less activity in the cerebellum during sentence completion[6], less activity in temporal regions during social reflection[7] and auditory oddball detection[8], but higher activity in the insula and medial prefrontal cortex during self-referential processing[9] and differential patterns of activity in these regions as well as the amygdala while viewing emotional pictures[10].

Such highly controlled paradigms enable precise inferences about evoked brain activity, but potentially at the expense of real-world validity. For example, brain response to social threat is often assessed with decontextualized static photographs of unfamiliar faces presented rapidly in series (see [11] for a review). Compare this to threat detection in the real world, which involves perceiving and interacting with both familiar and unfamiliar faces in a rich, dynamic social

context. Paranoid thoughts that eventually reach clinical significance usually have a slow, insidious onset, involving complex interplay between a person's intrinsic tendencies and his or her experiences in the world. In studying paranoia and other trait-level individual differences, then, is important to complement highly controlled paradigms with more naturalistic stimuli.

Narrative is an attractive paradigm for several reasons. First, narrative is an ecologically valid way to study belief formation in action. Theories of fiction posit that readers model narratives in a Bayesian framework in much the same way as real-world information[12], and story comprehension and theory-of-mind processes share overlapping neural resources[13]. Second, a standardized narrative stimulus provides identical input, so any variation in interpretation reflects individuals' intrinsic biases in how they assign salience, learn and form beliefs. Third, from a neuroimaging perspective, narrative listening is a continuous, engaging task that involves much of the brain[14] and yields data lending itself to innovative, data-driven analyses such as inter-subject correlation[15,16].

Previous work has shown that experimenters can manipulate patterns of brain activity during naturalistic stimuli by explicitly instructing participants to focus on different aspects of the stimulus. For example, Cooper et al. showed that activity patterns in temporal and frontal regions varied according to whether listeners were told to pay attention to action-, space- or time-related features of short stories[17]. Lahnakoski et al. showed participants the same movie twice, asking them to adopt different perspectives each time, and found differences in neural synchrony depending on which perspective had been taken; results were primarily in parahippocampal, parietal and occipital regions[18]. Most recently, Yeshurun et al. presented participants with a highly ambiguous story with at least two plausible—but very different—interpretations, and used explicit primes to strongly bias each participant toward one interpretation or the other. Responses in higher-order brain areas, including default mode, were more similar among participants who had received the same prime, indicating that shared beliefs have a powerful effect on how individuals perceive an identical stimulus[19]. However, while informative, these studies have all relied on an explicit prime or instruction; they cannot explain why individuals often spontaneously arrive at different interpretations of the same stimulus.

In this work, we use participants' intrinsic personality traits as an implicit prime, relating individual differences in trait-level paranoia to brain activity during a naturalistic task in which participants are faced with complex, ambiguous social circumstances. Using a novel narrative

stimulus, we show that while much of the brain is synchronized across all participants during story listening, stratifying participants based on trait-level paranoia reveals an additional set of regions showing stereotyped activity specifically among high-paranoia individuals; many of these are regions involved in theory-of-mind and mentalizing. An encoding model of the task suggests that these regions, including the temporal pole and medial prefrontal cortex, are particularly sensitive to "mentalizing events" when the main character is experiencing an ambiguous social interaction or explicitly reasoning about other characters' intentions. Finally, we measured participants' behavioral reactions to the narrative by analyzing their speech as they freely recalled the story, and identified semantic and syntactic features that vary parametrically with trait-level paranoia. Together, results indicate that trait-level paranoia modulates both neural and behavioral responses to a single stimulus across individuals.

RESULTS

Behavioral data and task performance

We created an original, fictional narrative to serve as the stimulus for this study. The narrative described a main character faced with a complex social scenario that was deliberately ambiguous with respect to the intentions of certain characters; it was designed such that different individuals would interpret the events as more nefarious and others as less so. A synopsis of the story is given in Box 1.

Twenty-two healthy participants listened to a pre-recorded audio version of the narrative (total duration = 21:50 min:sec, divided into three parts) during fMRI scanning. Following each of the three parts, participants answered three challenging multiple-choice comprehension questions to ensure they had been paying attention. Performance was very accurate (15 of the 22 subjects answered 9/9 [100%] questions correctly, while five answered 8/9 [89%] correctly and two answered 7/9 [78%] correctly). Self-report data indicated that subjects generally found the narrative engaging and easy to pay attention to (engagement rating on a scale of 1 to 5: mean = 3.8, s.d. = 0.96, median = 4, median absolute deviation [m.a.d.] = 0.72; attention rating: mean = 4.1, s.d. = 0.87, median = 4, m.a.d. = 0.66).

During a separate behavioral visit one week prior to the scan, participants completed several self-report questionnaires and behavioral tasks to assess personality traits and cognitive abilities (see Fig. 1a for a schematic of the experimental protocol). Our primary measure of interest was subscale A from the Green et al. Paranoid Thoughts Scale[20] (GPTS-A). We administered this scale on a different day, and placed it amongst other tasks unrelated to paranoia, to minimize any priming effects or demand characteristics that might influence participants' eventual reactions to the narrative. Possible scores on the GPTS-A range from 16 to 80; higher scores are generally observed only in clinical populations[20]. In our healthy sample, we observed a right-skewed distribution that nonetheless had some variance (range = 16-40, mean = 20.6, s.d. = 6.3; median = 18.5, m.a.d. = 4.0; see Fig. 1b for a histogram of the distribution). This is consistent with observations from much larger sample sizes that paranoia follows an exponential, rather than normal, distribution in the healthy population[4,5,20].

Story listening evokes widespread neural synchrony

Our primary approach for analyzing the fMRI data was inter-subject correlation (ISC), which is a model-free way to identify brain regions responding reliably to a naturalistic stimulus across subjects[15,16]. In this approach, the timecourse from each voxel in one subject's brain across the duration of the stimulus is correlated with the timecourse of the same voxel in a second subject's brain. Voxels that show high correlations in their timecourses across subjects are considered to have a stereotyped functional role in processing the stimulus. The advantage of this approach is that it does not require the investigator to have an *a priori* model of the task, nor to assume any fixed hemodynamic response function.

In a first-pass analysis, we calculated ISC at each voxel across the whole sample of n = 22 participants, using a recently developed statistical approach that relies on a linear mixed-effects model with crossed random effects to appropriately account for the correlation structure embedded in the data[21]. Results are shown in Fig. 2. As expected given the audio-linguistic nature of the stimulus, ISC was highest in primary auditory cortex and language regions along the superior temporal lobe, but we also observed widespread ISC in other parts of association cortex, including frontal, parietal, midline and temporal areas, as well as the posterior cerebellum. These results replicate previous reports that complex naturalistic stimuli induce

stereotyped responses across participants in not only the relevant primary cortex, but also higher-order brain regions[14,15,22].

Also as expected, ISC was generally lower or absent in primary motor and somatosensory cortex, although we did observe significant ISC in parts of primary visual cortex, despite the fact that there was no timecourse of visual input during the story. (To encourage imagery and engagement, we had participants fixate on a static photograph that was thematically relevant to the story during listening, so the observed ISC in visual cortex may reflect similarities in the timecourse of internally generated visualization across participants.)

Trait-level paranoia modulates neural response to the narrative

Having established that story listening evokes widespread neural synchrony across the entire sample, we next sought to determine if there were brain regions whose degree of ISC was modulated by trait-level paranoia. Using a median split of GPTS-A scores, we stratified our sample into a low-paranoia (GPTS-A $\leq$ 18, n = 11) and high-paranoia (GPTS-A $\geq$ 19, n = 11) group (Fig. 1b). We then used the same linear mixed-effects model described above formulated as a two-group contrast to reveal areas that are differentially synchronized across paranoia groups.

We were primarily interested in two contrasts. First, which voxels show greater ISC among pairs of high-paranoia participants versus low-paranoia participants, or vice versa? This contrast reveals regions that have a stereotyped response in one group but not the other, suggesting a more prominent role for that region in processing the stimulus. Second, which voxels show greater ISC among pairs of participants within the same paranoia group (i.e., high-high and low-low) than across groups (high-low)? This contrast reveals regions that have equally stereotyped responses within each group, but whose response timecourses are qualitatively different across groups, suggesting that such regions may be responding to different events within the stimulus depending on paranoia level. These contrasts are schematized in Fig. 1c.

Results from both contrasts are shown in Fig. 3. In the first contrast, several regions emerged as being more synchronized in the high-paranoia group relative to the low-paranoia group. Significant clusters were found in the left temporal pole (Talairach coordinates for center of mass: [+46.7, -10.0, -26.2]), left precuneus ([+10.8, +71.0, +35.9]), and two regions of the right medial prefrontal cortex (mPFC; one anterior [-8.1, -46.9, +16.3] and one dorsal [+2.9, -

14.8, +45.1]). Searches for these coordinates on Neurosynth, an automatic fMRI results synthesizer for mapping between neural and cognitive states[23], indicated that for the left temporal pole and right anterior mPFC clusters, top meta-analysis terms included “mentalizing”, “mental states”, “intentions”, and “theory mind”. There were no regions showing a statistically significant difference in the reverse direction (low-paranoia > high-paranoia).

In the second contrast, a set of regions also emerged as being more synchronized *within* paranoia groups than *across* paranoia groups. These were the left middle occipital gyrus (Talairach coordinates for center of mass: [+31.3, +86.1, +14.0], Neurosynth: “objects”, “scene”, “encoding”) and the right angular gyrus ([-44.8, +57.9, +37.9], Neurosynth: “beliefs”). There were no regions showing a statistically significant difference in the reverse direction (across groups > within groups), reinforcing the biological validity of stratifying based on trait-level paranoia.

Establishing specificity to paranoia: control analyses

We conducted several control analyses to rule out the possibility that the observed group differences were driven by a factor other than trait-level paranoia. (For all analyses in this section, we checked for both categorical and continuous relationships with paranoia; full results are reported in Table 1.)

For example, if the high-paranoia participants have better overall attentional and cognitive abilities, they might simply be paying closer attention to the story, inflating ISC values but not necessarily because of selective attention to ambiguous or suspicious details. However, there were no differences between high- and low-paranoia participants on any of the cognitive tasks we administered (verbal IQ, vocabulary, fluid intelligence or working memory), making it unlikely that observed differences are due to trait-level differences in attention or cognition. As for state-level attention during the story, there was no relationship between paranoia and number of comprehension questions answered correctly, total word count during the recall task, or self-report measures of engagement and attention. We also explored potential imaging-based confounds, and found that paranoia was not related to amount of head motion during the scan (as measured by mean framewise displacement), number of censored frames, or temporal signal-to-noise ratio (tSNR). In terms of demographics, paranoia groups did not differ in age or sex

breakdown. Thus we are reasonably confident that the observed effects are driven by true trait-level differences in paranoia between individuals.

Probing response to mentalizing events using an encoding model of the task

Results of the first contrast from the two-group ISC analysis indicated that certain brain regions showed a more stereotyped response in high-paranoia versus low-paranoia individuals. What features of the narrative were driving activity in these regions? In theory, ISC allows for reverse correlation, in which peaks of activation in a given region's timecourse are used to recover the stimulus events that evoked them[15]. In practice, this is often difficult. Especially with narrative stimuli, in which structure is built up over relatively long timescales[14], it is challenging to pinpoint exactly which event—word, phrase, sentence—triggered an increase in BOLD activity.

Rather than rely on reverse correlation, which is a purely data-driven *decoding* approach, we took an *encoding* approach: we created a model of the task based on events that we hypothesized would stimulate differing interpretations across individuals, and evaluated the degree to which certain regions of interest (ROIs) responded to such events, using a standard general linear model (GLM) analysis. Specifically, we labeled sentences in the story when the main character was experiencing an ambiguous (i.e., possibly suspicious) social interaction, and/or sentences when she was explicitly reasoning about the intentions of other characters. For brevity, we refer to these timepoints as "mentalizing events." In creating the regressor, all events were time-locked to the end of the last word of the labeled sentences, when participants are presumably evaluating information they just heard and integrating it into their situation model of the story.

We hypothesized that the two ROIs from the previous analysis known to be involved in theory-of-mind and mentalizing, specifically the left temporal pole and right medial PFC, would show higher evoked activity to mentalizing events in individuals with higher trait-level paranoia. We included two additional ROIs, the left posterior superior temporal sulcus and left Heschl's gyrus, as a positive and negative control, respectively. We selected the left posterior superior temporal sulcus as a positive control because of its well-established role in theory-of-mind and mentalizing processes, and the fact that it emerged as highly synchronized across all participants (cf. Fig. 2) but did not show a group difference (cf. Fig. 3); thus we hypothesized that this region

should respond to mentalizing events in all participants, regardless of paranoia score. Conversely, left Heschl's gyrus (primary auditory cortex) should only respond to low-level acoustic properties of the stimulus and not show preferential activation to mentalizing events in either group or the sample as a whole.

For each participant, we regressed the timecourse of each of these four ROIs against the mentalizing-events regressor and compared the resulting regression coefficients between groups (Fig. 4b). Compared to low-paranoia individuals, high-paranoia individuals showed stronger responses in both the left temporal pole (two-sample $t(20) = 2.71$, $p_{adj} = 0.014$) and right medial PFC ($t(20) = 3.36$, $p_{adj} = 0.007$). As hypothesized, responses in the left PSTS were strong across the whole sample (one-sample $t(21) = 8.73$, $p < 0.0001$), but there was no significant difference between groups in this ROI ($t(20) = 0.67$, $p_{adj} = 0.34$). Also as hypothesized, the sample as a whole did not show a significant response to these events in primary auditory cortex (one-sample $t(21) = 0.44$, $p = 0.66$), and there was no group difference ($t(20) = 0.47$, $p_{adj} = 0.34$).

To confirm that these results hold if paranoia is treated as a continuous variable, we also calculated the rank correlation between paranoia score and beta coefficient for all four ROIs (Fig. 4c). As expected, response to suspicious events was significantly related to trait-level paranoia in the left temporal pole ($r_s = 0.57$, $p = 0.005$) and right medial PFC ($r_s = 0.64$, $p = 0.001$), but not in the left posterior STS ($r_s = -0.04$, $p = 0.86$) or left Heschl's gyrus ($r_s = 0.02$, $p = 0.95$).

As an additional control, to check that this effect was specific to mentalizing events and not just any sentence offset, we created an inverse regressor comprising all *non*-mentalizing events (i.e., by flipping the binary labels from the mentalizing-events regressor, such that all sentences were labeled *except* those containing an ambiguous social interaction or explicit mentalizing as described above). As expected, there were no differences between paranoia groups in any of the four ROIs in response to non-mentalizing sentences (Fig. 4d), and no continuous relationships between regression coefficient and paranoia (Fig. 4e). This suggests that trait-level paranoia is associated with differential sensitivity of the left temporal pole and right medial PFC to not just any type of information, but specifically to socially ambiguous information that presumably triggers theory-of-mind processes.

Paranoia modulates behavioral response to the narrative: Speech analysis

Having established that paranoia levels modulate individual participants' brain responses to an ambiguous narrative, we next sought to determine if this trait also modulates behavioral responses to the narrative. In other words, does trait-level (intrinsic) paranoia bear upon state-related (stimulus-evoked) paranoia? If the observed differences in neural activity propagate up to conscious perception and interpretation of the stimulus, then participants' subjective experiences of the narrative should also bear a signature of trait-level paranoia.

Immediately following the scan, participants completed a post-narrative battery that consisted of free-speech prompts followed by multiple-choice items to characterize their beliefs and feelings about the story. For the first item, participants were asked to retell the story in as much detail as they could remember, and their speech was recorded. Participants were allowed to speak for as long as they wished on whatever aspects of the story they chose. Without guidance from the experimenter, participants recalled the story in rich detail, speaking an average of 1,081 words (range = 399-3,185, s.d. = 610).

Audio recordings of participants' speech were transcribed and submitted to the language analysis software Linguistic Inquiry and Word Count (LIWC; [24]. The output of LIWC is one vector per participant describing the percentage of speech falling into various semantic and syntactic categories. Example semantic categories are positive emotion ('love', 'nice'), money ('cash', 'owe'), and body ('hands', 'face'), while syntactic categories correspond to parts of speech such as pronouns, adjectives and prepositions; there are 67 categories in total.

