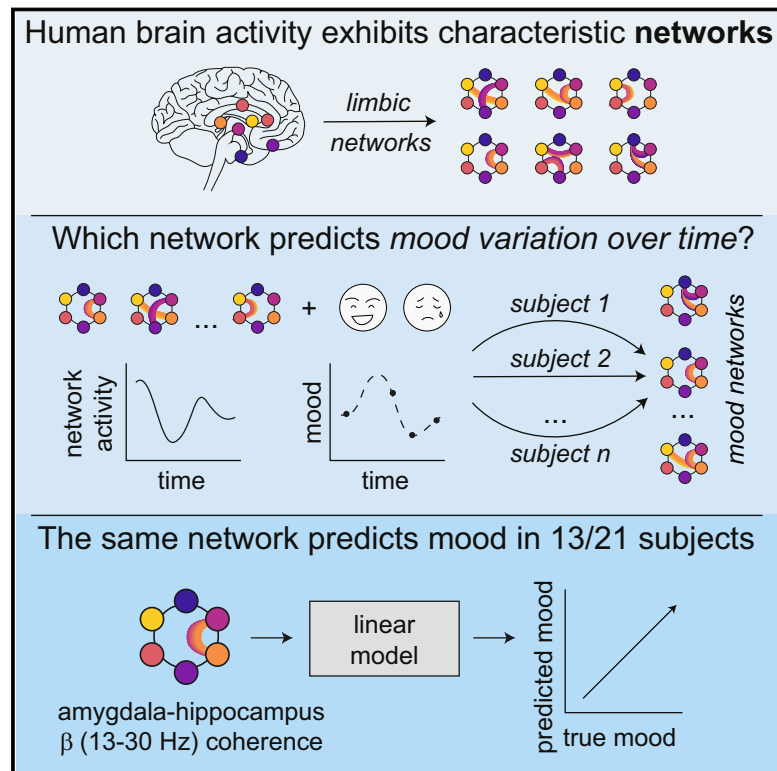


An Amygdala-Hippocampus Subnetwork that Encodes Variation in Human Mood

Graphical Abstract



Authors

Lowry A. Kirkby, Francisco J. Luongo, Morgan B. Lee, ..., Heather E. Dawes, Edward F. Chang, Vikaas S. Sohal

Correspondence

edward.chang@ucsf.edu (E.F.C.), vikaas.sohal@ucsf.edu (V.S.S.)

In Brief

Human brain activity associated with worsening mood is identified by integrating subjects' self-reported mood with intracranial recordings from the limbic system using machine learning.

Highlights

- Coordinated changes in coherence reveal conserved networks within the human brain
- A network defined by amygdala-hippocampus β -coherence predicts mood in 13 of 21 subjects
- Increased variability of coherence within this network predicts worse mood
- Subjects with higher baseline anxiety consistently have this mood-predictive network

An Amygdala-Hippocampus Subnetwork that Encodes Variation in Human Mood

Lowry A. Kirkby,^{1,2,3,4} Francisco J. Luongo,^{1,2,3} Morgan B. Lee,^{2,3,5} Mor Nahum,⁶ Thomas M. Van Vleet,⁶ Vikram R. Rao,⁷ Heather E. Dawes,^{2,3,5} Edward F. Chang,^{2,3,5,8,9,*} and Vikaas S. Sohal^{1,2,3,4,8,9,10,*}

¹Department of Psychiatry, University of California, San Francisco, San Francisco, CA 94143, USA

²Weill Institute for Neuroscience, University of California, San Francisco, San Francisco, CA 94143, USA

³Kavli Institute for Fundamental Neuroscience, University of California, San Francisco, San Francisco, CA 94143, USA

⁴Sloan-Swartz Center for Theoretical Neurobiology, University of California, San Francisco, San Francisco, CA 94143, USA

⁵Department of Neurological Surgery, University of California, San Francisco, San Francisco, CA 94143, USA

⁶Posit Science, Inc., San Francisco, CA 94111, USA

⁷Department of Neurology, University of California, San Francisco, San Francisco, CA 94143, USA

⁸These authors contributed equally

⁹Senior author

¹⁰Lead Contact

*Correspondence: edward.chang@ucsf.edu (E.F.C.), vikaas.sohal@ucsf.edu (V.S.S.)

<https://doi.org/10.1016/j.cell.2018.10.005>

SUMMARY

Human brain networks that encode variation in mood on naturalistic timescales remain largely unexplored. Here we combine multi-site, semi-chronic, intracranial electroencephalography recordings from the human limbic system with machine learning methods to discover a brain subnetwork that correlates with variation in individual subjects' self-reported mood over days. First we defined the subnetworks that influence intrinsic brain dynamics by identifying regions that showed coordinated changes in spectral coherence. The most common subnetwork, found in 13 of 21 subjects, was characterized by β -frequency coherence (13-30 Hz) between the amygdala and hippocampus. Increased variability of this subnetwork correlated with worsening mood across these 13 subjects. Moreover, these subjects had significantly higher trait anxiety than the 8 of 21 for whom this amygdala-hippocampus subnetwork was absent. These results demonstrate an approach for extracting network-behavior relationships from complex datasets, and they reveal a conserved subnetwork associated with a psychological trait that significantly influences intrinsic brain dynamics and encodes fluctuations in mood.

INTRODUCTION

Human emotion arises from interactions between brain regions within the limbic system, which includes the amygdala, hippocampus, insula, and cingulate cortex (LeDoux, 2003; Phillips et al., 2003). Much of what is known about emotion networks in the human brain comes from fMRI and positron emission tomography (PET) studies, in which specific emotions are pro-

voked using controlled stimuli while neural responses are observed non-invasively (Phan et al., 2002; Price and Drevets, 2010). These studies have shown, for example, that an insula-cingulate salience network is associated with affective experience (Seeley et al., 2007; Touroutoglou et al., 2012) and that intrinsic functional networks are altered in the setting of neuropsychiatric conditions, including depression (Greicius et al., 2007), Alzheimer's disease (Greicius et al., 2004), and schizophrenia (Calhoun et al., 2009). Importantly, PET imaging has identified changes in brain activity that are associated with antidepressant treatment response (Mayberg et al., 2000), leading to pioneering studies that have used deep brain stimulation to target specific structures to treat major depression (Mayberg et al., 2005).

Although these studies have revealed fundamental insights into the neural networks of emotion, they suffer from two major limitations. First, non-invasive imaging is restricted to short recording periods (1 to 2 hr) and often requires averaging across measurements to detect effects. As a result, these experiments cannot identify real-time neural correlates for slower changes in emotional state, such as changes in mood, that evolve over hours or days. Second, fMRI and PET are both indirect measures of neural activity with low temporal resolution on the order of seconds. As such, they cannot resolve sub-second oscillatory brain activity, which is thought to underlie information processing and cognitive function (Schnitzler and Gross, 2005). Thus, almost nothing is known about how rapid-timescale interactions between limbic brain regions contribute to changes in mood.

To address these two questions, we took advantage of a unique dataset: multi-site, semi-chronic intracranial electroencephalography (iEEG) recordings from the human limbic system, collected over several days as participants periodically rated their mood. These recordings were performed in patients with epilepsy for the primary clinical purpose of seizure localization and treatment.

Given that the amygdala is a major hub in brain networks that support emotional processing (Phelps and LeDoux, 2005), we

postulated that amygdala-containing subnetworks would also underlie mood variation. Subjects in our study were thus selected on the basis of amygdala electrode coverage in addition to three or more other limbic regions that are connected to the amygdala (Bickart et al., 2014). These included the hippocampus (ventral), cingulate cortex (inferior or superior), insular cortex, orbitofrontal cortex (anterior or posterior), and subtemporal cortex (anterior, middle, or posterior). However, not all subjects had electrode coverage in all of these regions because this was determined by clinical needs (Figure S1; Table S1).

Because iEEG measures neural activity directly from the brain, these recordings offer the potential to detect fast oscillatory network interactions that correlate in real time with changes in mood. However, they also have an inherent complication: because the number of network interactions is much greater than the number of mood ratings from a subject, there is a risk of detecting spurious correlations. To overcome this risk, we first used pattern recognition techniques (unsupervised machine learning) to define subnetworks that represent dominant interactions across the brain regions defined above. We next used multivariate regression (supervised machine learning) to determine how activity within these subnetworks related to subjects' mood ratings. Our study identifies a conserved subnetwork, driven by β -frequency (13–30 Hz) interactions between the amygdala (AMY) and hippocampus (HPC) that was present in 62% (13 of 21) of study subjects. Furthermore, when present, increased temporal variance of activity within this AMY-HPC β -frequency subnetwork consistently predicted worsening mood.

RESULTS

Limbic Activity Can Be Represented by ICNs

To identify limbic subnetworks, we first obtained a more tractable, low-dimensional data representation of our iEEG recordings by identifying a small number of patterns that explain the variation in network activity over time. Specifically, we identified intrinsic coherence networks (ICNs) that correspond to groups of brain regions that exhibit statistically significant, correlated fluctuations in rhythmic activity across recording sites (Bowyer, 2016). To identify ICNs, we first computed time series of coherence between all pairs of recording sites and in four frequency bands, which are thought to reflect distinct aspects of neural function and cognitive processing (Schnitzler and Gross, 2005): theta (θ , 4–8 Hz), alpha (α , 8–13 Hz), beta (β , 13–30 Hz), and gamma (γ , 30–70 Hz) (Figures 1A and 1B). We then performed principal component analysis to reduce dimensionality and noise, followed by independent component analysis to derive statistically significant, independent components that capture the variation in coherence across the network over time (Lopes-dos-Santos et al., 2013). Each independent component defines an ICN (Figure 1C). To visualize ICNs, we constructed connectivity diagrams showing pairs of electrodes whose temporal variance exceeds a threshold (Figure 1D). On average, we found 10 ± 4 (mean \pm SD) ICNs for each subject and frequency band, which explained $45\% \pm 10\%$ (θ), $47\% \pm 10\%$ (α), $67\% \pm 8\%$ (β), or $77\% \pm 7\%$ (γ) (mean \pm SD) of the total variance in coherence over time (Figure 1E).

ICNs from Different Subjects and Frequency Bands Cluster into Cliques

Common ICN motifs that are present in many subjects are likely to reflect strong, biologically conserved connectivity and functionality. Subjects had heterogeneous electrode placement, driven by clinical needs for seizure localization (Figure S1; Table S1), precluding a direct, one-to-one comparison of ICNs across subjects. To overcome this constraint, we projected each ICN onto a “common network template” (Figure S2A; STAR Methods). We then constructed a similarity matrix based on the Pearson correlation coefficient (CC) between all pairs of common network-projected ICNs to determine the spatial similarity between each ICN pair (Figure S2B). We visualized ICN similarity graphically, where each node corresponds to an ICN, and edges connect ICNs whose similarity exceeds a threshold ($CC > 0.55$; Figure S2B). To identify the most common ICN motifs, we considered the subgraph composed only of “core nodes” that were connected to at least 7 other nodes; this step corresponded to excluding idiosyncratic ICNs and retained 37% of all ICNs (Figure S3A). These thresholds therefore winnowed the number of connections to yield a sparse graph that is suitable for analysis. The resulting topological graph of ICNs clustered into nine “cliques” (Figure 2A; Figure S3B). Importantly, we verified that this clique structure was largely preserved using different correlation and core node thresholds that maintained a similar proportion ($\sim 40\%$) of ICNs (Figure S3C). Conversely, shuffling the similarity matrix completely eliminated cliques, confirming that the observed clique structure is non-random (Figure S3D).

Individual cliques comprised ICNs from many frequency bands and subjects (Figure 2B; Figure S3E). The most common ICN motif was dominated by interactions between electrodes in the AMY and HPC (clique 1) and present in 62% (13 of 21) of subjects (Figure 2B). Other conserved motifs were characterized by AMY-HPC interactions with other brain regions, including the insula (cliques 2 and 3, 52% of subjects), inferior cingulate cortex (clique 4, 48% of subjects), and subtemporal cortex (clique 5, 48% of subjects). The dominance of ICNs involving AMY and HPC electrodes is likely introduced by uneven electrode coverage, which favors these mesolimbic regions (Figure S1). Nonetheless, these observations suggest that there exist common modes of functional connectivity that give rise to conserved ICN motifs that are shared across individuals.

A Characteristic AMY-HPC Subnetwork Correlates with Mood in a Subset of Subjects

We hypothesized that mood-related patterns of activity should be major features of limbic system activity and, thus, be reflected in ICN dynamics. To measure subjects' subjective mood state, we used a custom-designed questionnaire (Nahum et al., 2017; Posit Science, Immediate Mood Scaler [IMS]). This consisted of a series of self-report items (Table S2) based on and validated against the Patient Health Questionnaire-9 (PHQ-9), Generalized Anxiety Disorder-7 (GAD-7), and Rumination scales (Trenynor et al., 2003). The final IMS score was found by summing individual responses to provide a single score, with higher values corresponding to a more positive instantaneous mood (Figures S4A and S4B). Note that, although we refer to this as a measure

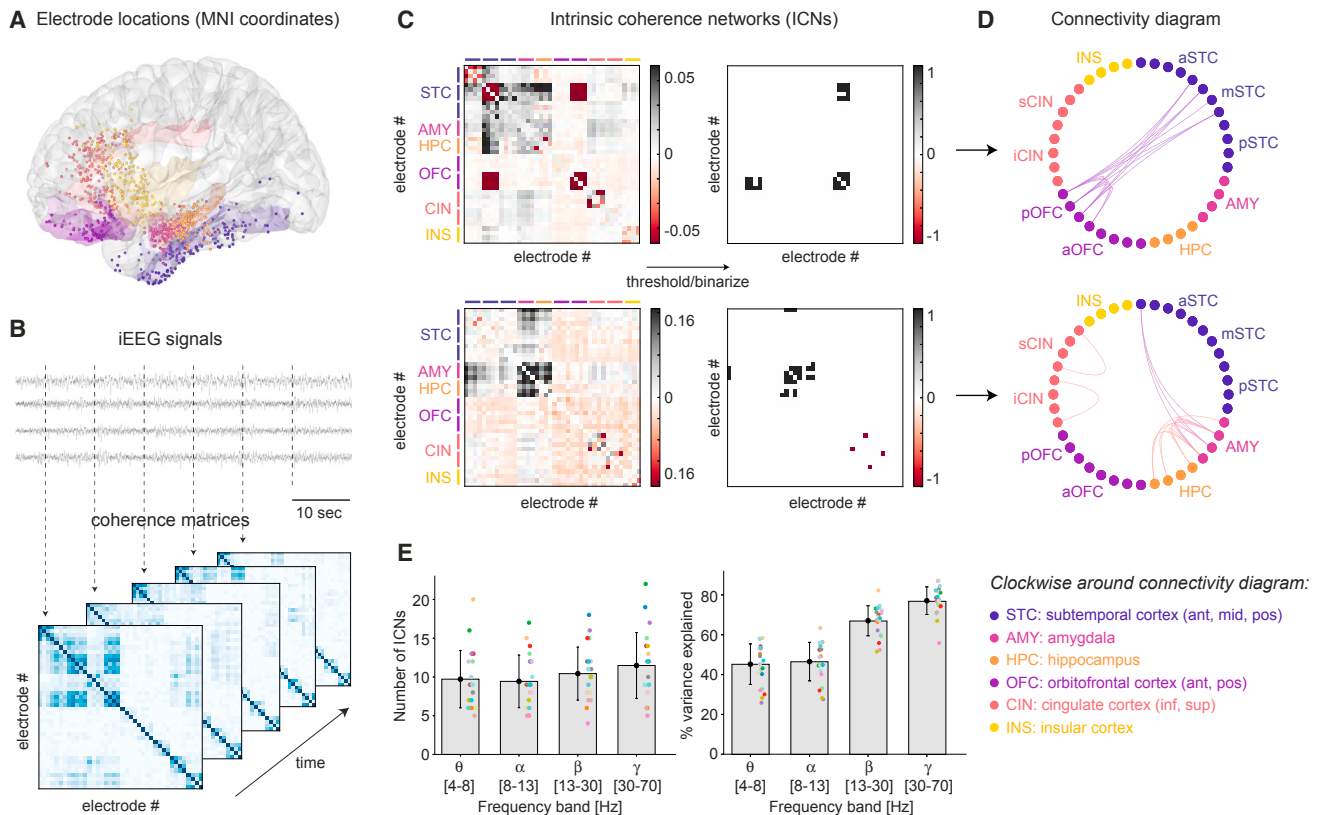


Figure 1. Limbic Activity Can Be Represented by ICNs

(A) Sagittal view of electrode locations for all subjects superimposed onto reconstructed brain in Montreal Neurological Institute (MNI) coordinate space. Electrodes are located in six brain regions: the subtemporal cortex (STC, blue), amygdala (AMY, magenta), hippocampus (HPC, orange), orbitofrontal cortex (OFC, purple), cingulate cortex (CIN, red), and insular cortex (INS, yellow).

(B) Top: example traces of intracranial electroencephalography (iEEG) signals collected from 4 STC electrodes. Bottom: example time series of coherence matrices in the β -frequency band. Each row and column correspond to a different electrode. Coherence was calculated across all electrode pairs in 10-s bins.

(C) Left: two example ICNs derived from β -frequency coherence matrices. Right: binarized ICNs derived by requiring a loading either above the 98th (black) or below the second (red) percentile.

(D) Chord representation of ICN connectivity. Each node corresponds to an electrode in the specified brain region; edges correspond to non-zero node-node pairs in the binarized ICN.

(E) Number of ICNs (left) and percent variance explained by all ICNs (right) across frequency bands for each subject (color-coded). Error bars correspond to mean \pm SD.

a, ant: anterior; m, mid: middle; p, pos: posterior; i, inf: inferior; s, sup: superior). See also [Figure S1](#).

of “mood,” it is more precisely a measure of “subjective well-being” that combines multiple aspects of mood and anxiety. This measure contrasts with scales such as the Beck Depression Inventory (BDI) or Beck Anxiety Inventory (BAI), which are weighted toward longer-timescale symptoms and designed to measure enduring trait levels of depression or anxiety rather than moment-to-moment variations in mood state (Kabacoff et al., 1997; Richter et al., 1998). We then used a model-based regression analysis (described below) to determine which ICNs correlated with variations in IMS score. This initial analysis was limited to subjects with 10 or more IMS data points and an IMS coefficient of variation of more than 5% because the remaining subjects had insufficient IMS points or variation to fit a model (6 subjects total; [Table S3](#)). The primary goal of this analysis was not to identify individualized mood-subnetwork relationships specific for each subject but, rather, to look for emergent

similarities in mood-predictive subnetworks across subjects. If we can identify conserved mood-predictive network(s) in this way, then we can then use cross-validation to test whether these generalize to a broader set of subjects.

To obtain traces of ICN activity over time, we projected the time series of coherence matrices onto each ICN. These traces exhibited a sharply peaked temporal structure, suggestive of mostly low coherence interspersed with transient periods of high coherence ([Figure 3A](#), blue traces). To capture this temporal structure, we computed the variance of ICN activity over time. For each individual subject, we then regressed the IMS scores against the average temporal variance of activity in the specific ICNs (i.e., ICNs discovered by principal component analysis [PCA] and independent component analysis [ICA] on that subjects’ dataset), using averaging windows (t_{ave}) that ranged between 5 and 30 min (in 5-min increments), centered on each