Using partial least-squares regression, we searched for relationships between speech features and trait-level paranoia. A first-pass analysis revealed that more than 72 percent of the variance in paranoia score could be accounted for by the first component of speech features; the loadings of semantic and syntactic categories for this component are visualized in Fig. 5a. The feature with the highest positive loading—indicating a positive relationship with paranoia—was affiliation, a category including words describing social and familial relationships (e.g., 'ally', 'friend', 'social'). Also associated with high trait-level paranoia were frequent use of adjectives as well as anxiety- and risk-related words (e.g., 'bad', 'crisis'); drives, a meta-category that includes words concerning affiliation, achievement, power, reward and risk; and health-related words (e.g., 'clinic', 'fever', 'infected'; recall that the story featured a doctor treating patients in a remote village; cf. Box 1). Features with strongly negative loadings—indicating an inverse relationship with paranoia—included male references (e.g., 'him', 'his', 'man', 'father'); anger-

related words (‘yell’, ‘annoyed’); function words (‘it’, ‘from’, ‘so’, ‘with’); and conjunctions (‘and’, ‘but’, ‘until’). See Fig. 5b for specific examples for selected categories from participants’ speech transcripts.

After the free-speech prompts, participants answered a series of multiple-choice questions. First, they were asked to rate the degree to which they were experiencing various emotions (suspicion, paranoia, sadness, happiness, confusion, anxiety, etc; 16 in total) on a scale from 1 to 5. Most of ratings skewed low—for example, the highest paranoia rating was 3, and only six subjects rated their paranoia level higher than 1. Interestingly, there was no significant correlation between trait-level paranoia score and self-reported paranoia (r_s = -0.02, p_{uncorr} = 0.91) or suspicion (r_s = 0.11, p_{uncorr} = 0.62) following the story. Neither were any of the other emotion ratings significantly correlated with trait-level paranoia (all uncorrected $p > 0.12$; see Fig. 5c).

Next, participants were asked to rate the three central characters on six personality dimensions (trustworthy, impulsive, considerate, intelligent, likeable, naïve). None of these personality ratings were significantly correlated with trait-level paranoia (all uncorrected $p > 0.09$; data not shown).

Finally, participants were asked to rate the likelihood of each of six scenarios. There was a trend such that individuals with higher trait-level paranoia were more likely to disagree with one of these scenarios (“Juan and the other villagers had not known anything about the disease before Carmen arrived”; r_s = -0.50, p_{uncorr} = 0.02), but this did not survive correction for multiple comparisons. None of the other scenario likelihood ratings significantly correlated with trait-level paranoia (all uncorrected $p > 0.19$; see Fig. 5d).

Overall, then, participants’ free speech was much more sensitive to individual differences in trait-level paranoia than their answers on the multiple-choice questionnaire. Self-report is a coarse measure that may suffer from response bias; behavior provides a richer feature set that allows for the discovery of more subtle associations. In studying nuanced individual differences, then, these results highlight the desirability of capturing behavior in naturalistic ways.

DISCUSSION

Here, we have shown that trait-level paranoia acts as a lens through which individuals perceive an ambiguous stimulus, yielding a spectrum of responses across participants at both the neural and behavioral levels. While our in-scanner narrative listening task evoked widespread inter-subject correlations in brain activity across all participants (Fig. 2), stratifying participants according to paranoia revealed differential inter-subject synchrony in various brain regions. Specifically, high-paranoia participants showed greater ISC in the left temporal pole, left precuneus, and two regions in the right medial prefrontal cortex (Fig. 3a). Participants more similar on trait-level paranoia—whether high or low—tended to have more similar activity patterns in left middle occipital gyrus and right angular gyrus (Fig. 3b). A follow-up event related analysis showed that the left temporal pole and right medial prefrontal cortex selectively responded to mentalizing events in high-paranoia participants, but not in low-paranoia participants (Fig. 4). Finally, analyzing participants' speech as they freely recalled the narrative revealed semantic and syntactic features that also scaled with trait-level paranoia (Fig. 5).

The fundamental advance of the present work is to show that an intrinsic personality trait can act as an "implicit prime" that colors how individuals perceive and interpret complex events. Previous work using naturalistic tasks has shown that brain activity and behavioral responses are sensitive to experimenter instructions, i.e., an explicit prime[18,19], or to the nature of the stimulus itself, i.e., whether it is more or less compelling or entertaining[25-27]. The present study extends these results in an important new direction, suggesting that in addition to such explicit modulations, there is substantial implicit variation in neural response to a naturalistic stimulus that stems from trait-level individual differences.

Our results have implications for the neural basis of both trait- and state-related paranoia. The relative hyperactivity of theory-of-mind regions in high-paranoia individuals fits with the conception of paranoia as "over-mentalizing", or the tendency to excessively attribute (malevolent) intentions to other people's actions[28]. In high-paranoia individuals, the story's ambiguous or mentalizing events triggered a response across more of the brain, activating not only the posterior superior temporal cortex as in low-paranoia individuals, but also an extended set of regions including the temporal pole and medial prefrontal cortex. Both of these regions are sometimes, but not always, reported in theory-of-mind tasks broadly construed; individual differences may at least partially explain the inconsistencies in the literature[29]. While their precise roles in such tasks remain unclear, it has been suggested that the medial PFC is

specifically involved in meta-cognitive processes of reflecting on feelings and intentions[30], and that the temporal pole is responsible for binding complex stimuli to emotional responses[31]. It is conceivable that both of these processes are more active in individuals with higher trait-level paranoia.

While the present study included only healthy controls with subclinical paranoia, it may provide a useful starting point for the study of paranoid or persecutory delusions in schizophrenia and related illnesses. Delusions with a persecutory theme account for roughly 70-80 percent of all delusions. This high prevalence is stable across time[32] and geo-cultural factors[33-36], suggesting a strong biological component. Persecutory delusions are also the type most strongly associated with anger and most likely to be acted upon, especially in a violent manner[37]. Thus, understanding the neurobiological basis of paranoid delusions is a critical problem in psychiatry.

But because delusions typically have a slow, insidious onset, it is nearly impossible to retrospectively recover specific triggering events in individual patients. A related challenge is that while thematically similar, each patient's delusion is unique in its details. Thus it is difficult to devise material that will evoke comparable responses across patients. One solution is to craft a model context using a stimulus that is ambiguous yet controlled—i.e., identical across participants, permitting meaningful comparisons of time-locked evoked activity— such as the one used in this work. Future work should study patient populations using paradigms such as this one, as they may shed light on mechanisms of delusion formation and/or provide eventual diagnostic or prognostic value.

While there is little work investigating how brain activity is altered during a naturalistic stimulus in psychiatric populations, a handful of studies have used such paradigms in autism, finding that ISC is lower among autistic individuals than typically developing controls while watching movies of social interactions[38-40]. Notably, the degree of asynchrony scales with autism-like phenotype severity in both the patient and control groups[39]. It is interesting to juxtapose these reports with the present results, in which individuals with a stronger paranoia phenotype were *more* synchronized during exposure to socially relevant material; ultimately, this fits with the notion of autism and psychosis as opposite ends of the same spectrum, involving hypo- and hyper-mentalization, respectively[41,42]. Eventually, it may be possible to combine a naturalistic stimulus with an ISC-based analysis that cuts across diagnostic labels to examine how the synchrony of neural response varies across both healthy and impaired populations.

From a methodological perspective, much of the research using fMRI to study individual differences has shifted focus in recent years from activation measured in task-based conditions to functional connectivity measured predominantly at rest[43-47]. Both suffer from limitations: traditional tasks are tightly controlled paradigms that often lack ecological validity; resting-state scans, on the other hand, are entirely unconstrained, making it difficult to separate signal from noise. Furthermore, very few resting-state scans use behavioral monitoring, so while they may detect individual or group differences, it is usually impossible to recover the nature of the mental events that give rise to such differences. Naturalistic tasks may be a happy medium for studying both group-level functional brain organization as well as individual differences[48,49]. We and others have argued that such tasks could serve as a "stress test" to draw out individual variation in brain and behaviors of interest[50-53], enhancing the signal-to-noise ratio in the search for neuroimaging-based biomarkers and permitting more precise inferences about the sources of individual differences in neural activity.

Acknowledgements

The authors thank Paul Taylor for advice on preprocessing pipelines for naturalistic task data, and Javier Gonzalez-Castillo, Daniel Handwerker, David Jangraw, Peter Molfese, Monica Rosenberg, and Tamara Vanderwal for helpful discussion.

Author contributions

E.S.F, P.R.C and R.T.C conceived the study. E.S.F. developed experimental materials and performed data collection. E.S.F. and G.C. analyzed the data. E.S.F., P.R.C., P.A.B. and R.T.C interpreted results. E.S.F. wrote the manuscript with comments from all other authors.

Competing financial interests

The authors have no competing financial interests to declare.

Carmen is a young American doctor who journeys to the Amazon to work in a small village health clinic. There, she meets Juan, a village leader, and Alba, a young girl whom she befriends. Soon after arriving, Carmen sees a series of patients with a very serious—and seemingly highly contagious—fever. At times it seems that Juan and the villagers are fully open and forthcoming with Carmen, while at other times their behavior is harder to interpret and it seems as if they may be hiding something. Carmen begins to wonder if the villagers had known about the disease before she arrived, and if she had somehow been deliberately lured to the remote location. The story ends abruptly, when Carmen discovers that Alba herself is sick, and that unbeknownst to her, Alba is Juan's daughter. Carmen fears she may have already been infected, and wonders what to do next.

Box 1. Synopsis of fictional narrative used as a stimulus.

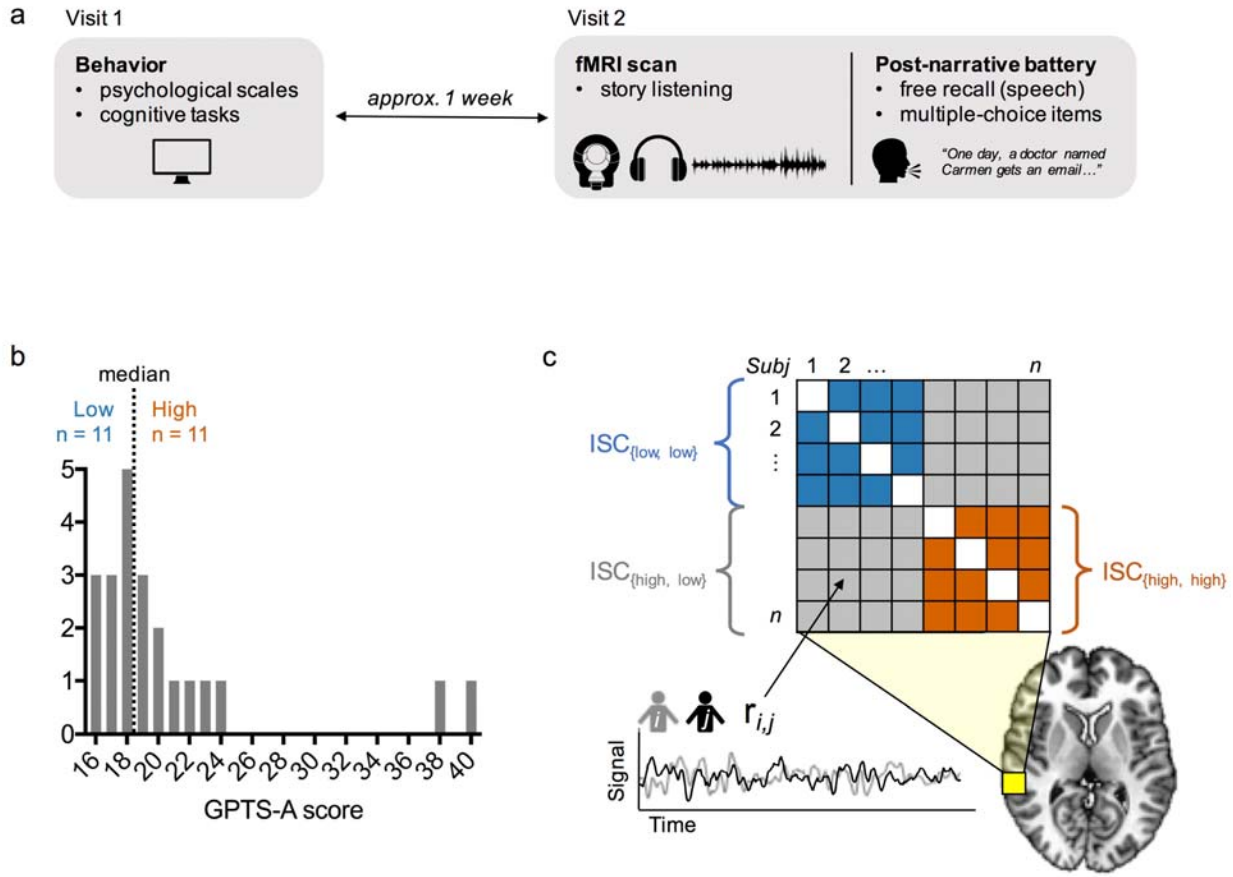

Figure 1. Experimental protocol, distribution of trait-level paranoia and inter-subject correlation analysis. a) Schematic of experimental protocol. Participants came to the laboratory for an initial behavioral visit, during which they completed several computerized cognitive tasks as well as self-report psychological scales, one of which was the Green et al. Paranoid Thoughts Scale(GPTS). To minimize demand characteristics and/or priming effects, the fMRI scan visit took place approximately one week later. During this visit, subjects listened to an ambiguous social narrative in the scanner and then completed an extensive post-narrative battery consisting of both free-speech prompts and multiple-choice items to characterize their recollections, feelings and beliefs toward the story. b) Distribution of scores on the GPTS-A subscale across n = 22 participants, and median split used to stratify participants into low (≤ 18, blue) and high (≥ 19, orange) trait-level paranoia. c) Schematic of inter-subject correlation (ISC) analysis. Following normalization to a standard template, the inter-subject correlation of activation timecourse during narrative listening was computed for each voxel (v, yellow square; enlarged relative to true voxel size for visualization purposes) for each pair of subjects (i,j), resulting in a matrix of pairwise correlation coefficients (r values). These values were then compared across paranoia groups using voxelwise linear mixed-effects models with crossed random effects to account for the non-independent structure of the correlation matrix[21].