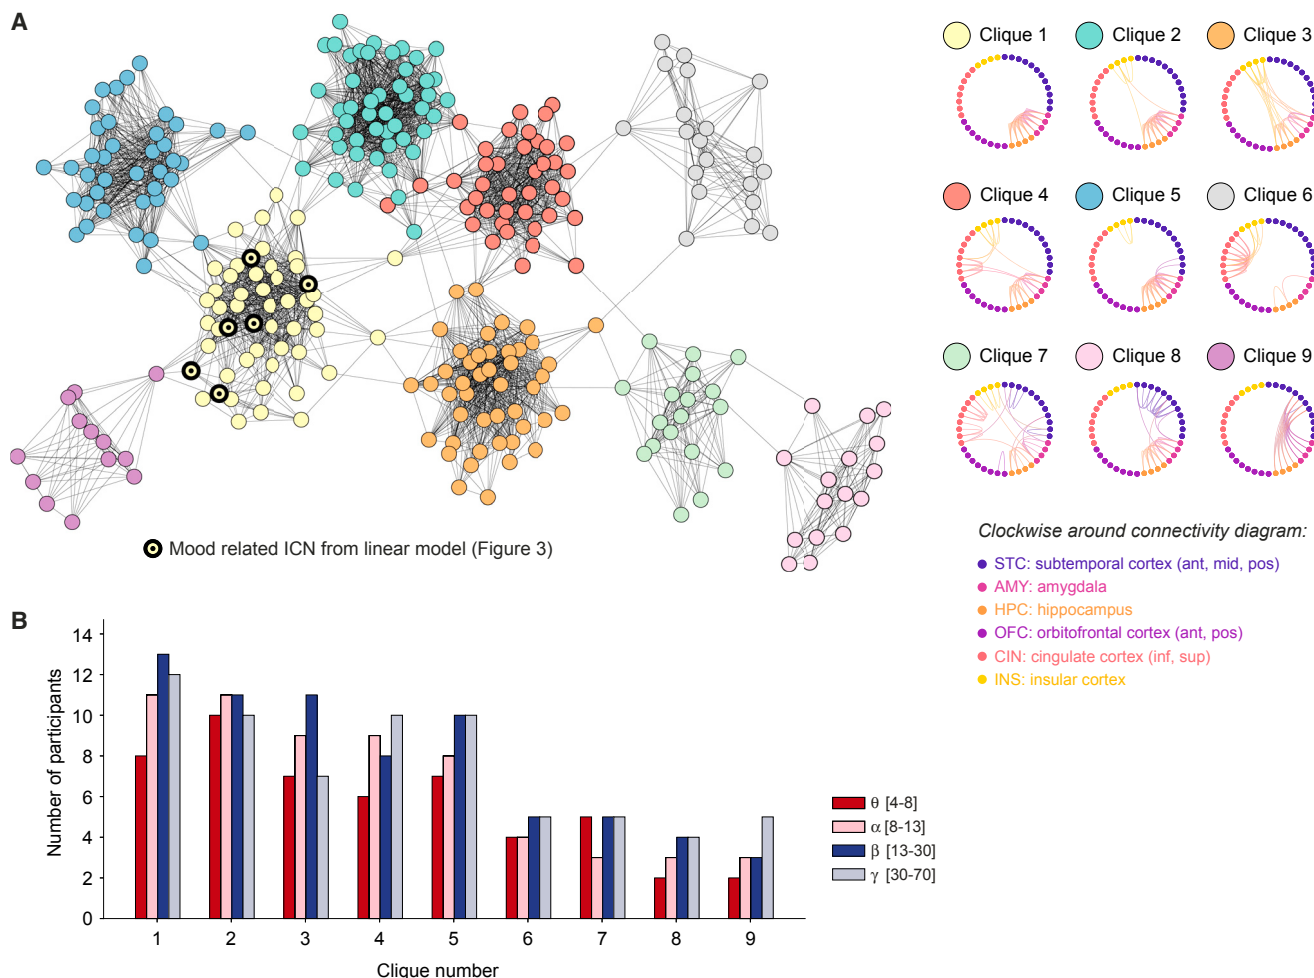


Figure 2. ICNs from Different Subjects and Frequency Bands Cluster into Cliques

(A) Left: graphical representation of ICNs that are most highly conserved across subjects. Each node corresponds to an ICN from a single subject in a single frequency band. ICN pairs that show a similar spatial connectivity pattern of brain regions are connected by an edge. Represented here are core ICNs that are connected to at least 7 other ICNs and, thus, represent spatial connectivity patterns that are most highly conserved across individuals. Color-coded clusters correspond to cliques; i.e., sets of at least 5 ICNs that are all connected to one another. ICNs highlighted with thick borders and a black dot in the node center in clique 1 correspond to ICNs with the highest coefficient for each subject from the elastic net linear model (Figure 3). Right: chord diagrams of average spatial connectivity patterns of all ICNs represented in a given clique (as in Figure 1D).

(B) Number of unique subjects per frequency band found in each clique.

See also Figures S2 and S3.

IMS time point (t_{IMS}). To mitigate the possibility of overfitting, we used an elastic net linear model (Zou and Hastie, 2005). This models IMS score based on a weighted sum of ICN variances and promotes sparsity by adding a penalty based on the sum of the weights, forcing weights for ICNs that are not correlated with IMS to zero. The best model parameters were selected for each subject using 3-fold cross-validation.

We found that using a 20-min averaging window resulted in all 6 subjects having a non-zero R^2 regression score; therefore, we chose this window for subsequent analysis. However, similar results could be obtained using a range of averaging windows, in particular for 15- and 25-min windows (Figure S4D), indicating that the model selection was robust to this temporal parameter. Across subjects, the average number of ICNs retained was

3.7 ± 2.4 (mean \pm SD; Figure 3B; Table S3). In contrast, for models using mean ICN activity (rather than variance), the average number of ICNs retained was 8.2 ± 6.3 ICNs. This suggests that mean-based models are less sparse than variance-based ones and, hence, potentially more prone to overfitting. To further validate that we identified statistically meaningful relationships and not spurious ones, we used the elastic net approach to fit linear models for IMS based on randomly shuffled time series of ICN activity (100 for each subject). For each subject, models fit to true data consistently outperformed those fit to shuffled data from the same subject; i.e., models fit to true data retained fewer ICNs and had higher R^2 ($*p < 0.005$; Figure 3B). This confirms that the ability to predict the IMS score based on ICN activity is not simply due to the number of degrees

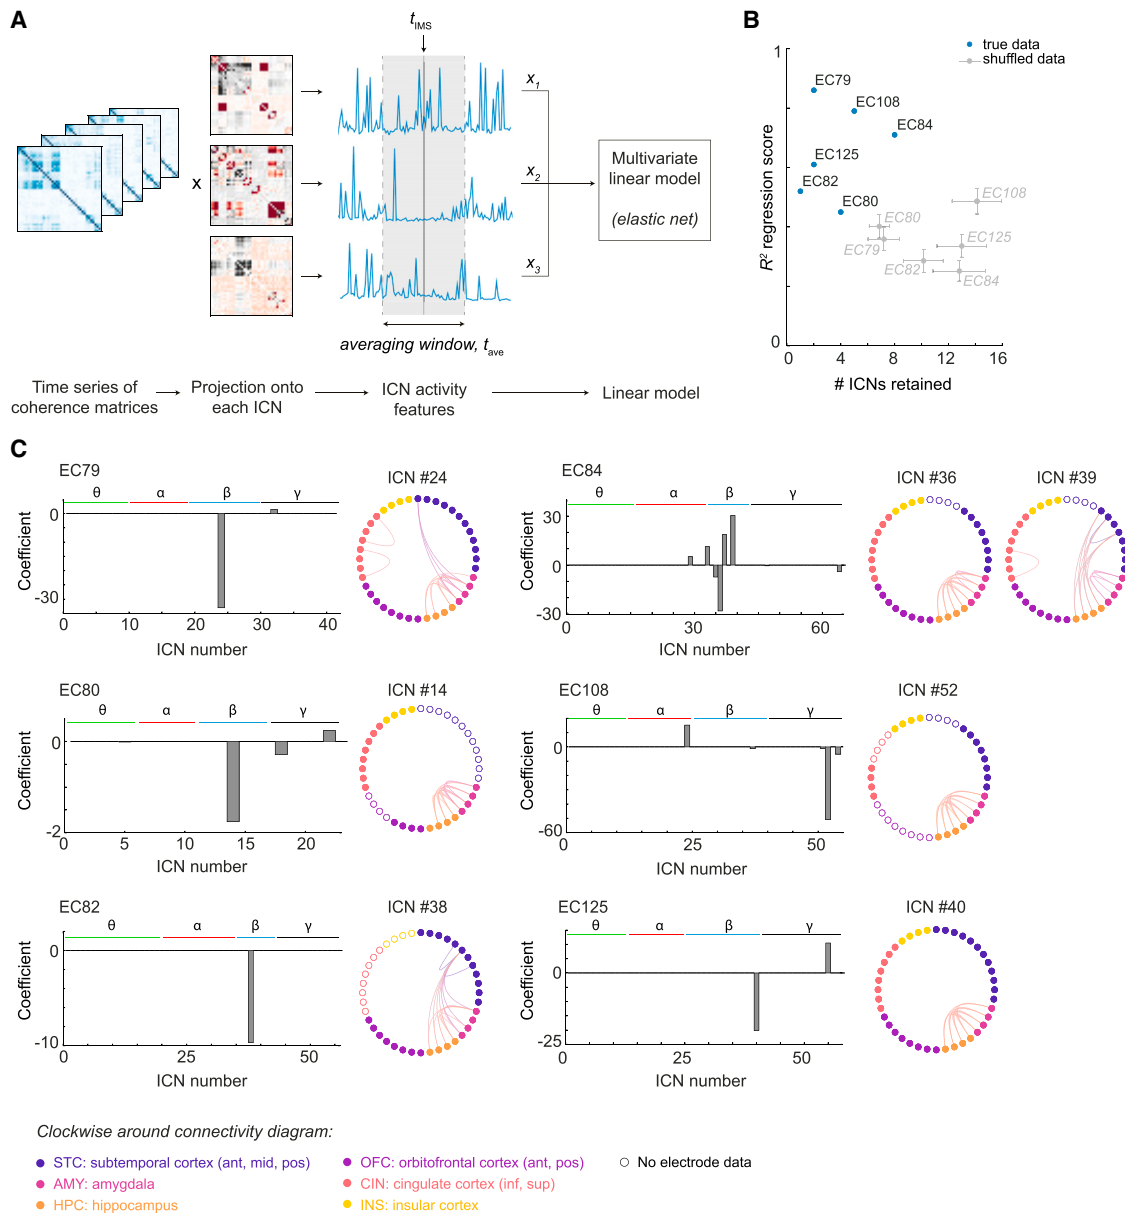


Figure 3. Linear ICN Model for IMS Prediction Converges on the Characteristic AMY-HPC Subnetwork across Subjects

(A) Schematic of the multivariate linear model used for IMS score prediction. ICN activity features (variance or mean of ICN trace) were averaged over an averaging window, t_{ave} , centered around each IMS time point, t_{IMS} .

(B) R^2 regression score versus number of ICNs retained (number of ICNs with non-zero coefficients) for models fit to each subject using a 20-min averaging window (blue data points). Grey data points correspond to R^2 regression scores for models fit to randomly shuffled time series of ICN activity for each subject. Error bars correspond to mean \pm SEM for 100 shuffled trials for each subject. All model metrics correspond to models fit using 3-fold cross-validation.

(C) Model output for 6 subjects with greater than 10 IMS data points and an IMS coefficient of variation (CV) of more than 5%. Left: ICN coefficients from the elastic net model using a 20-min averaging window. ICN numbers are grouped by frequency band (θ , α , β , and γ) and arbitrarily ranked. Right: chord diagram representation of ICN represented by largest coefficient (absolute value) (as in Figure 1D). Electrodes in subjects with no implantation in a given brain region are represented as hollow nodes.

See also Figure S4.

of freedom; if it were, then models based on shuffled data would have performed as well as those based on true data.

Although we constructed independent models for each subject, we observed a striking and highly improbable degree of

convergence. Across subjects, the single most IMS-predictive ICN (i.e., the ICN with the highest absolute magnitude coefficient) was consistently dominated by an interaction between electrodes located in the AMY and HPC and almost always in

the β -frequency band (apart from EC108, which was in the γ -band) (Figure 3C). Moreover, all of these ICNs clustered in the same clique in ICN topological space (Figure 2A, clique 1; nodes with thick borders represent the ICNs with highest coefficients). We estimated the probability that the most IMS-predictive ICN for these 6 subjects would converge in a single cluster (clique 1), with 5 of 6 of these in the same frequency band (β -band) to be 4×10^{-8} (STAR Methods). This highly significant degree of convergence indicates that the multivariate regression (performed independently for each subject) was not identifying spurious correlations driven by multiple comparisons (had this been the case, the ICNs identified by regression for each subject would more likely have been randomly distributed). Rather, the regression identified a mood-network relationship that was conserved across subjects. We next performed additional cross-validation of this relationship by determining whether it generalized to additional subjects who were not part of this initial elastic net analysis.

The Temporal Variance of AMY-HPC β -Coherence Predicts Mood across Subjects who Share the Characteristic AMY-HPC Subnetwork

Based on the convergence noted above, we hypothesized that activity in this characteristic ICN alone (the β -frequency ICN in clique 1), which we refer to as the “ β -AH ICN,” should correlate with IMS score across subjects, not just for the 6 for whom we fit a model. We found β -frequency AMY-HPC (β -AH) ICNs belonging to clique 1 in 13 of 21 subjects (Figure 2B). For each of these subjects, we regressed the IMS score against the temporal variance of activity in their corresponding β -AH ICN over a 20-min window centered around each IMS time point. In every case, we observed a negative trend between IMS score and β -AH ICN activity (Figure 4), consistent with the negative coefficients derived from the elastic net. Importantly, we are not claiming that this relationship was statistically significant on an individual subject basis—this is clearly not the case because many subjects had an insufficient number of IMS ratings to yield a significant relationship. Rather, to definitively test for a conserved relationship between activity in this subnetwork and mood, we used cross-validation; i.e., we split our population into a training set and a distinct, out-of-sample test set (Figure 5A). For each of the 13 subjects who had the β -AH ICN, we quantified the conserved feature of this subnetwork: AMY-HPC β -coherence. We computed the time series of AMY-HPC β -coherence for each subject by averaging across all electrode pairs within these two structures. Similar to the temporal structure of ICN activity, AMY-HPC β -coherence showed sharp peaks in its temporal structure, which we again captured using variance over time.

Our first training set comprised the 6 subjects used for the elastic net analysis. Before pooling across subjects, we normalized data on an individual subject basis by converting both IMS and the variance of AMY-HPC β -coherence to Z score values. For these 6 subjects, we regressed the IMS score against the temporal variance of AMY-HPC β -coherence (computed over a 20-min window centered on each IMS time point). This revealed a strong negative correlation between IMS and the variance of AMY-HPC β -coherence ($r^2 = 0.42$, $**p < 10^{-10}$; Figure 5B).

Thus, even though the elastic net framework led us to discover a mood-related subnetwork, the neural correlate for mood is present in less processed and more readily interpretable data (the temporal variance of coherence between two interconnected brain regions). We then tested whether the linear model derived from these 6 subjects generalized to the broader population by predicting IMS for the remaining subjects, from their (Z score-normalized) variance of AMY-HPC β -coherence. This model was strongly predictive of IMS for the 7 remaining subjects who shared the characteristic β -AH ICN ($R^2 = 0.49$, $**p < 10^{-3}$; Figure 5B). Interestingly, the model had no predictive power in the 8 subjects who lacked this ICN ($R^2 < 0$, $p = 0.87$; Figure 5B). Note that a negative R^2 in this case indicates worse prediction than a horizontal line).

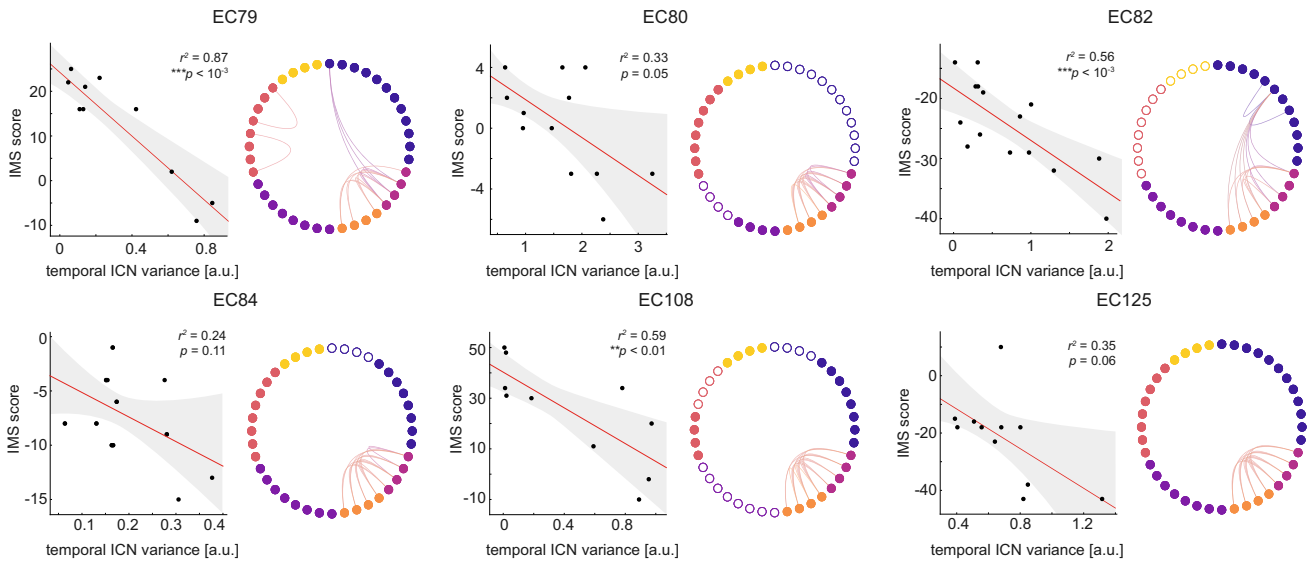
Next, to demonstrate that there is nothing exceptional about the 6 subjects we initially used for the elastic net analysis, we used the other 7 subjects who shared the characteristic β -AH ICN as the training set and fit a new linear model. Again, this model predicted IMS for the out-of-sample population, now the 6 subjects with the β -AH ICN who were used for the elastic net analysis ($R^2 = 0.42$, $**p < 10^{-6}$; Figure S5). This was expected based on Figure 5B and simply illustrates that we find the same IMS-coherence relationship regardless of which set of subjects we analyze (so long as they have the β -AH ICN). Again, this model did not predict IMS for the 8 subjects who lack the β -AH ICN ($R^2 < 0$, $p = 0.88$; Figure S5). This form of 2-fold cross-validation indicates that it is not simply the case that there is a relationship between the variance of AMY-HPC β -coherence and mood in all subjects who share the characteristic β -AH ICN. Rather, the same quantitative relationship holds across these individuals.

Variance of AMY-HPC β -Coherence Detects High-Coherence Outliers

To further characterize the temporal dynamics of the relationship between IMS and the variance of AMY-HPC β -coherence, we pooled data across all 13 subjects who shared a β -AH ICN and repeated our regression analysis using a range of averaging windows and window start positions (Figure 6A). Consistent with our elastic net findings, the correlation was strongest using a 20-min averaging window centered on each IMS time point. However, significant relationships were also present over a range of averaging windows and start positions, demonstrating that this relationship is robust for a range of temporal parameters (Figure 6B). Moreover, although we found that this correlation was strongest for β -frequency activity, a similar but slightly weaker one was also present for γ -frequency activity (Figure S6A).

We next tested whether mean AMY-HPC β -coherence correlated with IMS score (Figure 6C). Mean coherence also showed a negative trend with IMS, but the associated p value was 10^8 -fold weaker than for coherence variance. To understand how coherence values mapped onto coherence variance, we plotted cumulative distributions of coherence values corresponding to periods when the variance was above (“high variance”) or below (“low variance”) its median value. The tails of these two distributions were markedly different. By contrast, their means showed a less dramatic difference (Figure 6D). This confirms the intuition that variance should be more sensitive to high-coherence

A Participants used in elastic net analysis



B Participants not used in elastic net analysis with β -AH ICN

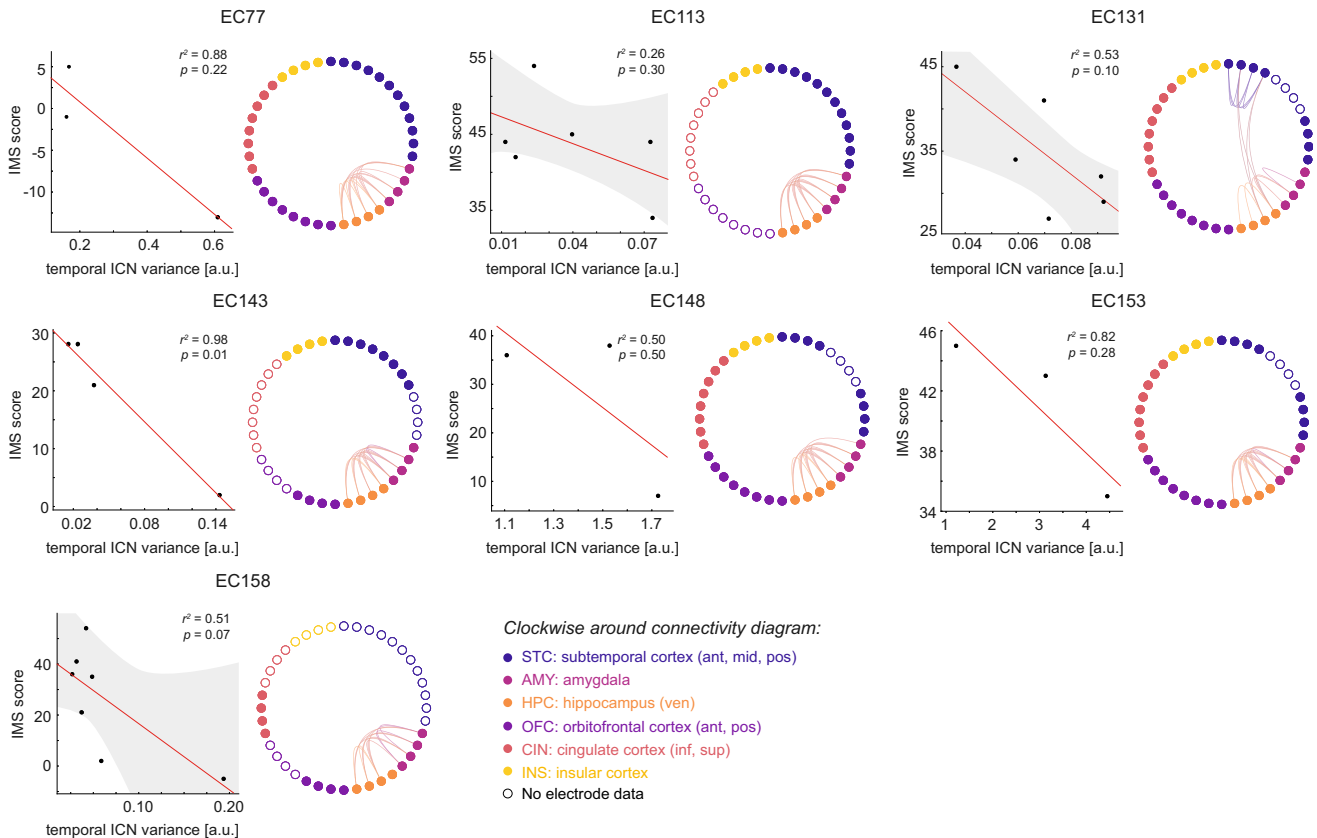
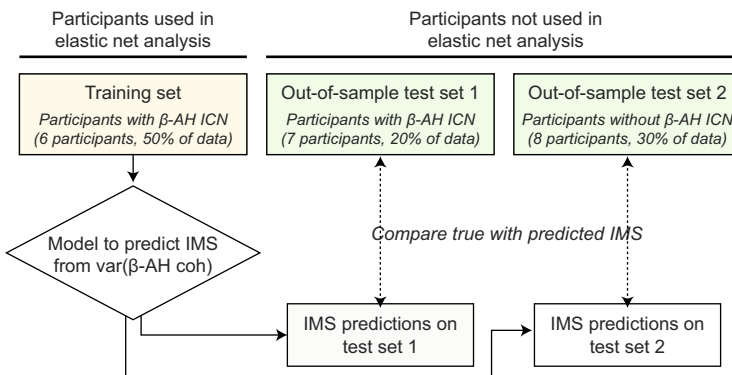