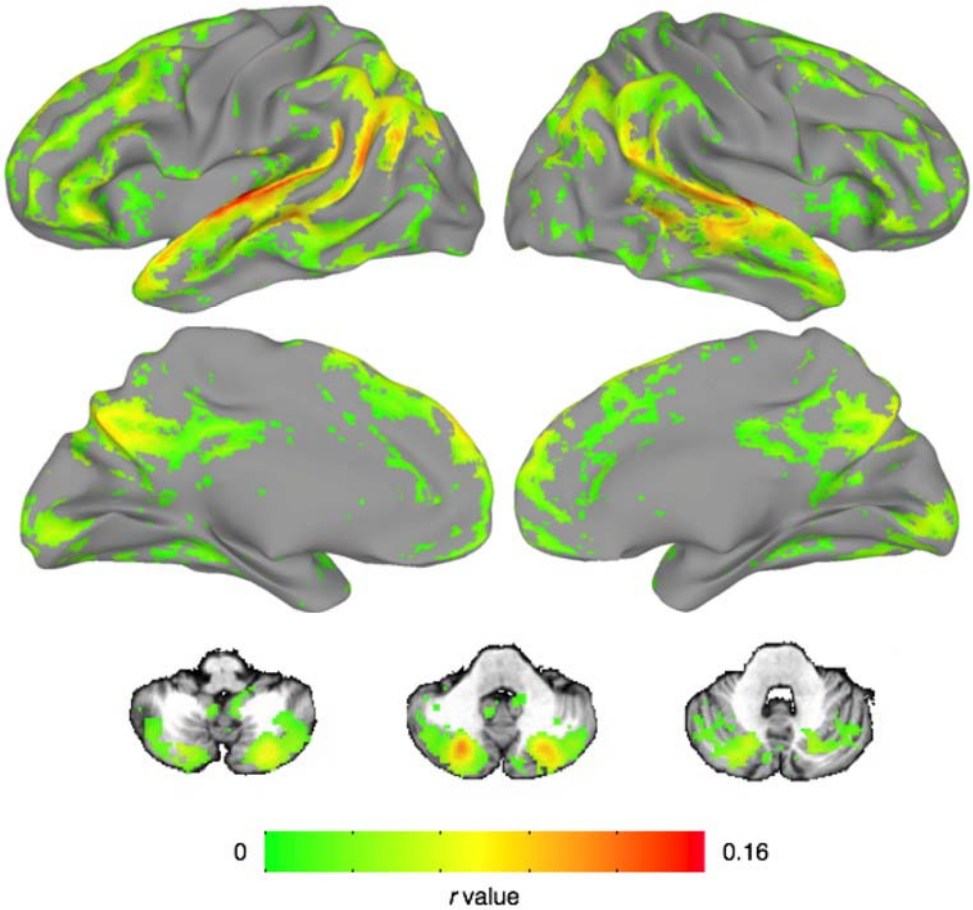

Figure 2. Narrative listening evokes widespread inter-subject correlation across the whole sample. Voxels showing significant inter-subject correlation (ISC) across the timecourse of narrative listening in all participants (n = 22). As expected, the highest ISC values were observed in auditory cortex, but several regions of association cortex in the temporal, parietal, frontal and cingulate lobes as well as the cerebellum also showed high synchrony. Also included are three representative axial slices from the cerebellum. Results are displayed at a voxelwise false-discovery rate (FDR) threshold of $q < 0.001$)

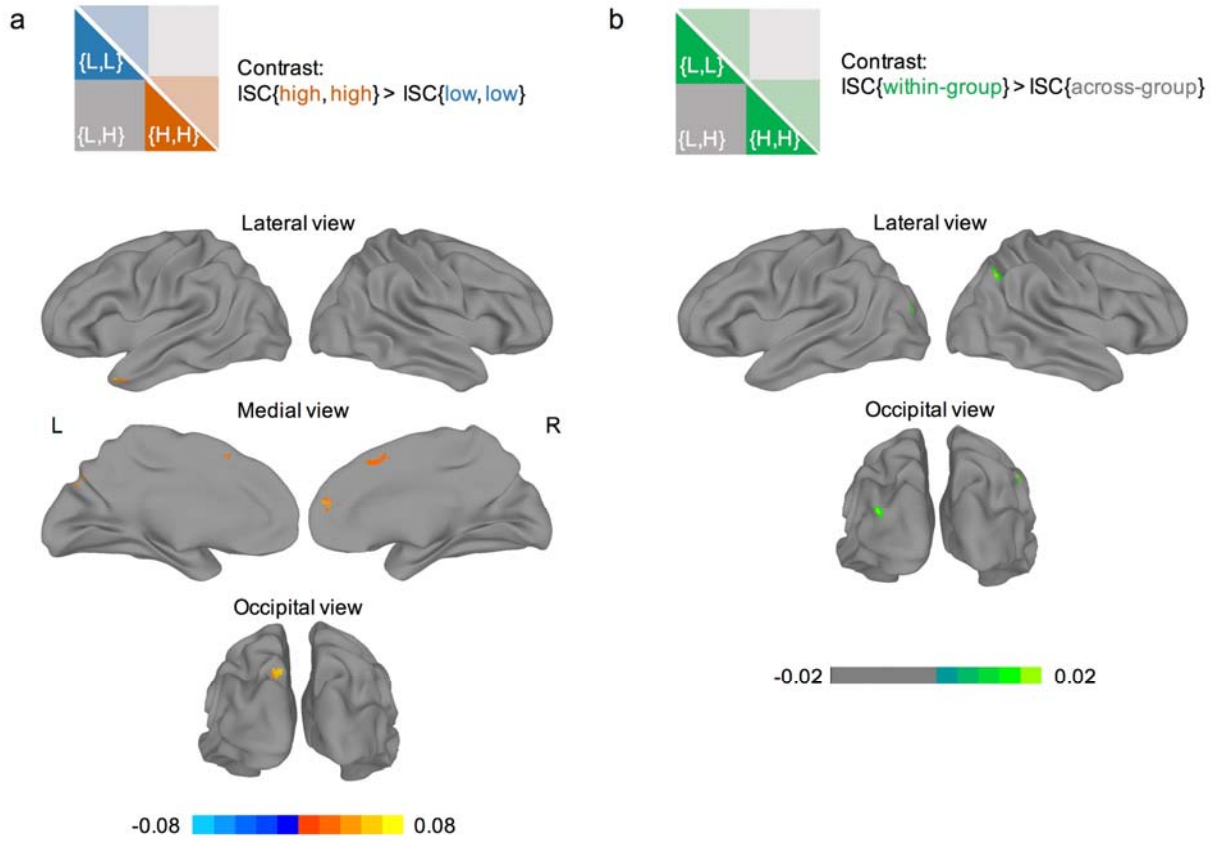

Figure 3. Trait-level paranoia modulates patterns of inter-subject correlation during narrative listening. a) Results from a whole-brain, voxelwise contrast revealing brain regions that are more synchronized between pairs of high-paranoia participants than pairs of low-paranoia participants (contrast schematized in top panel, cf. Fig. 1C). Significant clusters were detected in the left temporal pole, two regions in the right medial prefrontal cortex (one anterior and one dorsal and posterior), and the left precuneus. No clusters were detected in the opposite direction (low > high). b) Results from a whole-brain, voxelwise contrast revealing brain regions that are more synchronized within paranoia groups (i.e., high-high and low-low pairs) than across paranoia groups (i.e., high-low pairs; contrast schematized in top panel, cf. Fig 1C). Clusters were detected in the right angular gyrus and left lateral occipital cortex. For both contrasts, results are shown at an initial threshold of $p < 0.002$ with cluster correction corresponding to $p < 0.05$.

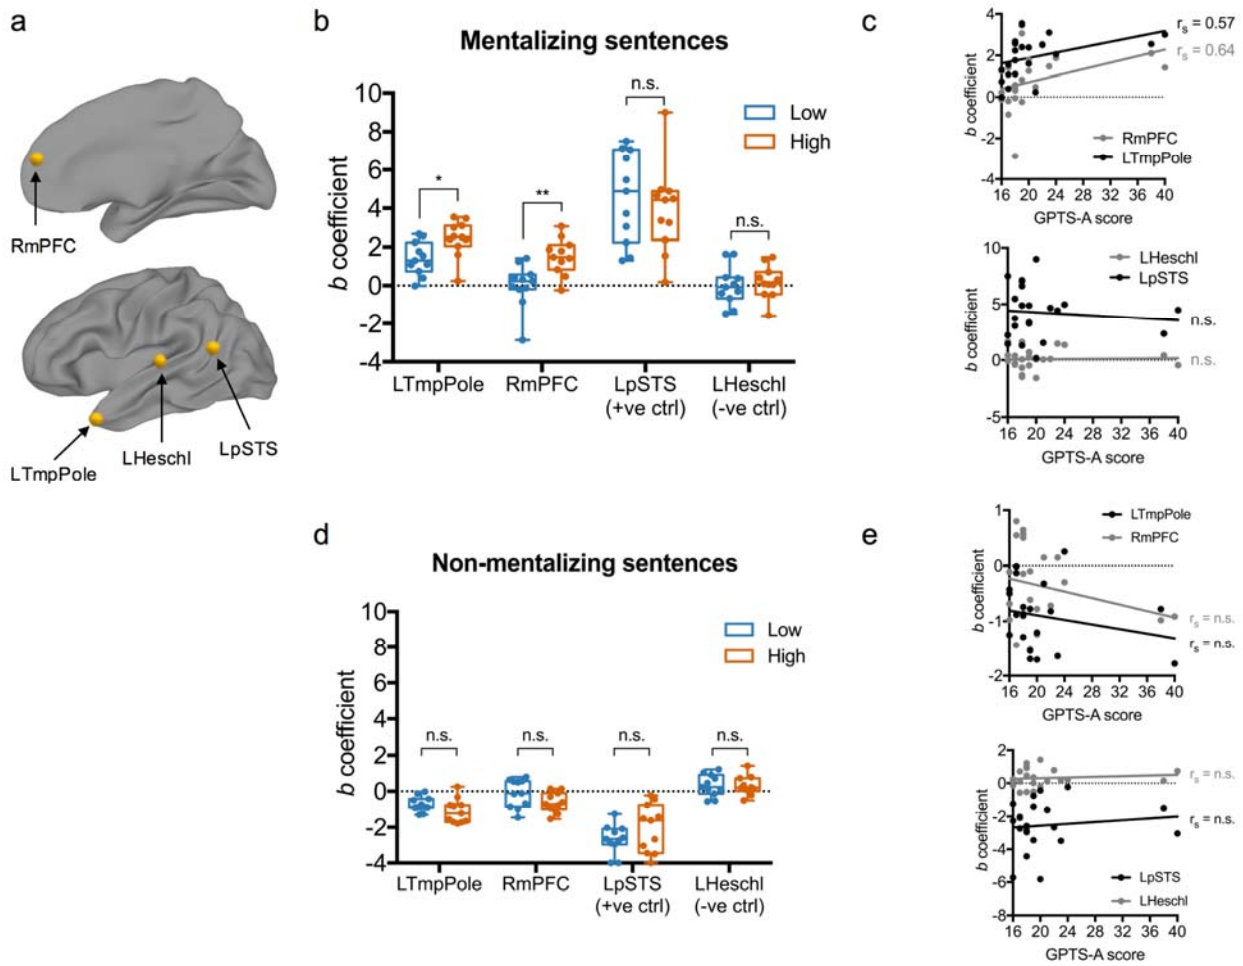

Figure 4. Response to mentalizing events is stronger in high- compared to low-paranoia individuals. a) Regions of interest (ROIs) for the event-related analysis. b) Comparison of beta coefficients for each ROI for the mentalizing-events regressor between paranoia groups (low, blue; high, orange). Each dot represents a subject. Boxes represent the median and $25^{th}/75^{th}$ percentiles, and whiskers represent the minimum and maximum. *p = 0.01; **p < 0.007; n.s., not significant (p-values adjusted for four comparisons). c) Beta coefficients for the mentalizing-events regressor plotted against paranoia score treated as a continuous variable (coefficients are the same as in (b)). Top panel: the two ROIs in which beta coefficient was hypothesized to scale with trait-level paranoia (left temporal pole and right medial PFC). Bottom panel: the two control ROIs (left posterior superior temporal sulcus and left Heschl's gyrus). Continuous relationships between these two variables were assessed with rank (Spearman) correlation, r_s. d) Comparison of beta coefficients for each ROI for the non-mentalizing-events regressor (the inverse of the mentalizing-events regressor shown in (b)). Each dot represents a subject. Boxes represent the median and $25^{th}/75^{th}$ percentiles, and whiskers represent the minimum and maximum. e) Beta coefficients for the non-mentalizing-events regressor plotted against paranoia score treated as a continuous variable (coefficients are the same as in (c)). Left and right panels as in (c).

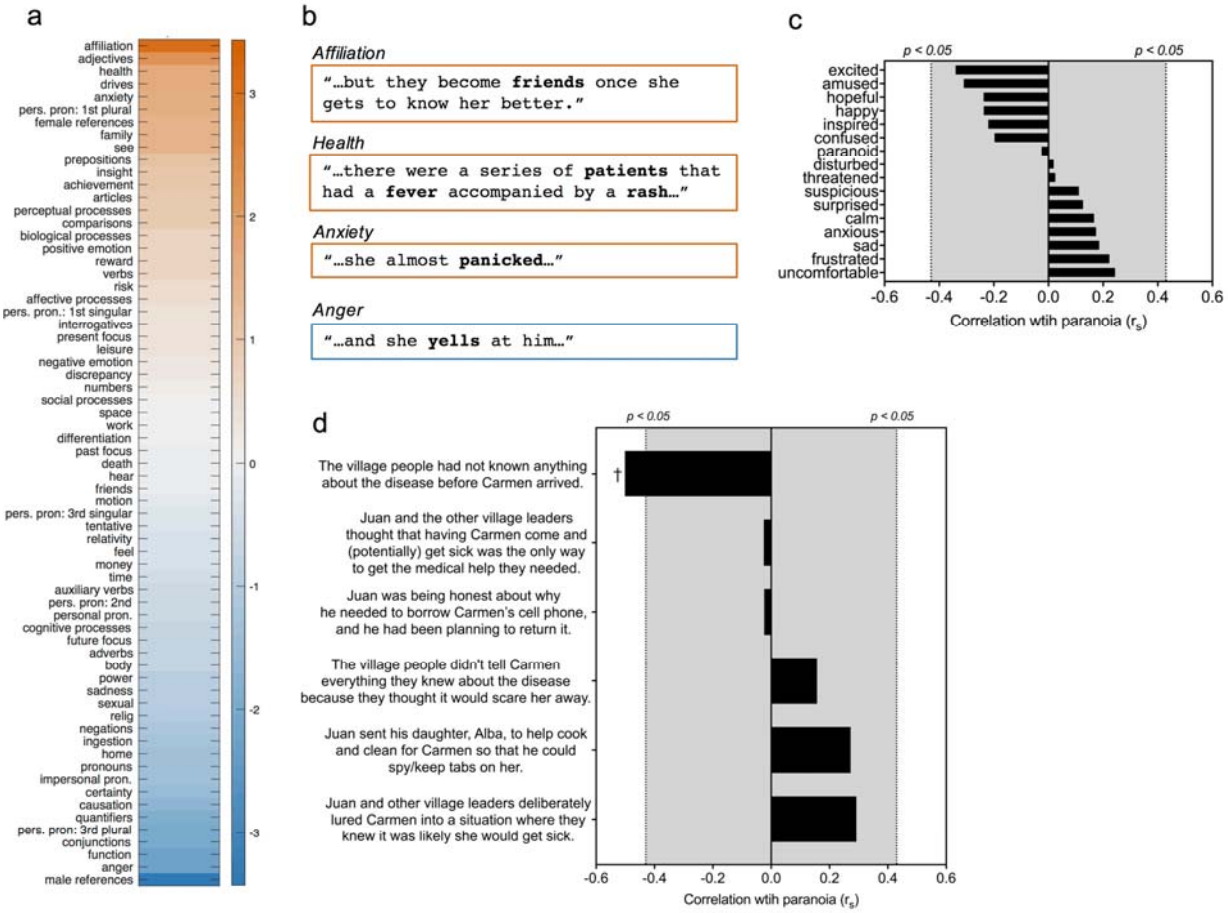

Figure 5. Speech analysis reveals a signature of trait-level paranoia in behavioral response to the narrative. a) Loadings of all semantic and syntactic categories for the first component from a partial least squares regression relating features of speech during narrative recall to trait-level paranoia score, sorted by strength and direction of association with paranoia (those positively related to paranoia at top in orange; those inversely related at bottom in blue). b) Example sentences from participant speech transcripts containing words falling into the three of the top positive categories (affiliation, health and anxiety) and one of the top negative categories (anger). c) Rank correlations between participants' trait-level paranoia and their self-report measures of 16 emotions following the narrative (self-report was based on a Likert scale from 1 to 5). Dotted lines represent approximate threshold for a significant correlation a $p < 0.05$ (uncorrected). Gray shaded area indicates non-significance. d) Rank correlations between participants' trait-level paranoia and their likelihood ratings of various potential scenarios following the narrative (based on a Likert scale from 1 to 5). Likelihood rating for one scenario (denoted with †) was significant at $p < 0.05$ but this did not survive correction for multiple comparisons.

		Categorical (low vs. high)		Continuous	
		t	*p*	Spearman *r*	*p*
Demographics	Age	0.81	0.43	-0.11	0.62
	Sex*	1.64	0.20	--	--
	Education (yrs)	-0.24	0.81	-0.15	0.49
Cognitive ability	Working memory: Letter n-back (precision)	-0.45	0.66	0.16	0.47
	Fluid intelligence: Raven's matrices (total correct of 9 items)	0.00	1.00	-0.03	0.89
	Vocabulary: WRAT Word Reading (total correct of 42 items)	-1.42	0.17	0.31	0.16
	Verbal IQ: Penn logical reasoning test (total correct of 8 items)	0.23	0.82	-0.01	0.96
	Words of 6+ letters (free recall)	-1.03	0.32	0.04	0.85
	Words per sentence (free recall)	0.31	0.76	-0.18	0.43
fMRI data quality	Head motion (mean FD; mm)	0.94	0.36	0.01	0.96
	No. frames censored	-0.70	0.49	-0.08	0.74
	Average tSNR	-1.12	0.28	0.23	0.30
Attention to stimulus	No. comprehension questions correct	-0.31	0.76	0.08	0.72
	Total word count, free recall	1.00	0.33	-0.26	0.24
	Self-reported attention	0.48	0.63	-0.02	0.95
	Self-reported engagement	0.89	0.39	-0.10	0.65

Table 1. Trait paranoia was unrelated to potential confounding variables. There were no significant differences between high- and low-paranoia participants in terms of demographics, cognitive abilities, fMRI data quality or attention to the stimulus. Categorical comparisons were carried out using Student's t-tests between the low and high paranoia groups as determined by median split (degrees of freedom for all t-tests = 20). Continuous comparisons were carried out using Spearman (rank) correlation between raw paranoia score and the variable of interest. All p-values are raw (uncorrected). *Measured with a chi-squared test. FD, framewise displacement; tSNR, temporal signal-to-noise ratio; WRAT, Wide Range Achievement Test.