Figure 4. Temporal Variance of Activity in the Characteristic β -AH Network Trends Negatively with IMS Score across Subjects

(A) Subject-specific regression of IMS score against temporal variance of β -AH ICN activity averaged over a 20-min interval around each IMS time point for 6 subjects used in the elastic net analysis. Grey confidence intervals represent 95% confidence on the regression estimate. Chord diagrams represent β -AH ICN for that subject, as in Figure 1D. Electrodes in subjects with no implantation in a given brain region are represented as hollow nodes.

(B) As in (A), but for the 7 subjects with the β -AH ICN who were not used in the elastic net analysis.

A Model schematic



B




Figure 5. Temporal Variance of AMY-HPC β -Coherence Predicts Mood across Subjects Who Share a β -AH Network

(A) Schematic of data partitioning into training and test datasets for IMS prediction based on variance of AMY-HPC β -coherence. A linear model was trained on data from the 6 subjects used in the elastic net analysis. The model was tested on two out-of-sample, cross-validation test sets of subjects not used in the elastic net analysis: (i) test data 1, 7 subjects with β -AH ICN; (ii) test data 2, 8 subjects without β -AH ICN. Predicted IMSs were compared with target (true) IMSs for each group of test data.

(B) Left: regression of IMS score against variance of AMY-HPC β -coherence over a 20-min averaging window centered around each IMS score time point and Z score normalized within each individual for the 6 subjects used in the elastic net analysis. The confidence interval represents 95% confidence on the regression estimate. A linear model was fit to these data and used for mood predictions for the remaining subjects. Center/right: IMS score predictions using a linear model for subjects with (center) and without (right) β -AH ICN. R^2 regression scores describe how well the model fits the data, where a negative score corresponds to a worse prediction than a horizontal line. The p values assess significance using a 10^6 -fold permutation test for residuals against the line $y = x$ (dashed line).

See also Figure S5.

outliers than the mean. We postulate that this increased sensitivity to high-coherence outliers drives the strong correlation we observed between the temporal variance of AMY-HPC β -coherence and IMS score. Moreover, the IMS score showed only a weak or no correlation with β -coherence or β -power within either the AMY or HPC alone (Figures S6B and S6C), suggesting that mood is encoded by AMY-HPC interactions rather than either individual region.

Epileptiform Activity in the AMY and HPC Does Not Correlate with Mood

The AMY and HPC are common epileptogenic regions. To control for potential effects of epileptiform activity on mood, we first detected AMY and HPC interictal discharges (transient spikes

in the raw voltage time series) around each IMS time point using a peak detection algorithm (Meliza and Margoliash, 2012; Figure S7A). The AMY discharge rate was only weakly correlated with IMS score or the temporal variance of AMY-HPC β -coherence ($r^2 = 0.04$; Figures S7B and S7C, left) but was more strongly correlated with mean AMY-HPC β -coherence ($r^2 = 0.19$; Figure S7D). However, the negative relationship between IMS score and AMY-HPC β -coherence remained even after we controlled for effects of the AMY discharge rate in our regression analysis (Figure S7E). The HPC discharge rate was not correlated with IMS score or variance of coherence and only weakly correlated with mean coherence (Figures S7B–S7D, right). As an additional control, we examined the relationship between anticonvulsant medication and mood for three subjects with the strongest

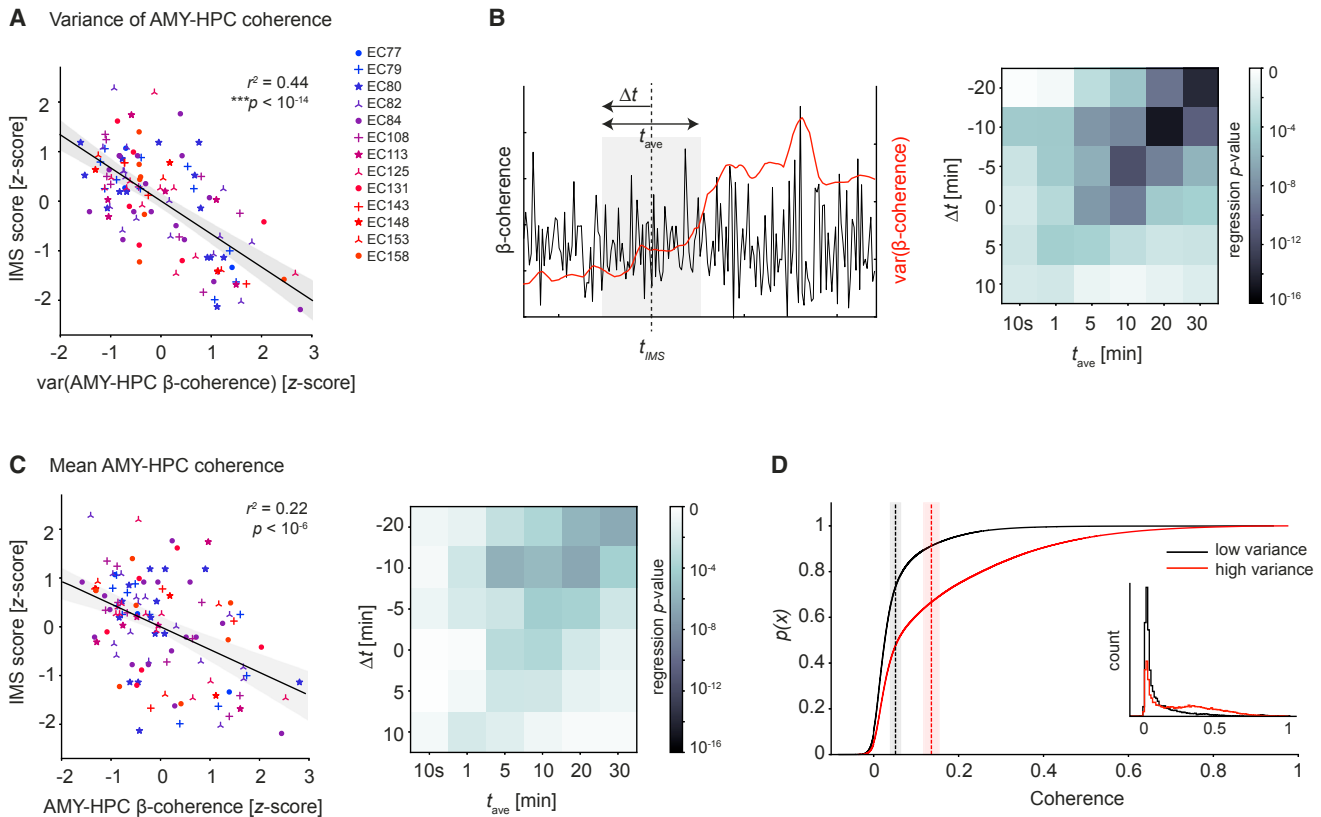


Figure 6. Variance of AMY-HPC β -Coherence Detects High-Coherence Outliers

(A) Regression of IMS score against variance of AMY-HPC β -coherence over a 20-min averaging window centered around each IMS score time point and Z score normalized within each subject for 13 of 21 subjects with β -AH ICN. The gray confidence interval represents 95% confidence on the regression estimate. (B) Right: schematic of AMY-HPC coherence calculated in 10-s contiguous bins (black) and coherence variance calculated over a 60-s sliding interval (red). For an IMS score recorded at time t_{IMS} , the mean variance was calculated over an averaging window, t_{ave} , at a time Δt from t_{IMS} . Left: heatmap showing regression p values for a range of values of t_{ave} and Δt for subjects with β -AH ICN. The regression plot in (A) corresponds to $t_{ave} = 20$ min and $\Delta t = -10$ min. (C) Left: regression of IMS score against mean AMY-HPC β -coherence over a 20-min averaging window centered around each IMS score time point and Z score normalized within each subject for 13 of 21 subjects with β -AH ICN. The gray confidence interval represents 95% confidence on the regression estimate. Right: heatmap of regression p values across a range of averaging windows, t_{ave} , and start times from each IMS data point, Δt (as in B). The regression plot on the left corresponds to $t_{ave} = 20$ min and $\Delta t = -10$ min. (D) Cumulative probability density distribution of coherence values for variance epochs below the median (low variance, black) versus above the median (high variance, red). Dashed lines and confidence intervals correspond to mean and SD of each distribution. Inset: example distribution of coherence values in low or high variance states for a single subject.

See also [Figures S6](#) and [S7](#).

IMS correlation: EC79, EC82, and EC108 ([Figure S7F](#)). Although changes in mood appeared to coincide with changes in anticonvulsant dosing for one subject (EC79), this was not the case for the other two: both increases and decreases in mood occurred after anticonvulsants had been discontinued for one subject (EC82) and while the other (EC108) was on a constant dose. These observations illustrate that the mood variation we observed was not solely due to the starting or stopping of anticonvulsants.