METHODS

Participants

A total of 23 healthy volunteers participated in this study. Data from one participant was excluded due to excessive head motion and self-reported falling asleep during the last third of the narrative. Thus, the final data set used for analysis contained 22 participants (11 females; age range = 19-35 years, mean = 27, s.d. = 4.4). All participants were right-handed, native speakers of English, with no history of neurological disease or injury, and were not on psychoactive medication at the time of scanning. All participants provided written informed consent in accordance with the Institutional Review Board of Yale University. The experiment took place over two visits to the laboratory. Participants were paid $25 upon completion of the first visit (behavioral assessments) and $75 upon completion of the second visit (MRI scan); all participants completed both visits.

Stimulus

An original narrative was written to serve as the stimulus for this experiment. For a synopsis of the story, see Box 1. To mitigate confounds associated with education level or verbal IQ, we wrote the narrative text to be easy to comprehend, with a readability level of 78.1/100 and a grade 5.5 reading level as calculated by the Flesch-Kinkaid Formula.

Audio recording. A male native speaker of English read the story aloud and his speech was recorded using high-quality equipment at Haskins Laboratories (New Haven, Conn.). The speaker was instructed to read in a natural, conversational tone, but without excess emotion. The final length of the audio recording was 21:50.

Experimental protocol

Session 1: Behavior. Approximately one week prior to the scan visit, participants came to the laboratory to complete a battery of self-report and behavioral tasks. While our primary measure of interest was the Green et al. Paranoid Thoughts Scale (GPTS)[20], we also administered several other psychological scales and cognitive assessments, in part to help reduce any demand characteristics that would allow participants to intuit the purpose of the study. We chose the

GPTS because it provides a meaningful assessment of trait-level paranoia in clinical, but crucially, also subclinical and healthy populations. In a previous study, score on this scale best predicted feelings of persecution following immersion in a virtual-reality environment[54]. The full GPTS contains two subscales, A and B, which pertain to ideas of social reference and ideas of persecution, respectively. We focused on subscale A, as it produces a wider range of scores in subclinical populations[20].

Session 2: MRI scan. The full audio recording was divided into three segments of length 8:46, 7:32, and 5:32, respectively; each of these segments was delivered in a continuous functional run while participants were in the scanner. To ensure attention, after each run, subjects answered three challenging multiple-choice comprehension questions regarding the content of the part they had just heard, for a total of nine questions. Immediately upon exiting the scanner, participants completed a post-narrative questionnaire that consisted of open-ended prompts to elicit free speech, followed by multiple-choice items.

MRI data acquisition and preprocessing

Scans were performed on a 3T Siemens TimTrio system at the Yale Magnetic Resonance Research Center. After an initial localizing scan, a high-resolution 3D volume was collected using a magnetization prepared rapid gradient echo (MPRAGE) sequence (208 contiguous sagittal slices, slice thickness = 1 mm, matrix size 256 × 256, field of view = 256 mm, TR = 2400 ms, TE = 1.9 ms, flip angle = 8°). Functional images were acquired using a multiband T2*-sensitive gradient-recalled single shot echo-planar imaging pulse sequence (TR = 1000 ms, TE = 30 ms, voxel size = 2.0mm^3, flip angle = 60°, bandwidth = 1976 Hz/pixel, matrix size = 110 × 110, field of view = 220 mm × 220 mm, multiband factor = 4).

We acquired the following functional scans: 1) an initial eyes-open resting-state run (6:00/360 TRs in duration) during which subjects were instructed to relax and think of nothing in particular; 2) a movie-watching run using Inscapes [55] (7:00/420 TRs); 3) three narrative-listening runs corresponding to parts I, II and III of the story (21:50/1310 TRs in total); and 4) a post-narrative, eyes-open resting-state run (6:00/360 TRs) during which subjects were instructed to reflect on the story they had just heard. The present work focuses exclusively on data acquired during narrative listening. The narrative stimulus was delivered through MRI-compatible audio

headphones and a short “volume check” scan was conducted just prior to the first narrative run to ensure that participants could adequately hear the stimulus above the scanner noise. To promote engagement, during the three narrative runs, participants were asked to fixate on a static image of a jungle settlement and to actively imagine the story events as they unfolded.

Following conversion of the original DICOM images to NIFTI format, AFNI (Cox 1996) was used to preprocess MRI data. The functional time series went through the following preprocessing steps: despiking, head motion correction, affine alignment with anatomy, nonlinear alignment to a Talairach template (TT_N27), and smoothing with an isotropic FWHM of 5 mm. A ventricle mask was defined on the template and intersected with the subject’s cerebrospinal fluid (CSF) mask to make a subject-specific ventricle mask. Regressors were created from the first three principal components of the ventricles, and fast ANATICOR (Jo et al 2010) was implemented to provide local white matter regressors. Additionally, the subject's 6 motion time series, their derivatives and linear polynomial baselines for each of the functional runs were included as regressors. Censoring of time points was performed whenever the per-time motion (Euclidean norm of the motion derivatives) was ≥0.3 or when ≥10% of the brain voxels were outliers. Censored time points were set to zero rather than removed altogether (this is the conventional way to do censoring, but especially important for inter-subject correlation analyses, to preserve the temporal structure across participants). The final output of this preprocessing pipeline was a single functional run concatenating data from the three story runs (total duration = 21:50, 1310 TRs). All analyses were conducted in volume space and projected to the surface for visualization purposes.

We used mean framewise displacement (MFD), a per-participant summary metric, to assess the amount of head motion in the sample. MFD was overall relatively low (after censoring: mean = 0.075 mm, s.d. = 0.026, range = 0.035-0.14). Number of censored time points during the story was overall low but followed a right-skewed distribution (range = 0-135, median = 4, median absolute deviation = 25). All 22 participants in the final analysis retained at least 89 percent of the total time points in the story, so missing data was not a substantial concern. Still, we performed additional control analyses to ensure that number of censored timepoints and amount of head motion were not associated with paranoia score in any way that would confound interpretation of the results (see Table 1).

Inter-subject correlation

Following preprocessing, inter-subject correlation (ISC) during the story was computed across all possible pairs of subjects (*i*,*j*) using AFNI's 3dTcorrelate function, resulting in 231 (n*(n-1)/2, where n = 22) unique ISC maps, where the value at each voxel represents the Pearson's correlation between that voxel's timecourse in subject *i* and its timecourse in subject *j*.

To identify voxels demonstrating statistically significant ISC across all 231 subject pairs, we performed inference at the single-group level using a recently developed linear mixed-effects (LME) model with a crossed random-effects formulation to accurately account for the correlation structure embedded in the ISC data [21]. This approach has been characterized extensively, including a comparison to non-parametric approaches, and found to demonstrate proper control for false positives and good power attainment [21]. The resulting map was corrected for multiple comparisons and thresholded for visualization using a voxelwise false discovery rate threshold of $q < 0.001$ (Fig. 2).

In a second analysis, we stratified participants according to a median split of scores on the GPTS-A subscale. We used these groups to identify voxels that had higher ISC values within one paranoia group or the other, or higher ISC values within rather than across paranoia groups. To this end, we used a two-group formulation of the LME model. This model gives the following outputs: voxelwise population ISC values within group 1 (G_{11}); voxelwise population ISC values within group 2 (G_{22}); voxelwise population ISC values between the two groups that reflect the ISC effect between any pair of subjects with each belonging to different groups (G_{12}). These outputs can be compared to obtain several possible contrasts. Here, we were primarily interested in two of these contrasts: 1) G_{11} versus G_{22}; and 2) $G_{11}G_{22}$ versus G_{12}. The maps resulting from each of these contrasts were thresholded using an initial voxelwise threshold of $p < 0.002$ and controlled for family-wise error (FWE) using a cluster size threshold of 50 voxels, corresponding to a corrected *p*-value of 0.05. We opted for a particularly stringent initial p-threshold in light of recent concerns about false positives arising from performing cluster correction on maps with more lenient initial thresholds [56].

Event-related analysis

Creating the regressor. One of the authors (E.S.F.) manually labeled sentences containing either an ambiguous social interaction or an instance of the main character

mentalizing about other characters' intentions using a binary scoring system (1 = ambiguous social interaction or mentalizing present in sentence, 0 = neither ambiguous social interaction nor mentalizing present). Four additional, independent raters previously naïve to the narrative listened to the same version that was played to participants in the scanner. They were then given a written version of the narrative broken down by sentence and asked to label each sentence as described above. Sentences that were labeled by at least three of the five raters were included in the final set of events, which were timestamped based on the offset of the last word of the sentence. These timestamps were convolved with a canonical hemodynamic response function (HRF) to create the mentalizing-events regressor.

As a control analysis, we also created a regressor that was the inverse of the above regressor, by flipping the binary labels (0 or 1) for all sentences and convolving the corresponding sentences offset timestamps with the HRF; we refer to this as the non-mentalizing-events regressor.

ROI definition and GLM. For the left temporal pole and right medial PFC, ROIs were defined based on the cluster-corrected group-comparison map for the contrast $ISC_{high} > ISC_{low}$ (cf. Fig. 3a). For the left posterior superior temporal sulcus and left Heschl's gyrus, spherical ROIs were created by placing a sphere with radius 4 mm around a central coordinate. In the case of the left pSTS, this was the peak voxel in this region identified by the whole-sample ISC analysis (cf. Fig. 2; Talairach *xyz*, [+53, +55, +18]). In the case of left Heschl's gyrus, this was selected anatomically (Talairach *xyz*, [-41, -24, +9]; as in [57]).

Timecourses for each ROI were extracted from each participant's preprocessed functional data using AFNI's 3dmaskave function and regressed against both the mentalizing- and non-mentalizing-events regressors to obtain a regression coefficient for each participant for each ROI. These regression coefficients were then compared across groups using two-sample t-tests corrected for four multiple comparisons. In the case of the two control ROIs (left pSTS and Heschl's gyrus) for the mentalizing-events regressor, these coefficients were also pooled across both groups and submitted to a one-sample t-test to test for a significant deviation from zero.

Analysis of speech features

Audio recordings of participants' retelling of the story were professionally transcribed by a third-party company. We submitted the resulting transcripts to Linguistic Inquiry and Word

Count LIWC; [24, liwc.net], a software program that takes as input a given text and counts the percentage of words falling into different syntactic and semantic categories. Because LIWC was developed by researchers with interests in social, clinical, health, and cognitive psychology, the language categories were created to capture people's social and psychological states.

We restricted LIWC output to the 67 linguistic (syntactic and semantic) categories, excluding categories relating to metadata (e.g., percentage of words found in the LIWC dictionary), as well as categories irrelevant to spoken language (e.g., punctuation). Thus, our final LIWC output was a 22x67 matrix where each row corresponds to a participant and each column to a category.

These categories can be scaled very differently from one another. For example, words in the syntactic category "pronoun" accounted for between 10.3-20.5 percent of speech transcripts, while words in the semantic category "leisure" accounted for only 0-1.09 percent. To give approximately equal weight to all categories, we standardized each category (to have zero mean unit variance) across participants before performing partial least squares regression (PLSR) as described in the next section. This ensures that the resulting PLS components are not simply dominated by variance in categories that are represented heavily in all human speech.

Relating speech features to paranoia

To determine which speech features were most related to trait-level paranoia, we submitted the data to a partial least squares regression (PLSR) with the z-scored speech features as X (predictors) and GPTS-A score as Y (response), implemented in Matlab as *plsregress*. PLSR is a latent variable approach to modeling the covariance structure between two matrices, which seeks to find the direction in X space that explains the maximum variance in Y space. It is well suited to the current problem, because it can handle a predictor matrix with more variables than observations, as well as multi-collinearity among the predictors.

In a first-pass analysis, we ran a model with 10 components to determine the number of components needed to explain most of the variance in trait-level paranoia. Results of this analysis indicated that the first component was sufficient to explain 72.3 percent of the total variance in GPTS-A score, so we selected just this component for visualization and interpretation. Predictor loadings for this component are visualized in Fig. 5a.

REFERENCES

1 Davis, B., Anderson, R. & Walls, J. *Rashomon Effects: Kurosawa, Rashomon and Their Legacies*. (Taylor & Francis, 2015).

2 Johns, L. C. & Van Os, J. The Continuity of Psychotic Experiences in the General Population. *Clin. Psychol. Rev.* **21**, 1125-1141, (2001).

3 Insel, T. *et al.* Research Domain Criteria (Rdoc): Toward a New Classification Framework for Research on Mental Disorders. *Am. J. Psychiatry.* **167**, 748-751, (2010).

4 Freeman, D. *et al.* Psychological Investigation of the Structure of Paranoia in a Non-Clinical Population. *The British Journal of Psychiatry* **186**, 427-435, (2005).

5 Bebbington, P. E. *et al. The Structure of Paranoia in the General Population.* Vol. 202 (2013).

6 Whalley, H. C. *et al.* Correlations between Fmri Activation and Individual Psychotic Symptoms in Un-Medicated Subjects at High Genetic Risk of Schizophrenia. *BMC Psychiatry* **7**, 61, (2007).

7 Brent, B. K. *et al.* Subclinical Delusional Thinking Predicts Lateral Temporal Cortex Responses During Social Reflection. *Soc. Cogn. Affect. Neurosci.*, nss129, (2012).

8 Sumich, A., Castro, A. & Kumari, V. N100 and N200, but Not P300, Amplitudes Predict Paranoia/Suspiciousness in the General Population. *Pers. Individ. Dif.* **61**, 74-79, (2014).

9 Modinos, G., Renken, R., Ormel, J. & Aleman, A. Self-Reflection and the Psychosis-Prone Brain: An Fmri Study. *Neuropsychology* **25**, 295-305, (2011).

10 Modinos, G. *et al.* Multivariate Pattern Classification Reveals Differential Brain Activation During Emotional Processing in Individuals with Psychosis Proneness. *Neuroimage* **59**, 3033-3041, (2012).

11 Bishop, S. J. Neural Mechanisms Underlying Selective Attention to Threat. *Ann. N. Y. Acad. Sci.* **1129**, 141-152, (2008).

12 Kukkonen, K. Bayesian Narrative: Probability, Plot and the Shape of the Fictional World. *Anglia* **132**, 720-739, (2014).

13 Mar, R. A. The Neural Bases of Social Cognition and Story Comprehension. *Annu. Rev. Psychol.* **62**, 103-134, (2011).

14 Lerner, Y., Honey, C. J., Silbert, L. J. & Hasson, U. Topographic Mapping of a Hierarchy of Temporal Receptive Windows Using a Narrated Story. *J. Neurosci.* **31**, 2906-2915, (2011).

15 Hasson, U., Nir, Y., Levy, I., Fuhrmann, G. & Malach, R. Intersubject Synchronization of Cortical Activity During Natural Vision. *Science* **303**, 1634-1640, (2004).

16 Hasson, U., Malach, R. & Heeger, D. J. Reliability of Cortical Activity During Natural Stimulation. *Trends in Cognitive Sciences* **14**, 40-48, (2010).

17 Cooper, E. A., Hasson, U. & Small, S. L. Interpretation-Mediated Changes in Neural Activity During Language Comprehension. *Neuroimage* **55**, 1314-1323, (2011).

18 Lahnakoski, J. M. *et al.* Synchronous Brain Activity across Individuals Underlies Shared Psychological Perspectives. *Neuroimage* **100**, 316-324, (2014).

19 Yeshurun, Y. *et al.* Same Story, Different Story: The Neural Representation of Interpretive Frameworks. *Psychol. Sci.* **28**, 307-319, (2017).

20 Green, C. *et al.* Measuring Ideas of Persecution and Social Reference: The Green Et Al. Paranoid Thought Scales (Gpts). *Psychol. Med.* **38**, 101-111, (2008).

21 Chen, G., Taylor, P. A., Shin, Y.-W., Reynolds, R. C. & Cox, R. W. Untangling the Relatedness among Correlations, Part Ii: Inter-Subject Correlation Group Analysis through Linear Mixed-Effects Modeling. *NeuroImage* **147**, 825-840, (2017).
22 Yarkoni, T., Speer, N. K. & Zacks, J. M. Neural Substrates of Narrative Comprehension and Memory. *Neuroimage* **41**, 1408-1425, (2008).
23 Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C. & Wager, T. D. Large-Scale Automated Synthesis of Human Functional Neuroimaging Data. *Nature Methods* **8**, 665, (2011).
24 Linguistic Inquiry and Word Count: Liwc 2015 (Pennebaker Conglomerates, Austin, TX, 2015).
25 Schmälzle, R., Häcker, F. E. K., Honey, C. J. & Hasson, U. Engaged Listeners: Shared Neural Processing of Powerful Political Speeches. *Soc. Cogn. Affect. Neurosci.* **10**, 1137-1143, (2015).
26 Dmochowski, J. P. *et al.* Audience Preferences Are Predicted by Temporal Reliability of Neural Processing. *Nature communications* **5**, 4567, (2014).
27 Ki, J. J., Kelly, S. P. & Parra, L. C. Attention Strongly Modulates Reliability of Neural Responses to Naturalistic Narrative Stimuli. *J. Neurosci.* **36**, 3092-3101, (2016).
28 Brüne, M. "Theory of Mind" in Schizophrenia: A Review of the Literature. *Schizophr. Bull.* **31**, 21-42, (2005).
29 Carrington, S. J. & Bailey, A. J. Are There Theory of Mind Regions in the Brain? A Review of the Neuroimaging Literature. *Hum. Brain Mapp.* **30**, 2313-2335, (2009).
30 Amodio, D. M. & Frith, C. D. Meeting of Minds: The Medial Frontal Cortex and Social Cognition. *Nat. Rev. Neurosci.* **7**, 268-277, (2006).
31 Olson, I. R., Plotzker, A. & Ezzyat, Y. The Enigmatic Temporal Pole: A Review of Findings on Social and Emotional Processing. *Brain* **130**, 1718-1731, (2007).
32 Cannon, B. J. & Kramer, L. M. Delusion Content across the 20th Century in an American Psychiatric Hospital. *Int. J. Soc. Psychiatry* **58**, 323-327, (2012).
33 Tateyama, M., Asai, M., Hashimoto, M., Bartels, M. & Kasper, S. Transcultural Study of Schizophrenic Delusions. Tokyo Versus Vienna and Tubingen (Germany). *Psychopathology* **31**, 59-68, (1997).
34 Stompe, T. *et al.* Comparison of Delusions among Schizophrenics in Austria and in Pakistan. *Psychopathology* **32**, 225-234, (1998).
35 Gutiérrez-Lobos, K., Schmid-Siegel, B., Bankier, B. & Walter, H. Delusions in First-Admitted Patients: Gender, Themes and Diagnoses. *Psychopathology*, (2001).
36 Brakoulias, V. & Starcevic, V. A Cross-Sectional Survey of the Frequency and Characteristics of Delusions in Acute Psychiatric Wards. *Australasian Psychiatry* **16**, 87-91, (2008).
37 Coid, J. W., Ullrich, S., Kallis, C. & et al. The Relationship between Delusions and Violence: Findings from the East London First Episode Psychosis Study. *JAMA Psychiatry* **70**, 465-471, (2013).
38 Hasson, U. *et al.* Shared and Idiosyncratic Cortical Activation Patterns in Autism Revealed under Continuous Real-Life Viewing Conditions. *Autism Research* **2**, 220-231, (2009).
39 Salmi, J. *et al.* The Brains of High Functioning Autistic Individuals Do Not Synchronize with Those of Others. *NeuroImage: Clinical* **3**, 489-497, (2013).

40 Byrge, L., Dubois, J., Tyszka, J. M., Adolphs, R. & Kennedy, D. P. Idiosyncratic Brain Activation Patterns Are Associated with Poor Social Comprehension in Autism. *J. Neurosci.* **35**, 5837-5850, (2015).

41 Crespi, B. & Badcock, C. Psychosis and Autism as Diametrical Disorders of the Social Brain. *Behav. Brain Sci.* **31**, 241-261, (2008).

42 Ciaramidaro, A. *et al.* Schizophrenia and Autism as Contrasting Minds: Neural Evidence for the Hypo-Hyper-Intentionality Hypothesis. *Schizophr. Bull.* **41**, 171-179, (2015).

43 Kelly, C., Biswal, B. B., Craddock, R. C., Castellanos, F. X. & Milham, M. P. Characterizing Variation in the Functional Connectome: Promise and Pitfalls. *Trends in Cognitive Sciences* **16**, 181-188, (2012).

44 Castellanos, F. X., Di Martino, A., Craddock, R. C., Mehta, A. D. & Milham, M. P. Clinical Applications of the Functional Connectome. *Neuroimage* **80**, 527-540, (2013).

45 Finn, E. S. *et al.* Functional Connectome Fingerprinting: Identifying Individuals Using Patterns of Brain Connectivity. *Nat. Neurosci.* **18**, 1664-1671, (2015).

46 Smith, S. M. *et al.* A Positive-Negative Mode of Population Covariation Links Brain Connectivity, Demographics and Behavior. *Nat. Neurosci.* **18**, 1565-1567, (2015).

47 Rosenberg, M. D. *et al.* A Neuromarker of Sustained Attention from Whole-Brain Functional Connectivity. *Nat. Neurosci.* **advance online publication**, (2015).

48 Ren, Y., Nguyen, V. T., Guo, L. & Guo, C. C. Inter-Subject Functional Correlation Reveal a Hierarchical Organization of Extrinsic and Intrinsic Systems in the Brain. *Sci. Rep.* **7**, 10876, (2017).

49 Vanderwal, T. *et al.* Individual Differences in Functional Connectivity During Naturalistic Viewing Conditions. *bioRxiv*, (2016).

50 Finn, E. S. & Constable, R. T. Individual Variation in Functional Brain Connectivity: Implications for Personalized Approaches to Psychiatric Disease. *Dialogues Clin. Neurosci.* **18**, 277-287, (2016).

51 Dubois, J. & Adolphs, R. Building a Science of Individual Differences from Fmri. *Trends in Cognitive Sciences* **20**, 425-443, (2016).

52 Koyama, M. S. *et al.* Imaging the at-Risk Brain: Future Directions. *J. Int. Neuropsychol. Soc.* **22**, 164-179, (2016).

53 Finn, E. S. *et al.* Can Brain State Be Manipulated to Emphasize Individual Differences in Functional Connectivity? *Neuroimage*, (2017).

54 VALMAGGIA, L. R. *et al.* Virtual Reality and Paranoid Ideations in People with an 'at-Risk Mental State' for Psychosis. *The British Journal of Psychiatry* **191**, s63-s68, (2007).

55 Vanderwal, T., Kelly, C., Eilbott, J., Mayes, L. C. & Castellanos, F. X. Inscapes: A Movie Paradigm to Improve Compliance in Functional Magnetic Resonance Imaging. *Neuroimage* **122**, 222-232, (2015).

56 Eklund, A., Nichols, T. E. & Knutsson, H. Cluster Failure: Why Fmri Inferences for Spatial Extent Have Inflated False-Positive Rates. *Proc. Natl. Acad. Sci. USA*, 201602413, (2016).

57 Schönwiesner, M. *et al.* Heschl's Gyrus, Posterior Superior Temporal Gyrus, and Mid-Ventrolateral Prefrontal Cortex Have Different Roles in the Detection of Acoustic Changes. *J. Neurophysiol.* **97**, 2075-2082, (2007).

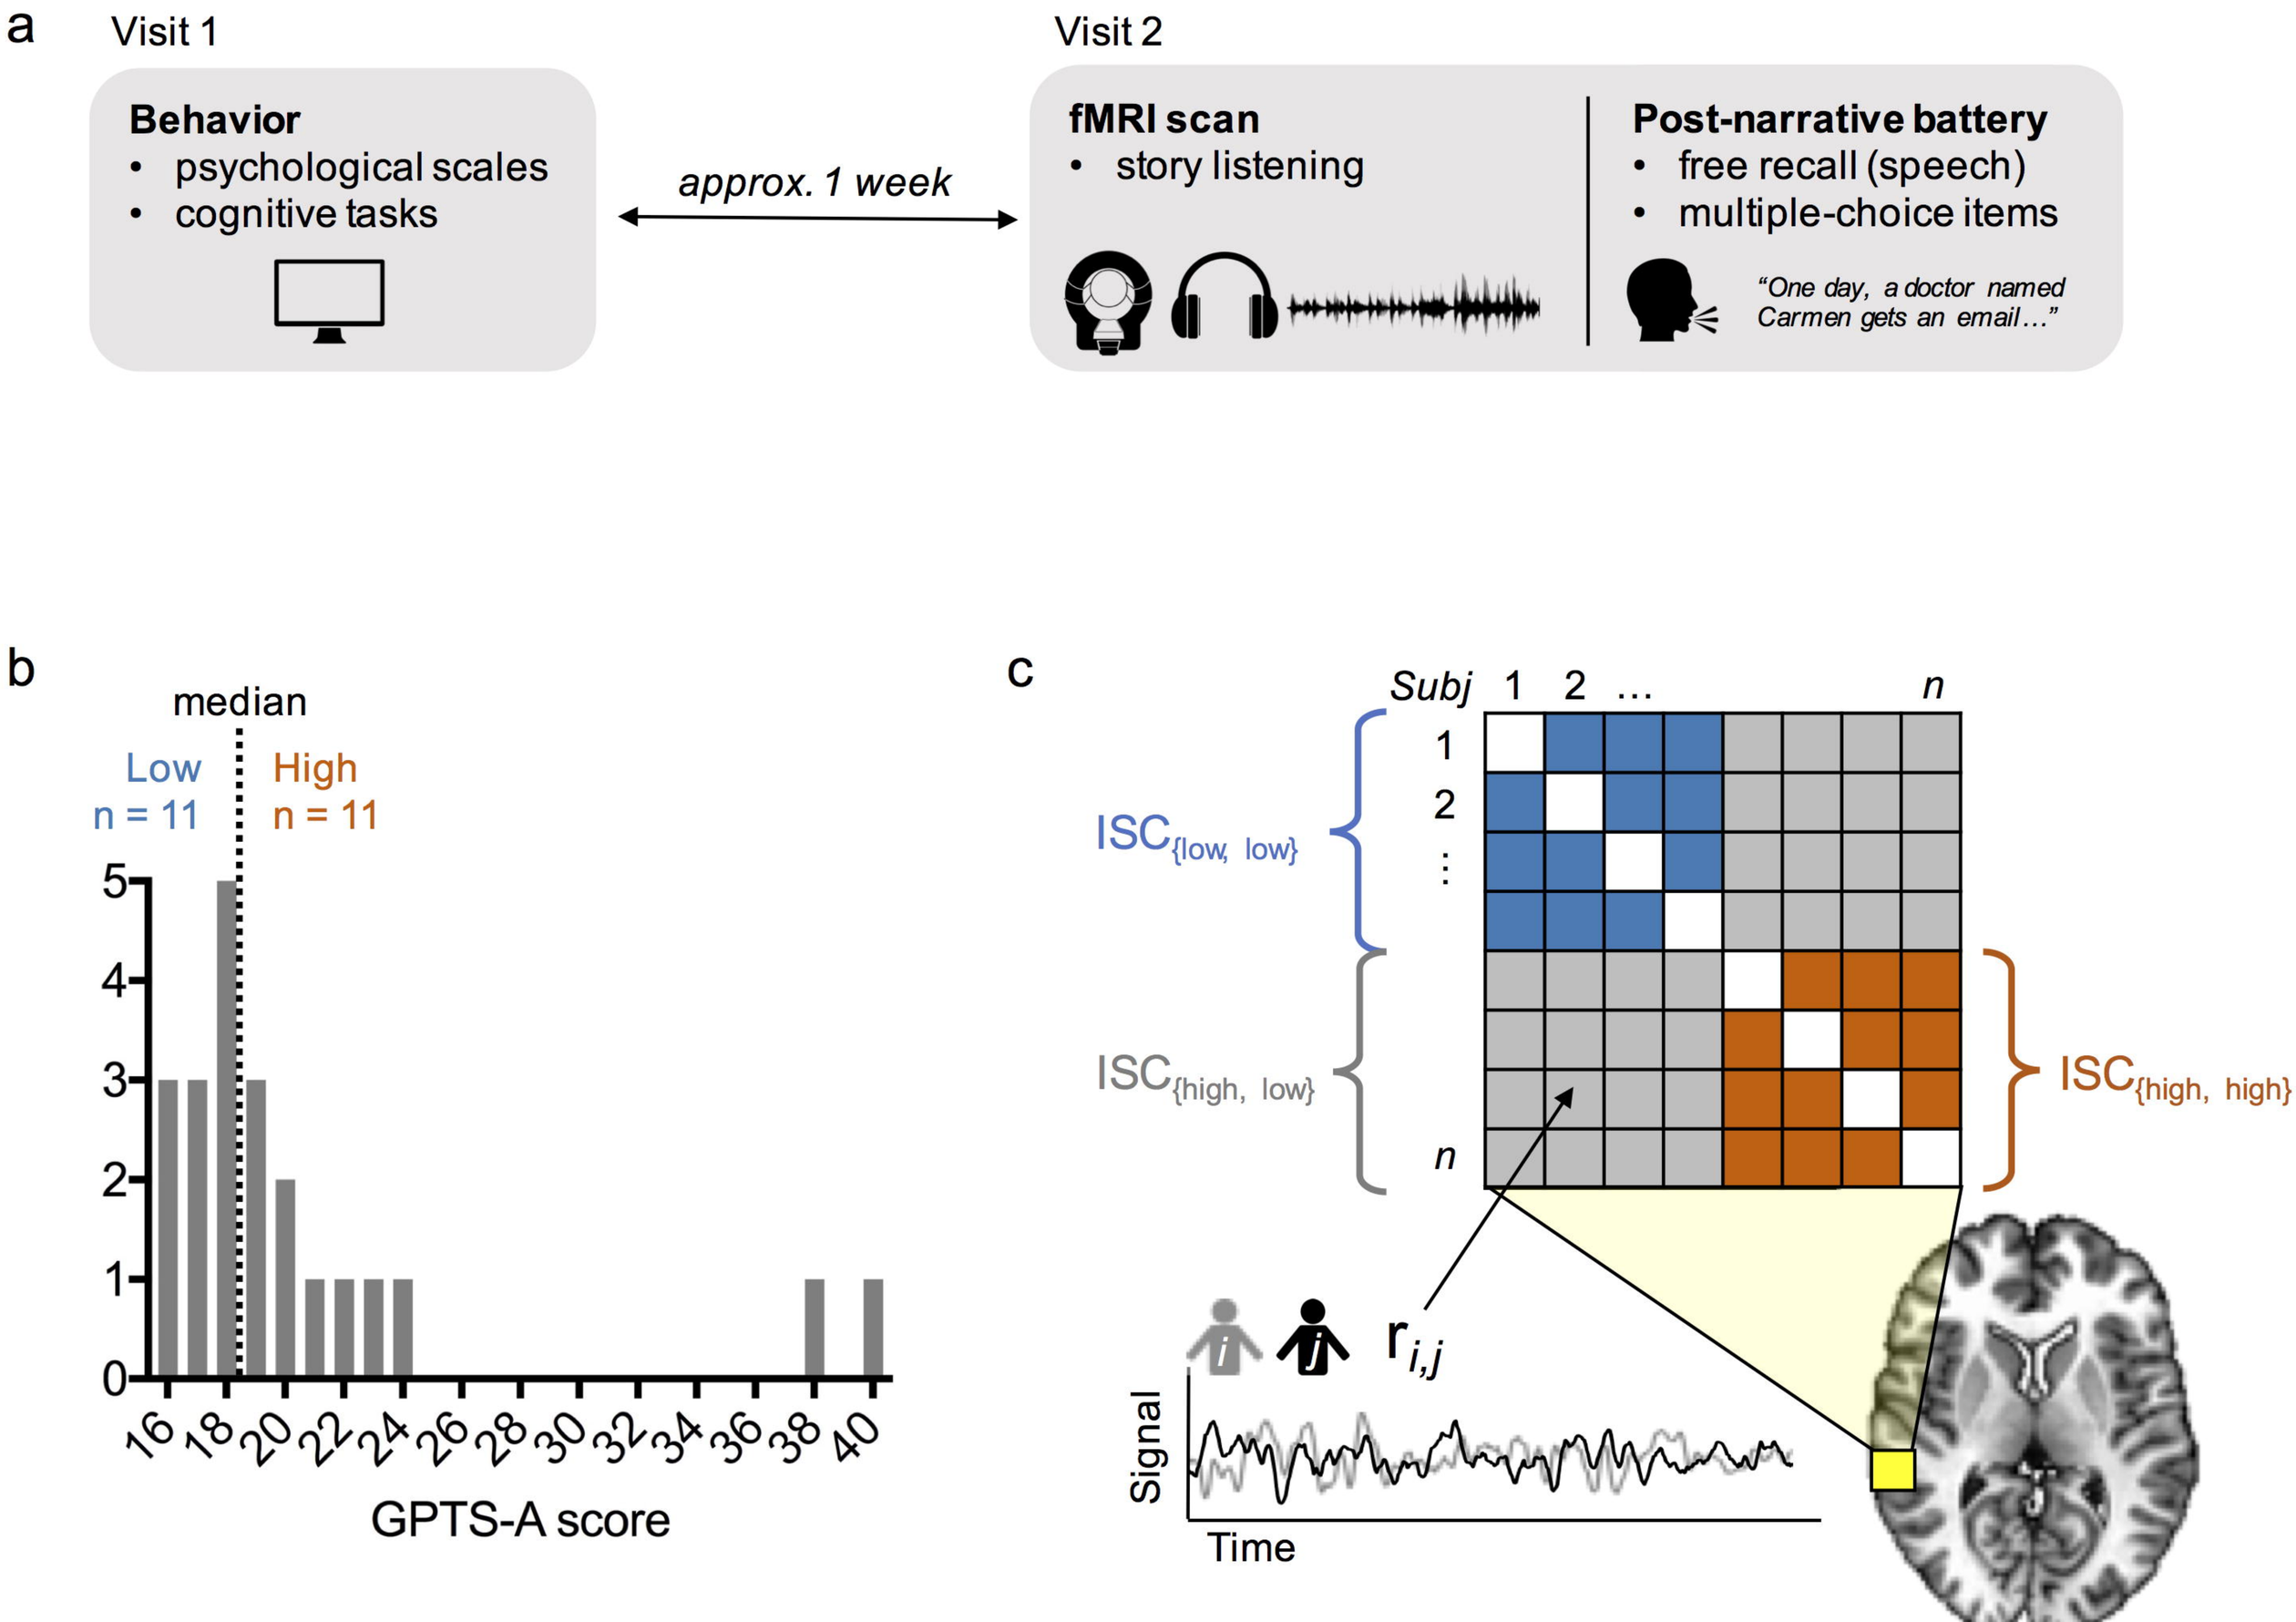
a
Visit 1
Behavior
• psychological scales
• cognitive tasks
approx. 1 week
Visit 2
fMRI scan
• story listening
Post-narrative battery
• free recall (speech)
• multiple-choice items
"One day, a doctor named Carmen gets an email…"
b
median
Low
n = 11
High
n = 11
GPTS-A score
c
Subj 1 2 … n
ISC{low, low}
ISC{high, low}
ISC{high, high}
r i,j
Signal
Time