Subjects who Share the Characteristic AMY-HPC ICN Have Higher Trait Anxiety than Those who Lack It

Why do only $\sim 60\%$ of subjects have the mood-predictive β -AH ICN? We did not find any difference in AMY or HPC electrode placement between subjects who do or do not have this ICN

([Figure S1B](#)). Therefore, we tested whether differences in psychological traits might predict which subjects have this subnetwork. Although subjects used IMS to rate their current sense of well-being during their hospitalization, concurrent with iEEG data collection, they also filled out two questionnaires related to psychological traits during a pre-op visit 4.5 ± 2.5 months (mean \pm SD) before their hospital stay: BDI and BAI, which are designed to measure enduring trait levels of depression or anxiety, respectively ([Kabacoff et al., 1997](#); [Richter et al., 1998](#)). Intriguingly, subjects with the conserved β -AH ICN had significantly higher BAI scores on average than subjects who lacked it ([Figure 7](#)). In particular, we found that the majority of subjects who had this subnetwork were at or near the “moderate to severe” range of the BAI, whereas all those who lacked it were all in the “minimal” range ($*p < 0.02$). BDI scores also trended lower




Figure 7. Trait Anxiety Differences Separate Subjects with and without a β -AH Network

Beck Anxiety Inventory (BAI) and Beck Depression Inventory (BDI) scores for subjects with (black) and without (red) β -AH ICN ($^*p < 0.02$, two-sided t test).

in subjects who lacked a β -AH ICN, but this difference was not significant (Figure 7). Importantly, neither BAI nor BDI showed a relationship with a subject’s mean IMS, demonstrating a dissociation between these two “trait” measures and our IMS-based measure of “mood state” (Figure S4).

DISCUSSION

Our study identifies a characteristic AMY-HPC subnetwork related to mood variation over time. Specifically, within this subnetwork, increased temporal variance of AMY-HPC β -coherence correlates with worsening mood. This subnetwork is present in 62% (13 of 21) of study subjects, including all subjects with significant trait anxiety. Conversely, all subjects in whom this subnetwork is absent have minimal trait anxiety. The direction of causality for the last observation remains unclear. Specifically, the absence of this subnetwork may render certain individuals resilient to anxiety, leaving other factors to drive variation in their mood. Alternatively, this subnetwork may be present in those individuals, but their low anxiety may cause it to be fixed in a quiescent state, again leaving other factors to drive changes in their mood. It will be important for future studies to determine the direction of causality between network activity and mood; for example, by perturbing either the temporal variance of coherence between AMY and HPC or mood state.

Extracting Network-Behavior Relationships from High-Dimensional Datasets

Simply regressing a high-dimensional dataset against a behavioral measure runs the risk of detecting spurious correlations. We directly addressed this risk in three ways. First, we reduced the dimensionality of our datasets by identifying ICNs—temporally coordinated patterns of spectral coherence—and found those that were shared across subjects. Second, for each of the 6 subjects with the most IMS points, we used an elastic net to predict their IMS score from a sparse subset of ICNs to

mitigate the possibility of overfitting. Third, we looked for convergence across subjects. Although there is no *a priori* reason why the most mood-predictive ICNs should be similar across different subjects, we found that these clustered into the same ICN clique and frequency band—a degree of convergence highly unlikely to occur by chance ($p < 10^{-7}$). Moreover, using 2-fold cross-validation, we found that the quantitative relationship derived from one set of subjects (e.g., the 6 subjects used for the elastic net analysis) predicted 40%–50% of the variation in mood for a distinct, out-of-sample set of subjects (e.g., the other 7 subjects who also possessed the β -AH network).

This convergence validates the robustness of our findings and raises the question of whether specific aspects of our approach facilitated the discovery of this conserved, mood-related AMY-HPC subnetwork. Indeed, many studies have identified intrinsic brain networks, such as the default mode (Greicius et al., 2003; Raichle et al., 2001), central executive (Fox et al., 2006), and salience (Seeley et al., 2007) networks, within fMRI or PET data by searching for regions whose activity co-varied with that of other regions to define networks. Importantly, identification of these networks required that the average correlation between activity in the different regions was high. In contrast, our approach of performing ICA on the time series of coherence matrices identified brain regions connected by coherent interactions that are not consistently strong but, rather, tended to fluctuate in concert with coherent interactions between other regions. This is a key difference because interactions with a wide dynamic range may be better suited to uncover variation in equally dynamic behavioral measures such as mood. This concept is closely related to the recently proposed “chronnectome” for identifying re-occurring, time-varying connectivity networks from fMRI data (Calhoun et al., 2014).

Additional patterns of brain activity that encode subjective mood are likely to emerge from future studies that examine complementary iEEG activity features; for example, power features within single brain regions. Moreover, these granular types of features may identify idiosyncratic or subject-specific neural signatures for mood, in contrast to the more composite and dominant network-level features described here, which are likely to be better suited for identifying neural signatures for mood that are shared across individuals.

AMY-HPC in Mediating Mood Variability

AMY and HPC share strong anatomical connections, and both regions have been implicated in multiple aspects of emotion (Felix-Ortiz et al., 2013; Fournier and Duman, 2013; Zheng et al., 2017). In particular, it is thought that interactions between these two regions link emotion processing in the AMY to memory storage in the HPC, mediating the consolidation and retrieval of emotional memory (Girardeau et al., 2017; Goshen et al., 2011; Smith et al., 2006). Our findings support an intimate link between mood and emotional memory, where the encoded signal may reflect rumination on negative thoughts (Milazzo et al., 2016), context-dependent anxiety (Felix-Ortiz and Tye, 2014; Felix-Ortiz et al., 2013), or the recollection of negative experiences (Fanselow and Dong, 2010). Our findings may also relate to previous studies that have associated β -frequency activity with anxious thinking and active concentration (Güntekin and Başar,

2010) and, more recently, with depression (Clark et al., 2016). Computational models of β -frequency oscillations in networks relevant to Parkinson's disease have suggested that β -synchrony may reduce network flexibility (Brittain et al., 2014), potentially promoting the resident negative emotional state or memory over the formation of a new one.

Although our study did not explicitly study mood in the context of depression, we did find correlations between IMS score and the variance of AMY-HPC β -coherence in subjects with BDI values consistent with both depression and euthymia. This suggests that this AMY-HPC subnetwork may encode mood in both depressed and non-depressed individuals, although it does not address whether it plays a causal role in depression. This observation complements the large body of functional imaging work that has identified networks of brain regions associated with depression and antidepressant responses (Drysedale et al., 2017; Ferenczi et al., 2016; Hultman et al., 2016; 2018; Ressler and Mayberg, 2007). It will be particularly important for future studies to explore how the AMY-HPC interaction identified here may be altered in the setting of major depression and how this interaction may be regulated by regions such as the subgenual cingulate that are potential targets for deep brain stimulation in depression (Mayberg et al., 2005). This may lead to novel pathophysiological insights, whereas identifying mechanisms that can maintain this interaction within an optimal range may yield novel therapeutic strategies.

It is important to note here that we have looked for a neural relationship with IMS score as a whole. IMS is based on a range of questions, some of which might more classically be associated with DSM (Diagnostic and Statistical Manual of Mental Disorders) diagnoses of depression or anxiety. However, given that mood is a multifaceted emotional state, we decided to focus on this more composite measure here.

Relationship between Seizure Activity and AMY-HPC Activity

An important potential confounding factor of our study is the overlap between the mood-related AMY-HPC subnetwork and brain regions commonly implicated in seizures (refer to Table S1 for subjects' seizure foci). This currently represents a limitation of studies based on semi-chronic iEEG recordings, which, for the most part, are only available from patients with epilepsy. Nonetheless, iEEG recordings from patients with epilepsy have led to significant advances in our understanding of human neurophysiology (for example, in memory processing) despite the fact that memory-related networks also overlap with seizure foci (Burke et al., 2014; Kucewicz et al., 2017; Serruya et al., 2014). We looked for—and failed to find—a correlation between epileptiform activity and either IMS score or the variance of AMY-HPC β -coherence that might explain the observed correlation with mood. An added confounding factor is that patients undergoing intracranial monitoring for seizure localization often have their seizure medications tapered and eventually restarted over the course of their hospitalization, raising the possibility that some variation in mood may not be entirely “natural” but, rather, reflect medication-related mood changes. Although this would be interesting, it would not alter our key finding that the biomarker indicates variations in subjective mood. While we acknowledge

that some mood fluctuations may have been driven by medication changes, our analysis of a subset of subjects indicated that starting or stopping anticonvulsants did not generally account for the mood variation we observed. Understanding the temporal dynamics of medication-related changes in mood and neural biomarkers represents an important future direction but is beyond the scope of the current study, particularly because subjects received different combinations of medications. The fact that we observed a conserved neural biomarker-mood relationship across subjects with diverse medication regimens provides additional evidence that this relationship is not trivially mediated by medications.

Although the generalization of our findings to non-epilepsy populations remains to be determined, studying mood in this context is valid given the strong comorbidity between epilepsy and mood and anxiety disorders, in particular depression (Hermann et al., 2000). In fact, one of the most effective treatments for depression is electroconvulsive therapy (ECT), which was discovered based on observations from individuals with epilepsy (Faedda et al., 2010) but also has profound relevance for the non-epilepsy population (Read and Bentall, 2010).

In summary, our large-scale electrical recordings of the human limbic system have revealed an AMY-HPC subnetwork and a specific spatiotemporal neural signature that encode a large portion (~40%–50%) of the variation in mood over time. This subnetwork is shared across more than 60% of individuals and is consistently present in individuals with elevated anxiety. Our findings contribute to a deeper understanding of the neural encoding of mood and anxiety and reveal a biomarker that might be useful for diagnosis and treatment of mood and anxiety disorders, in particular innovative treatments using closed-loop deep brain stimulation.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
 - Participants
 - Inclusion criteria
- METHOD DETAILS
 - Neural recordings
 - Cortical surface extraction and electrode visualization
 - Psychological testing
 - Region of interest abbreviations
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Preprocessing of electrophysiological recordings
 - Coherence matrices
 - Power spectra
 - Intrinsic coherence networks (ICNs)
 - Common network analysis
 - Regression analysis
 - Clique-convergence probability
 - Interictal activity analysis
 - Medication analysis

● DATA AND SOFTWARE AVAILABILITY

- Data availability
- Code availability

SUPPLEMENTAL INFORMATION

Supplemental Information includes seven figures and three tables and can be found with this article online at <https://doi.org/10.1016/j.cell.2018.10.005>.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive comments. This research was partially funded by the Defense Advanced Research Projects Agency (DARPA) under cooperative agreement W911NF-14-2-0043 issued by the Army Research Office contracting office in support of DARPA's SUBNETS program. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the United States government. This work was also supported by a Swartz Foundation postdoctoral fellowship 127713C (to L.A.K.).

AUTHOR CONTRIBUTIONS

All authors helped to conceive and design parts of the study. E.F.C. conceived and designed the overall approach for measuring subjective mood while recording iEEG signals from human subjects. E.F.C. also conceived and designed the approach for obtaining iEEG recordings from multiple distributed limbic regions, performed the implantations, and collected the iEEG data. M.N. and T.M.V.V. designed and created the application to measure IMS. M.B.L. collected IMS data. V.R.R. performed neuropsychiatric evaluation and epilepsy monitoring of subjects and confirmed electrode localizations. L.A.K., F.J.L., and V.S.S. conceived and designed the analysis framework, including the methods for identifying ICNs. L.A.K. performed all analyses. L.A.K. and V.S.S. wrote the paper. H.E.D., E.F.C., and V.S.S. oversaw and guided all aspects of the project.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: May 10, 2018

Revised: June 26, 2018

Accepted: September 28, 2018

Published: November 8, 2018

REFERENCES

Ashburner, J., and Friston, K.J. (1997). Spatial transformation of images (Human Brain Function).

Bickart, K.C., Dickerson, B.C., and Barrett, L.F. (2014). The amygdala as a hub in brain networks that support social life. *Neuropsychologia* 63, 235–248.

Bowyer, S.M. (2016). Coherence a measure of the brain networks: past and present. *Neuropsychiatric Electrophysiology* 2, 1.

Brittain, J.S., Sharott, A., and Brown, P. (2014). The highs and lows of beta activity in cortico-basal ganglia loops. *Eur. J. Neurosci.* 39, 1951–1959.

Burke, J.F., Long, N.M., Zaghoul, K.A., Sharan, A.D., Sperling, M.R., and Kahana, M.J. (2014). Human intracranial high-frequency activity maps episodic memory formation in space and time. *Neuroimage* 85, 834–843.

Calhoun, V.D., Eichele, T., and Pearlson, G. (2009). Functional brain networks in schizophrenia: a review. *Front. Hum. Neurosci.* 3, 17.

Calhoun, V.D., Miller, R., Pearlson, G., and Adalı, T. (2014). The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* 84, 262–274.

Clark, D.L., Brown, E.C., Ramasubbu, R., and Kiss, Z.H.T. (2016). Intrinsic Local Beta Oscillations in the Subgenual Cingulate Relate to Depressive

Symptoms in Treatment-Resistant Depression. *Biol. Psychiatry* 80, e93–e94.

Drysdale, A.T., Grosenick, L., Downar, J., Dunlop, K., Mansouri, F., Meng, Y., Fetcho, R.N., Zebley, B., Oathes, D.J., Etkin, A., et al. (2017). Resting-state connectivity biomarkers define neurophysiological subtypes of depression. *Nat. Med.* 23, 28–38.

Faedda, G.L., Becker, I., Baroni, A., Tondo, L., Aspland, E., and Koukopoulos, A. (2010). The origins of electroconvulsive therapy: Prof. Bini's first report on ECT. *J. Affect. Disord.* 120, 12–15.

Fanselow, M.S., and Dong, H.-W. (2010). Are the dorsal and ventral hippocampus functionally distinct structures? *Neuron* 65, 7–19.

Felix-Ortiz, A.C., and Tye, K.M. (2014). Amygdala inputs to the ventral hippocampus bidirectionally modulate social behavior. *J. Neurosci.* 34, 586–595.

Felix-Ortiz, A.C., Beyeler, A., Seo, C., Leppla, C.A., Wildes, C.P., and Tye, K.M. (2013). BLA to vHPC inputs modulate anxiety-related behaviors. *Neuron* 79, 658–664.

Ferenczi, E.A., Zalocusky, K.A., Liston, C., Grosenick, L., Warden, M.R., Amatya, D., Katovich, K., Mehta, H., Patenaude, B., Ramakrishnan, C., et al. (2016). Prefrontal cortical regulation of brainwide circuit dynamics and reward-related behavior. *Science* 351, aac9698.

Fischl, B. (2012). FreeSurfer. *Neuroimage* 62, 774–781.

Fischl, B., Sereno, M.I., Tootell, R.B., and Dale, A.M. (1999). High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.* 8, 272–284.

Fournier, N.M., and Duman, R.S. (2013). Illuminating hippocampal control of fear memory and anxiety. *Neuron* 77, 803–806.

Fox, M.D., Corbetta, M., Snyder, A.Z., Vincent, J.L., and Raichle, M.E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proc. Natl. Acad. Sci. USA* 103, 10046–10051.

Girardeau, G., Inema, I., and Buzsáki, G. (2017). Reactivations of emotional memory in the hippocampus-amygdala system during sleep. *Nat. Neurosci.* 20, 1634–1642.

Goshen, I., Brodsky, M., Prakash, R., Wallace, J., Gradinaru, V., Ramakrishnan, C., and Deisseroth, K. (2011). Dynamics of retrieval strategies for remote memories. *Cell* 147, 678–689.

Greicius, M.D., Krasnow, B., Reiss, A.L., and Menon, V. (2003). Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proc. Natl. Acad. Sci. USA* 100, 253–258.

Greicius, M.D., Srivastava, G., Reiss, A.L., and Menon, V. (2004). Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *Proc. Natl. Acad. Sci. USA* 101, 4637–4642.

Greicius, M.D., Flores, B.H., Menon, V., Glover, G.H., Solvason, H.B., Kenna, H., Reiss, A.L., and Schlaggar, A.F. (2007). Resting-state functional connectivity in major depression: abnormally increased contributions from subgenual cingulate cortex and thalamus. *Biol. Psychiatry* 62, 429–437.

Güntekin, B., and Başar, E. (2010). Event-related beta oscillations are affected by emotional eliciting stimuli. *Neurosci. Lett.* 483, 173–178.

Hermann, B.P., Seidenberg, M., and Bell, B. (2000). Psychiatric comorbidity in chronic epilepsy: identification, consequences, and treatment of major depression. *Epilepsia* 41 (Suppl 2), S31–S41.

Hultman, R., Mague, S.D., Li, Q., Katz, B.M., Michel, N., Lin, L., Wang, J., David, L.K., Blount, C., Chandry, R., et al. (2016). Dysregulation of Prefrontal Cortex-Mediated Slow-Evolving Limbic Dynamics Drives Stress-Induced Emotional Pathology. *Neuron* 91, 439–452.

Hultman, R., Ulrich, K., Sachs, B.D., Blount, C., Carlson, D.E., Ndubuizu, N., Bagot, R.C., Parise, E.M., Vu, M.T., Gallagher, N.M., et al. (2018). Brain-wide Electrical Spatiotemporal Dynamics Encode Depression Vulnerability. *Cell* 173, 166–180.e14.

Hyvärinen, A., and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural Netw.* 13, 411–430.

- Jerbi, K., Freyermuth, S., Dalal, S., Kahane, P., Bertrand, O., Berthoz, A., and Lachaux, J.-P. (2009). Saccade related gamma-band activity in intracerebral EEG: dissociating neural from ocular muscle activity. *Brain Topogr.* *22*, 18–23.
- Kabacoff, R.I., Segal, D.L., Hersen, M., and Van Hasselt, V.B. (1997). Psychometric properties and diagnostic utility of the Beck Anxiety Inventory and the State-Trait Anxiety Inventory with older adult psychiatric outpatients. *J. Anxiety Disord.* *11*, 33–47.
- Kucewicz, M.T., Berry, B.M., Kremen, V., Brinkmann, B.H., Sperling, M.R., Jobst, B.C., Gross, R.E., Lega, B., Sheth, S.A., Stein, J.M., et al. (2017). Dissecting gamma frequency activity during human memory processing. *Brain* *140*, 1337–1350.
- LeDoux, J.E. (2003). Emotion Circuits in the Brain. *Annu. Rev. Neurosci.* *23*, 155–184.
- Lopes-dos-Santos, V., Ribeiro, S., and Tort, A.B.L. (2013). Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* *220*, 149–166.
- Mayberg, H.S., Brannan, S.K., Tekell, J.L., Silva, J.A., Mahurin, R.K., McGinnis, S., and Jerabek, P.A. (2000). Regional metabolic effects of fluoxetine in major depression: serial changes and relationship to clinical response. *Biol. Psychiatry* *48*, 830–843.
- Mayberg, H.S., Lozano, A.M., Voon, V., McNeely, H.E., Seminowicz, D., Hamani, C., Schwab, J.M., and Kennedy, S.H. (2005). Deep brain stimulation for treatment-resistant depression. *Neuron* *45*, 651–660.
- Meliza, C.D., and Margoliash, D. (2012). Emergence of selectivity and tolerance in the avian auditory cortex. *J. Neurosci.* *32*, 15158–15168.
- Milazzo, A.-C., Ng, B., Jiang, H., Shirer, W., Varoquaux, G., Poline, J.B., Thirion, B., and Greicius, M.D. (2016). Identification of Mood-Relevant Brain Connections Using a Continuous, Subject-Driven Rumination Paradigm. *Cereb. Cortex* *26*, 933–942.
- Nahum, M., Van Vleet, T.M., Sohal, V.S., Mirzabekov, J.J., Rao, V.R., Wallace, D.L., Lee, M.B., Dawes, H., Stark-Inbar, A., Jordan, J.T., et al. (2017). Immediate Mood Scaler: Tracking Symptoms of Depression and Anxiety Using a Novel Mobile Mood Scale. *JMIR Mhealth Uhealth* *5*, e44.
- Phan, K.L., Wager, T., Taylor, S.F., and Liberzon, I. (2002). Functional neuroanatomy of emotion: a meta-analysis of emotion activation studies in PET and fMRI. *Neuroimage* *16*, 331–348.
- Phelps, E.A., and LeDoux, J.E. (2005). Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* *48*, 175–187.
- Phillips, M.L., Drevets, W.C., Rauch, S.L., and Lane, R. (2003). Neurobiology of emotion perception I: The neural basis of normal emotion perception. *Biol. Psychiatry* *54*, 504–514.
- Price, J.L., and Drevets, W.C. (2010). Neurocircuitry of mood disorders. *Neuropsychopharmacology* *35*, 192–216.
- Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., and Shulman, G.L. (2001). A default mode of brain function. *Proc. Natl. Acad. Sci. USA* *98*, 676–682.
- Read, J., and Bentall, R. (2010). The effectiveness of electroconvulsive therapy: a literature review. *Epidemiol. Psychiatr. Soc.* *19*, 333–347.
- Ressler, K.J., and Mayberg, H.S. (2007). Targeting abnormal neural circuits in mood and anxiety disorders: from the laboratory to the clinic. *Nat. Neurosci.* *10*, 1116–1124.
- Richter, P., Werner, J., Heerlein, A., Kraus, A., and Sauer, H. (1998). On the validity of the Beck Depression Inventory. A review. *Psychopathology* *31*, 160–168.
- Schnitzler, A., and Gross, J. (2005). Normal and pathological oscillatory communication in the brain. *Nat. Rev. Neurosci.* *6*, 285–296.
- Schult, D.A., and Swart, P. (2008). Exploring network structure, dynamics, and function using NetworkX. In *Proceedings of the 7th Python in Science conference (SciPy 2008)*, G. Varoquaux, T. Vaught, and J. Millman, eds., pp. 11–15.
- Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., and Greicius, M.D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *J. Neurosci.* *27*, 2349–2356.
- Serruya, M.D., Sederberg, P.B., and Kahana, M.J. (2014). Power shifts track serial position and modulate encoding in human episodic memory. *Cereb. Cortex* *24*, 403–413.
- Smith, A.P.R., Stephan, K.E., Rugg, M.D., and Dolan, R.J. (2006). Task and content modulate amygdala-hippocampal connectivity in emotional retrieval. *Neuron* *49*, 631–638.
- Srinath, R., and Ray, S. (2014). Effect of amplitude correlations on coherence in the local field potential. *J. Neurophysiol.* *112*, 741–751.
- Touroutoglou, A., Hollenbeck, M., Dickerson, B.C., and Feldman Barrett, L. (2012). Dissociable large-scale networks anchored in the right anterior insula subservice affective experience and attention. *Neuroimage* *60*, 1947–1958.
- Treynor, W., Gonzalez, R., and Nolen-Hoeksema, S. (2003). Rumination Reconsidered: A Psychometric Analysis. *Cognit. Ther. Res.* *27*, 247–259.
- Zheng, J., Anderson, K.L., Leal, S.L., Shestyuk, A., Gulsen, G., Mnatsakanyan, L., Vadera, S., Hsu, F.P.K., Yassa, M.A., Knight, R.T., and Lin, J.J. (2017). Amygdala-hippocampal dynamics during salient information processing. *Nat. Commun.* *8*, 14413.
- Zou, H., and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Series B Stat. Methodol.* *67*, 301–320.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited Data		
Raw and analyzed iEEG data	This paper	N/A
IMS data	This paper	N/A
Software and Algorithms		
Python packages: cohere_pairs, fastICA, ElasticNetCV	Python libraries: matplotlib, sklearn, scipy, numpy	https://www.anaconda.com/download/#macos
Pygraphviz package	https://pygraphviz.github.io	https://github.com/pygraphviz/pygraphviz
Networkx package	Schult and Swart, 2008	https://networkx.github.io
Quickspikes package	Meliza and Margoliash, 2012	https://github.com/melizalab/quickspikes
Custom code and algorithms	This paper	N/A
Other		
Human participants with treatment-resistant epilepsy, undergoing surgical treatment at the UCSF Medical Center	This paper	N/A
Intracranial electrodes	Ad-Tech subdural strip and SEEG depth electrodes	https://adtechmedical.com/about-us
Neural recordings and amplifier	Natus XLTex EEG clinical recording system and Quantum LTM amplifier	https://neuro.natus.com/products-services/natus-quantum-ltm-amplifier