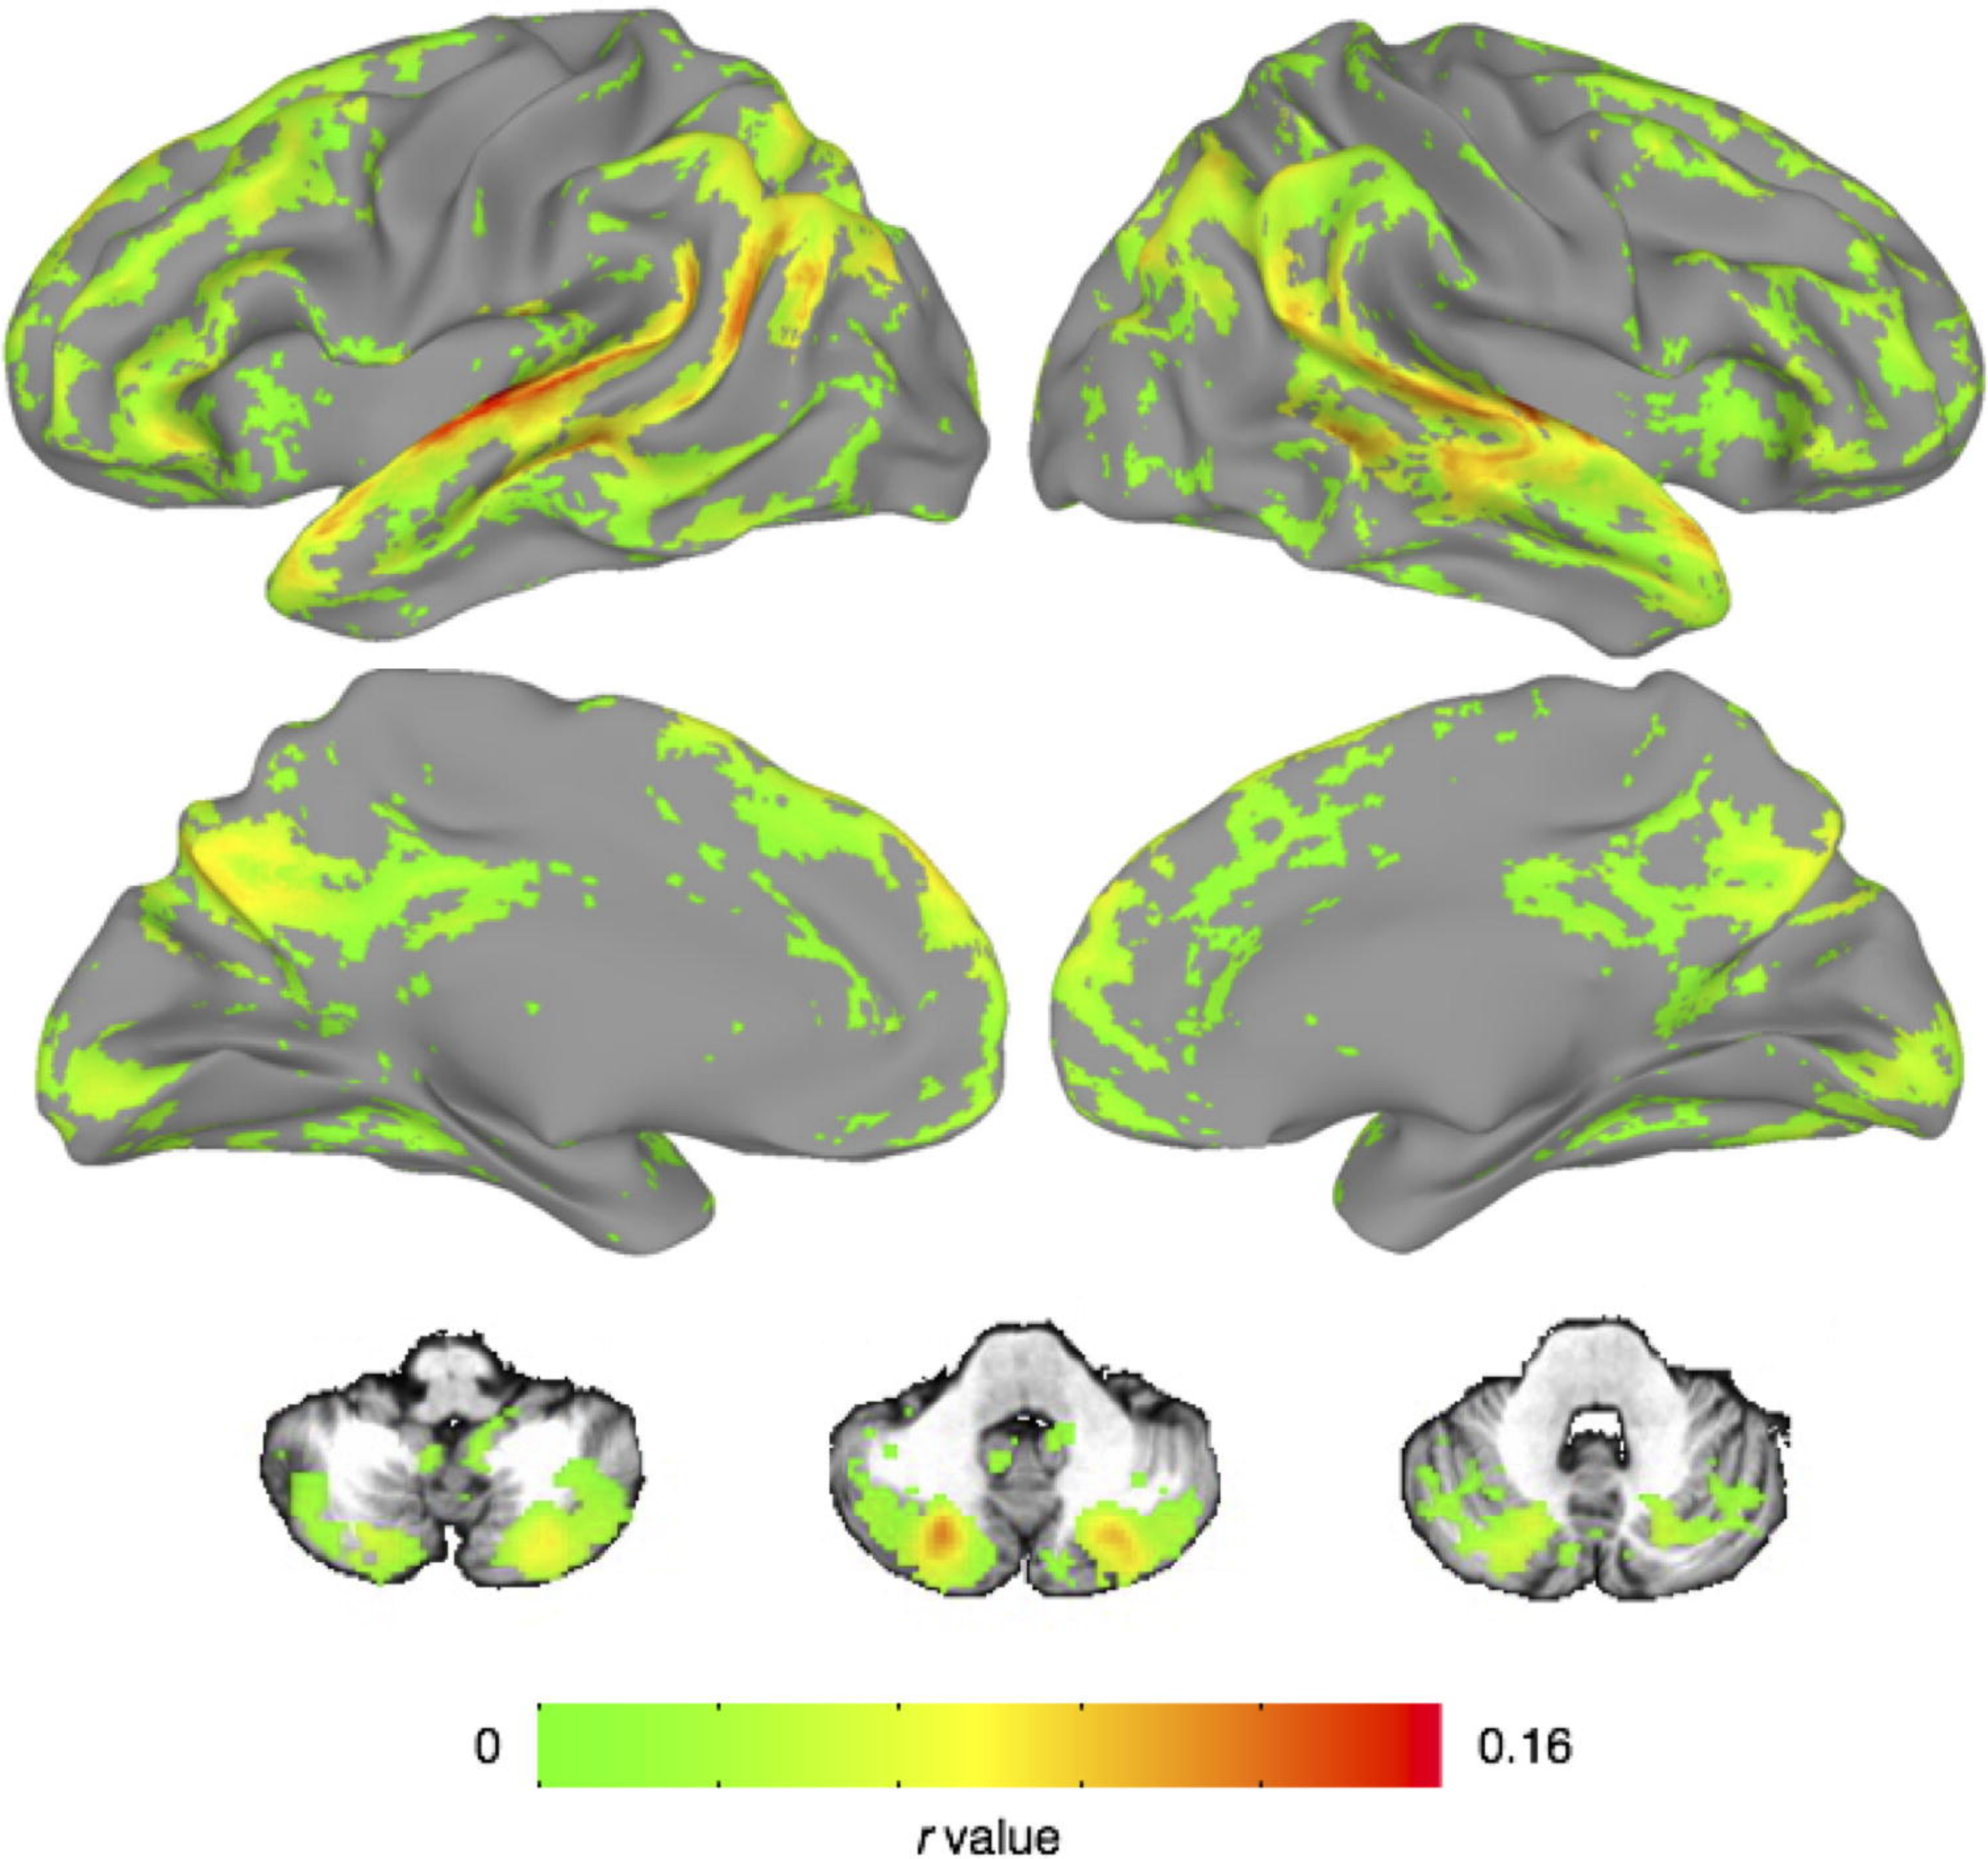
0
0.16
r value

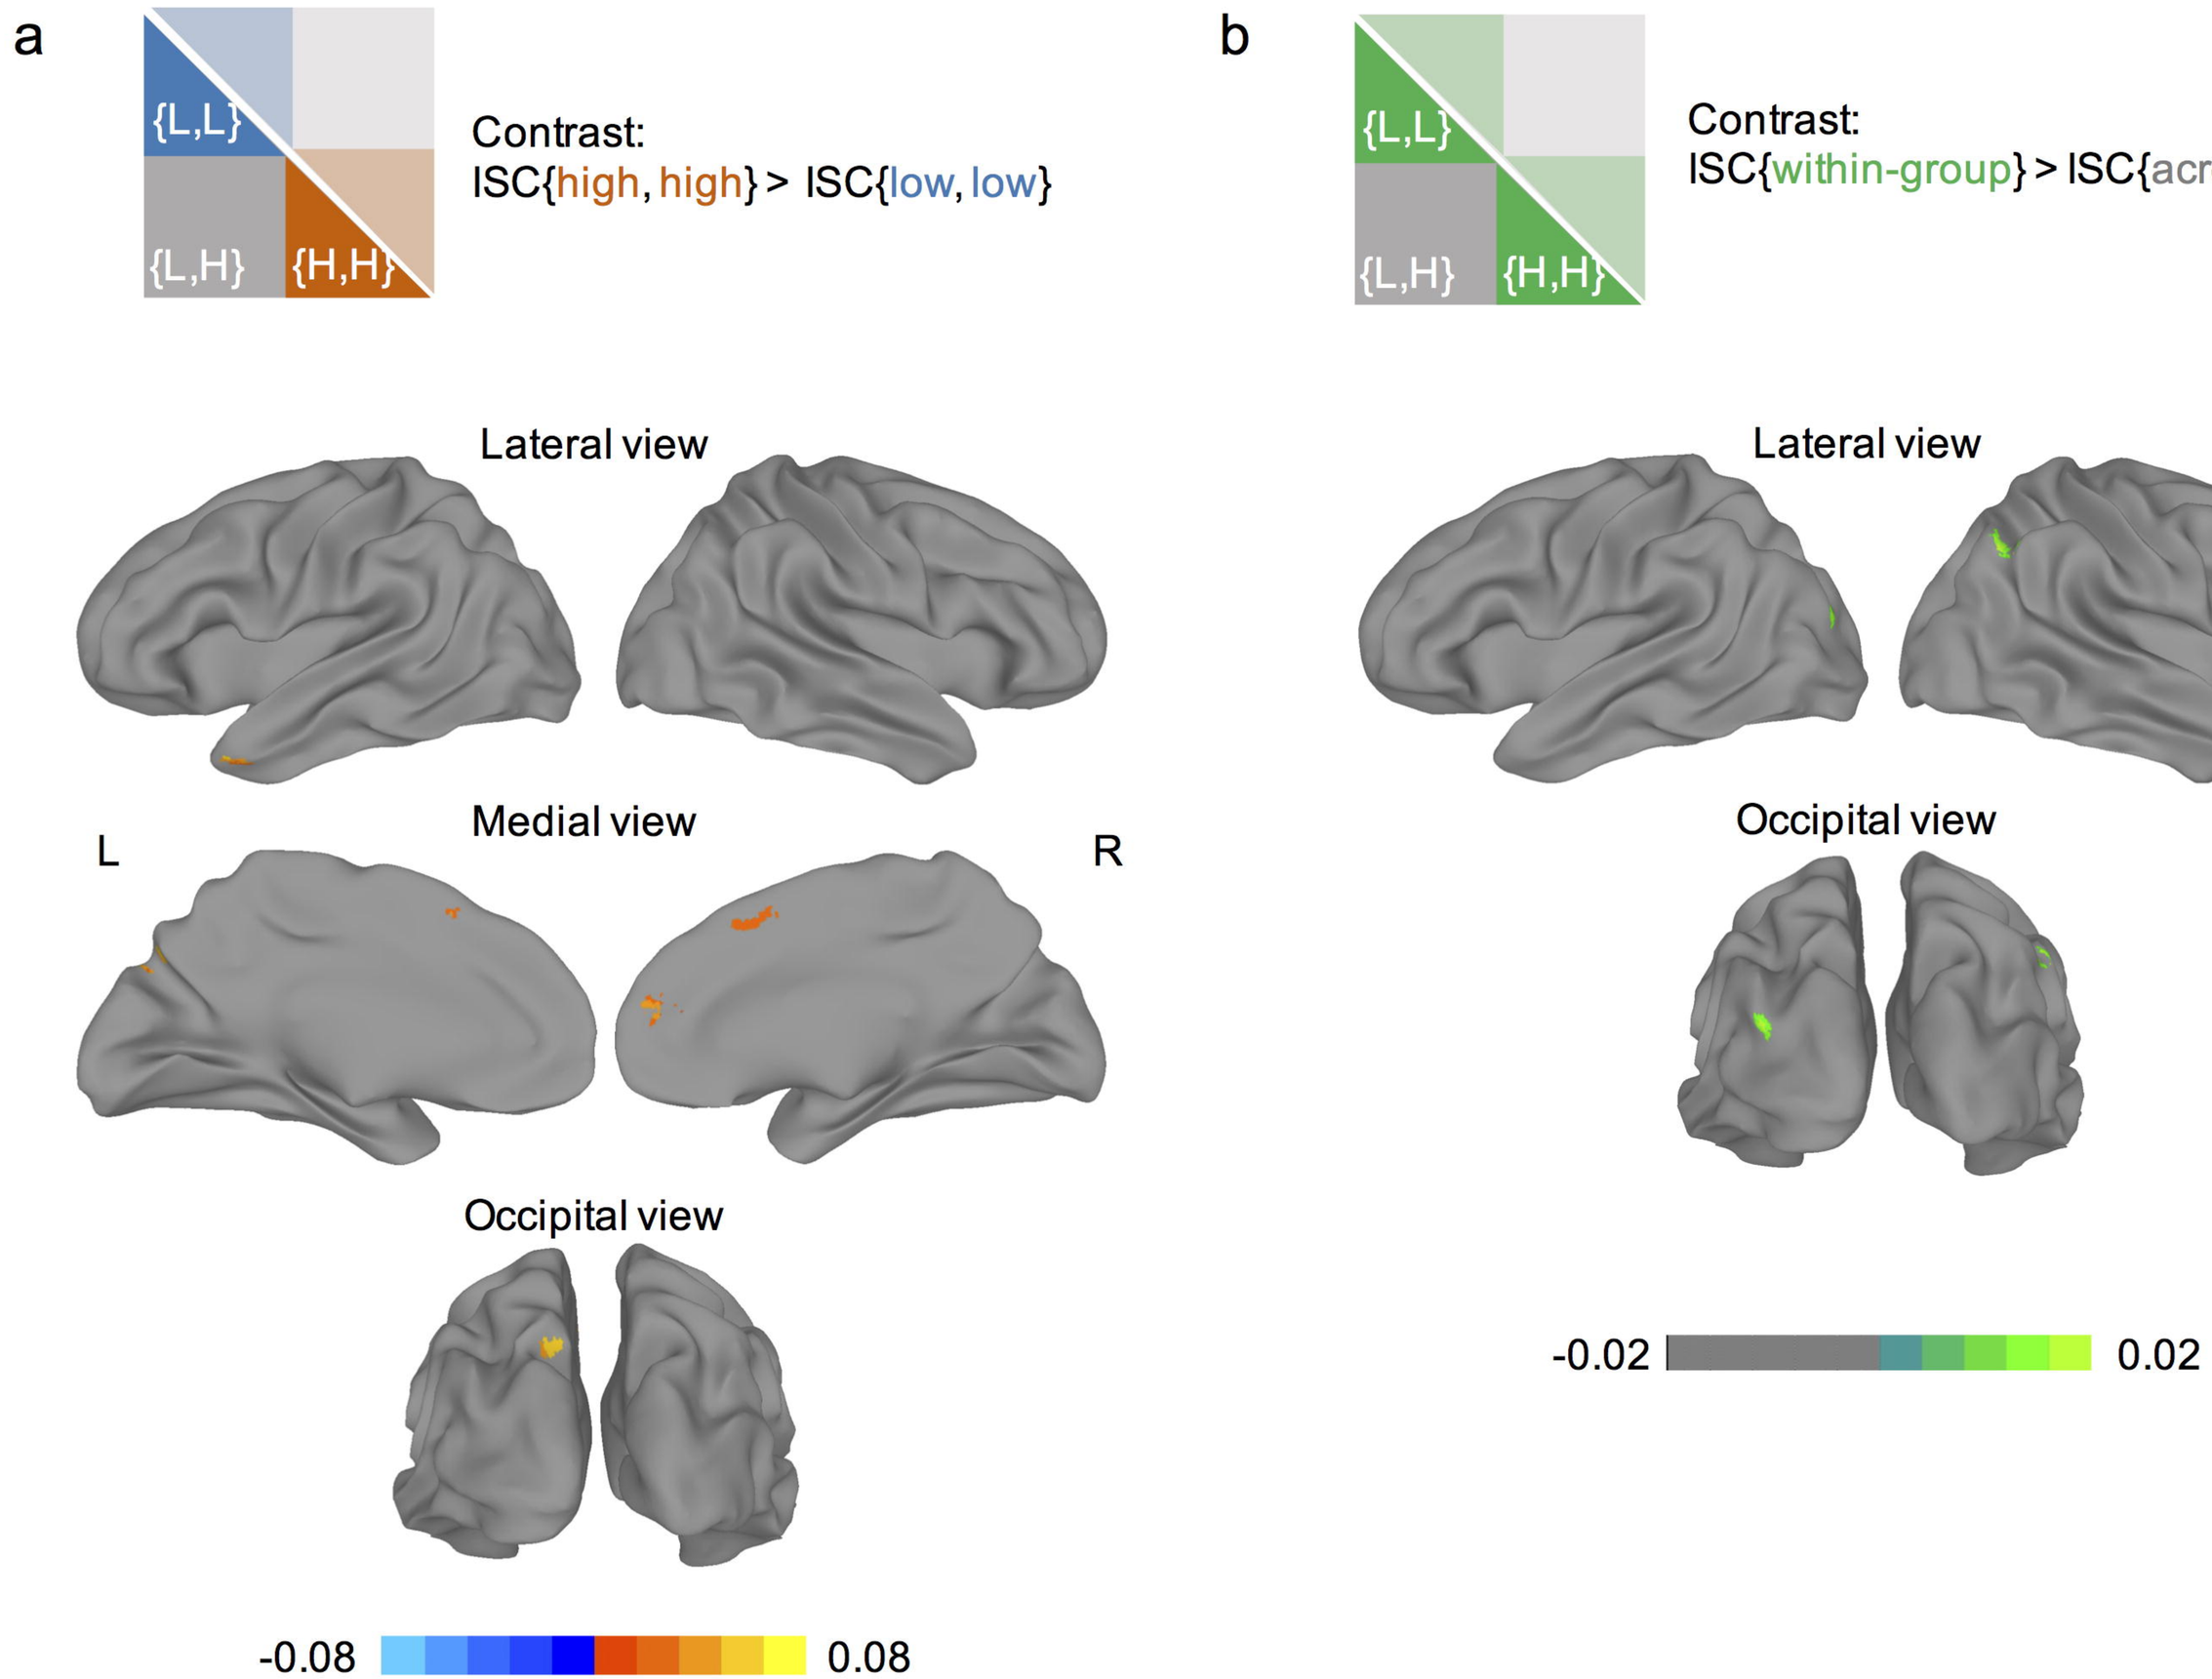
a
{L,L}
{L,H}
{H,H}
Contrast:
ISC{high, high} > ISC{low, low}
Lateral view
Medial view
L
R
Occipital view
-0.08
0.08
b
{L,L}
{L,H}
{H,H}
Contrast:
ISC{within-group} > ISC{across-group}
Lateral view
Occipital view
-0.02
0.02

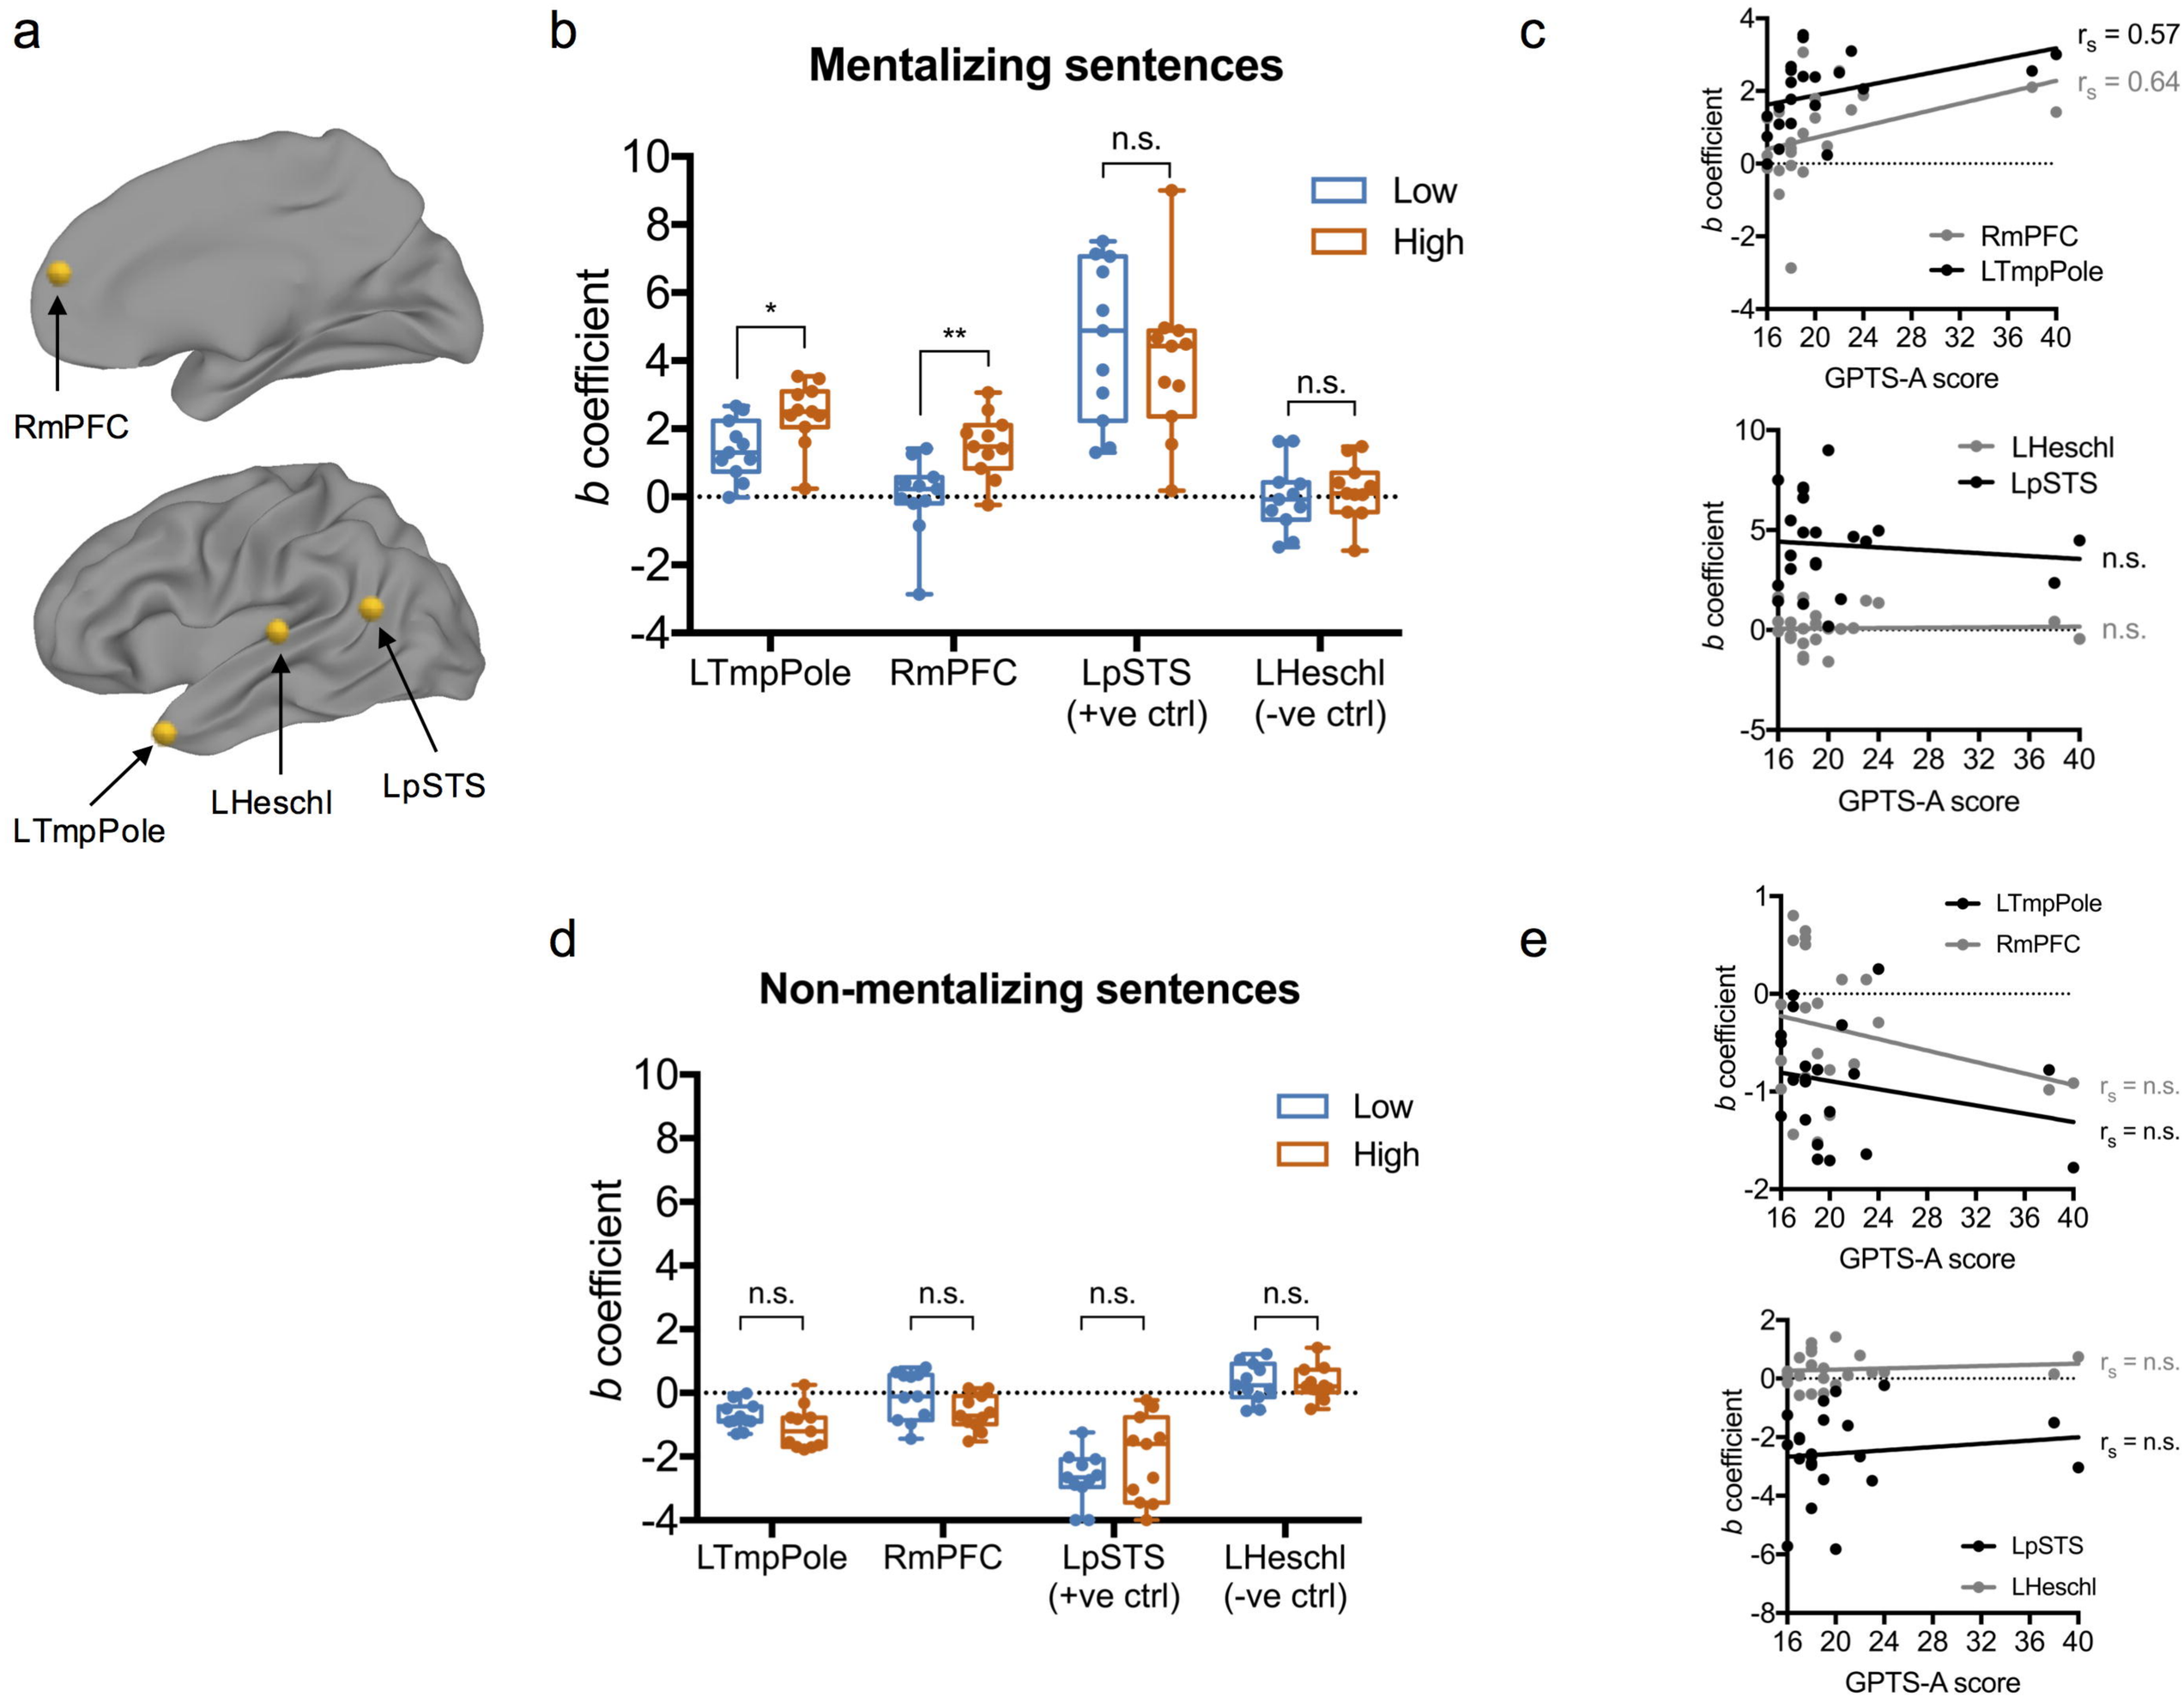
a
RmPFC
LTmpPole
LHeschl
LpSTS
b
Mentalizing sentences
b coefficient
Low
High
*
**
n.s.
n.s.
LTmpPole
RmPFC
LpSTS (+ve ctrl)
LHeschl (-ve ctrl)
c
r_s = 0.57
r_s = 0.64
RmPFC
LTmpPole
GPTS-A score
LHeschl
LpSTS
n.s.
n.s.
d
Non-mentalizing sentences
n.s.
e
LTmpPole
RmPFC
r_s = n.s.
r_s = n.s.
LpSTS
LHeschl

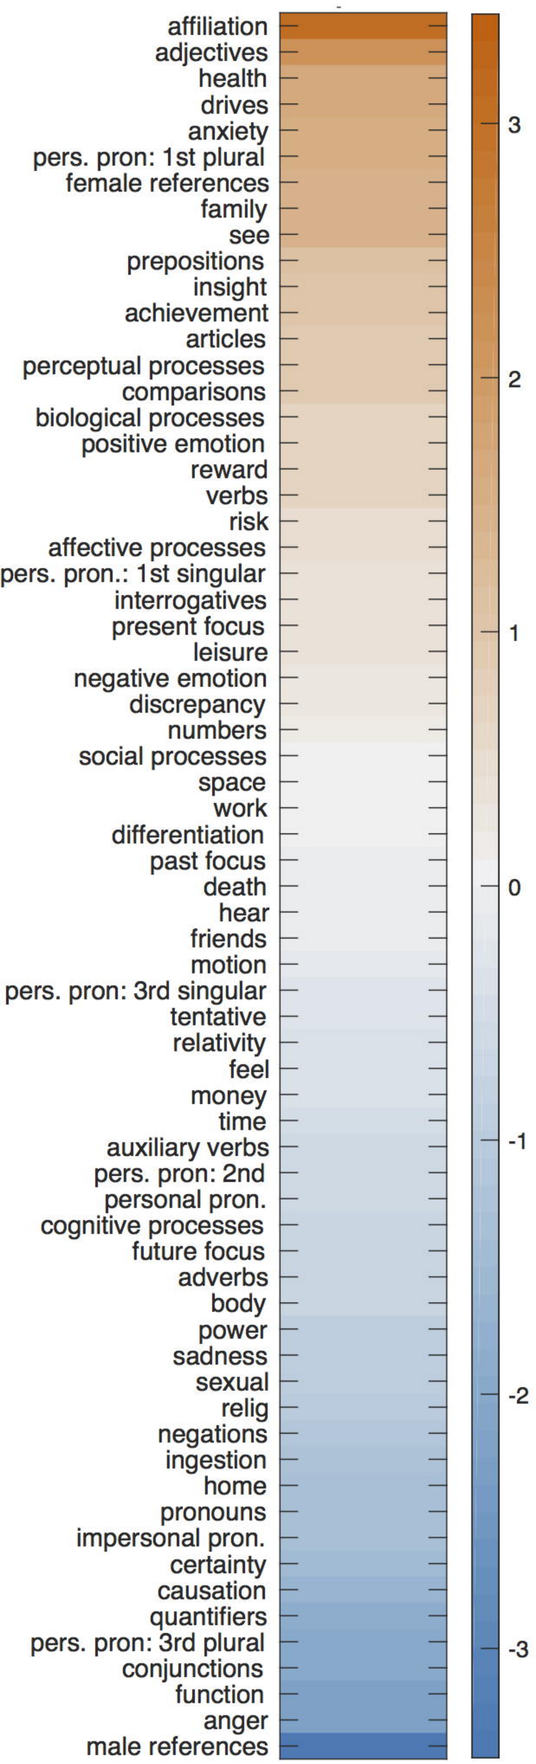

b

Affiliation

"…but they become **friends** once she gets to know her better."

Health

"…there were a series of **patients** that had a **fever** accompanied by a **rash**…"

Anxiety

"…she almost **panicked**…"

Anger

"…and she **yells** at him…"

c

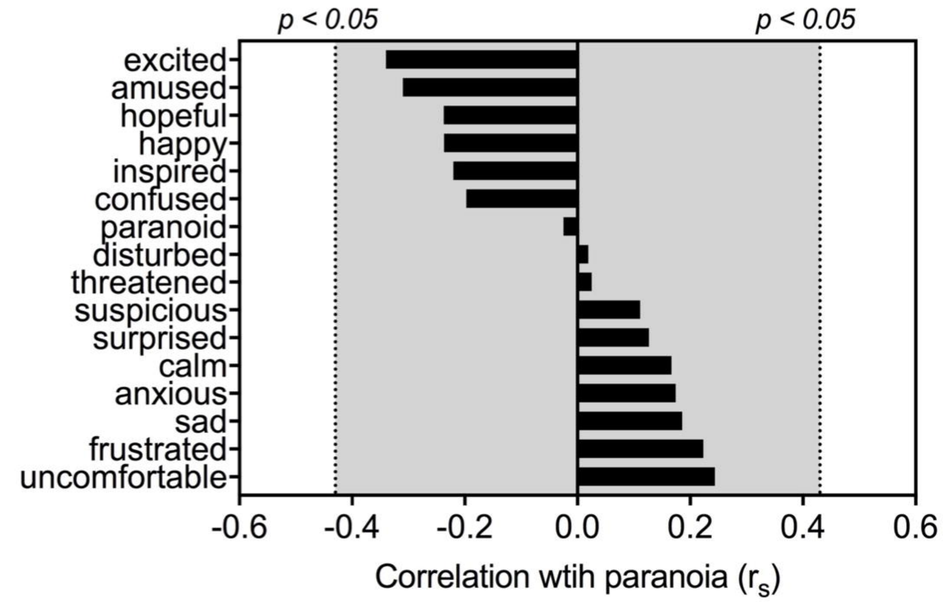
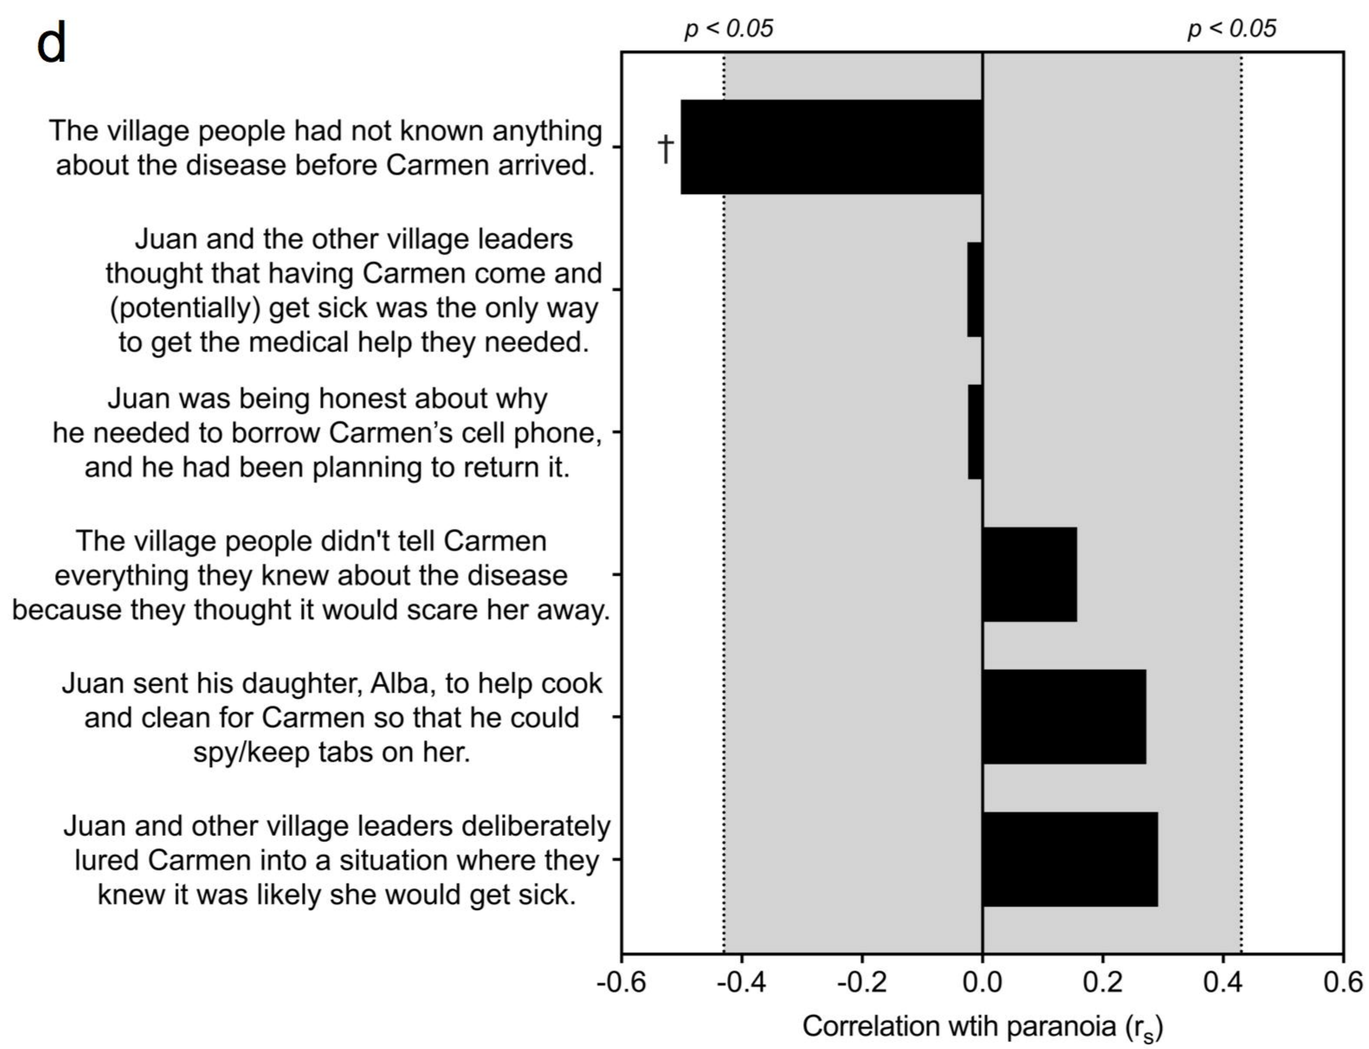
d