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Vikaas S. Sohal (vikaas.sohal@ucsf.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Participants

21 human subjects (16 male, 5 female, aged 33 ± 10 years) with treatment-resistant epilepsy were surgically implanted with semi-chronic intracranial electrodes (for sites included in this study: Ad-Tech 4-contact strip or depth, 10mm center-to-center spacing, 2.3mm exposed diameter; Ad-Tech 10-contact depth, 6 or 5mm center-to-center spacing) for the clinical purpose of seizure localization. All procedures were approved by the University of California, San Francisco Institutional Review Board. All subjects gave written informed consent to participate in the study before surgery. 12 subjects had left hemisphere implantations, 9 had right hemisphere implantations. Subjects' clinical information is provided in [Table S1](#).

Results of the study (i.e., the presence of a mood-predictive AMY-HPC subnetwork) generalized to both sexes. Specifically, the mood-predictive AMY-HPC subnetwork was present in 4/5 female and 9/16 male subjects. There is not a significant difference between these proportions ($p = 0.61$ by Fisher's exact test), however, additional female subjects would be necessary to examine possible sex differences more rigorously.

Inclusion criteria

Subjects were included in our study based on the following criteria: (i) AMY electrode implantation (the dominant brain region for emotion) (ii) minimum coverage of 4 brain regions (to assess large-scale, cross-regional subnetworks) (iii) unilateral electrode implantation, and (iv) minimum of 3 IMS data points. Subjects were included in elastic net regression analysis based on the following criteria: (i) minimum of 10 IMS data points (ii) minimum IMS coefficient of variance of 5%.

METHOD DETAILS

Neural recordings

Electrophysiological recordings were acquired at a sampling rate of either 512 Hz or 1.02 kHz using the Nicolet Natus (EC77, EC79, EC80) or Natus XLTex (remaining subjects) EEG clinical recording systems, with either the EEG128FS amplifier and Neuroworks 8.0 software (EC82, EC84) or Quantum amplifier and Neuroworks 8.1 software (EC108-EC136). Recordings were performed

continuously over a 7–14 day hospitalization period. Electrode coverage varied for each subject, depending on their epileptic pathology. To maximize generalizability across subjects, we focused our analysis on depth electrodes located in limbic structures that were most common across subjects: subtemporal cortex (STC, 4-contact strip in anterior, middle and/or posterior STC), amygdala (AMY, single 4- or 10- contact depth), hippocampus (HPC, 4- or 10-contact depth in ventral HPC), orbitofrontal cortex (OFC, 4-contact strip in anterior and/or posterior OFC), cingulate cortex (CIN, 4- or 10-contact depth in inferior and/or superior CIN) and insular cortex (INS, single 4- or 10-contact depth). Electrode locations were confirmed by visual inspection of brain image reconstructions.

All subjects primarily remained laying down in the hospital bed throughout the duration of their hospitalization, minimizing contamination of the recordings by motor artifacts related to locomotion or ambulation. Moreover, there should be minimal contamination from eye movement artifacts, as there is evidence that these do not contaminate intracranial EEG from depth electrodes within the brain parenchyma (Jerbi et al., 2009).

Cortical surface extraction and electrode visualization

Electrodes on each subject's brain were localized by co-registering the pre-operative T1 MRI with a post-operative CT scan containing the electrode locations, using a normalized mutual information routine in the Statistical Parametric Mapping software SPM12 (Ashburner and Friston, 1997). Pial surface reconstructions were created using Freesurfer (Fischl, 2012). For visualization of electrode coordinates in MNI-coordinates, a nonlinear surface registration was performed using a spherical sulcal-based alignment in Freesurfer, aligning to the cvs avg35 in MNI152 template (Fischl et al., 1999). Centroid locations for AMY and HPC electrodes were computed in MNI coordinate space using the center of mass across all electrodes anatomically verified to be located within either brain structure.

Psychological testing

Subjects' psychological state was evaluated up to four times daily during their hospitalization period prompted by a research assistant, using a tablet-based, custom-designed questionnaire called the Immediate Mood Scaler (Nahum et al., 2017) (IMS). This consisted of a set of either 5 (EC84), 23 (EC77, EC79, EC80) or 24 (remaining subjects) questions, in which the subject self-reported their instantaneous mood on a scale between -3 and $+3$ (Table S2). The final composite score provided a single value between ± 15 (5-question), ± 69 (23-question) or ± 72 (24-question), representative of the subject's overall psychological state at that time point. Higher values correspond to more positive instantaneous mood state. IMS questions were based on the Patient Health Questionnaire-9 (PHQ-9), Generalized Anxiety Disorder-7 (GAD-7), and Rumination scale (Nahum et al., 2017). Subjects were also administered PHQ-9, GAD-7 and Rumination scale on the first day of their hospitalization period. IMS data points taken within 1.5hr of the previous data point were considered to not be independent and excluded from analysis (12/160 data points total; inter-IMS-interval across all subjects = 9.2 ± 7.0 hr). When comparing across subjects, IMS scores were first z-scored relative to a subject's mean and standard deviation.

In the months prior to hospitalization, subjects were administered the Beck Depression Inventory (BDI) and Beck Anxiety Inventory (BAI) to assess trait levels of depression and anxiety.

Region of interest abbreviations

Subtemporal cortex (STC), amygdala (AMY), hippocampus (HPC), orbitofrontal cortex (OFC), cingulate cortex (CIN), insular cortex (INS), anterior (ant, a), middle (mid, m), posterior (pos, p), ventral (ven, v), inferior (inf, i), superior (sup, s).

QUANTIFICATION AND STATISTICAL ANALYSIS

Preprocessing of electrophysiological recordings

Voltage signals from each electrode were band-passed filtered between 0.5–256 Hz and downsampled to 512 Hz using an 8th order chebyshev type I filter. Signals were notch-filtered at 60, 120, 180 and 240 Hz with 4 Hz bandwidth using a 5th order butterworth filter, to reduce line-noise-related artifacts. Signals were then re-referenced to the common average across channels sharing the same lead. Noise and artifact occurring over several leads, for example, due to disturbances to cables or connectors, was removed from signals recorded on the Nicolet Natus recording system (EC77, EC79 and EC80) using an Independent Components Analysis (ICA) approach. Briefly, signals were divided into ~ 1 hr time segments and decomposed with ICA using Python's *fastICA* algorithm, which maximizes nongaussianity of prewhitened components using an iterative, fixed-point rotation scheme. This separated the signals into a combination of artifact and neural independent components (ICs). The artifact components showed large voltage deflections and were distinguishable from the neural components. A subset of ICs ($\sim 10\%$ per dataset, from a total of over 6000 ICs per dataset) were manually labeled as either artifact or neural. We then used the labeled dataset to train a logistic classifier to distinguish between artifact versus neural ICs based on two features: power spectra and amplitude distribution. ICs identified as artifact using the classifier were subtracted from the original signals.

Coherence matrices

Voltage traces from electrodes located in STC, AMY, HPC, OFC, CIN and INS were split into contiguous 10 s segments. Signal coherence, $C_{xy}(f)$, is a measure of the normalized cross-spectral density and is given by:

$$C_{xy}(f) = \frac{|G_{xy}(f)|^2}{G_{xx}(f)G_{yy}(f)}$$

where $G_{xy}(f)$ corresponds to the cross-spectral density between two signals x and y , and $G_{xx}(f)$ and $G_{yy}(f)$ correspond to the auto-spectral densities of x and y , respectively. Signal coherence was calculated between all pairs of electrodes for each 10 s segment using Python's *cohere_pairs* function, which computes power spectral density using Welch's method with a non-overlapping Hanning window. We repeated the calculation using phase randomized surrogate signals (i.e., signals with the same power spectra as the original signals but reconstructed with randomized phases) and subtracted these values from the coherence of the original signal. This step subtracted out a noise floor resulting from co-variations in amplitude that are not a result of true phase coherence, which can significantly bias coherence measure (Srinath and Ray, 2014). We constructed 4 time-series of coherence matrices by averaging across four frequency bands: theta θ , [4-8 Hz]; alpha, α [8-13 Hz]; beta, β [13-30 Hz]; and, gamma γ [30-70 Hz].

Power spectra

Voltage traces from electrodes located in AMY and HPC were z-scored relative to the mean and standard deviation of activity across the entire recording duration, and hence split into contiguous 10 s segments. The power spectral density of each 10 s segment was calculated using Python's *welch* function in the SciPy library, which computes power spectral density using Welch's method with a non-overlapping Hanning window. Time series of mean power between 13 and 30 Hz (β -frequency) were obtained for each electrode and averaged across all electrodes in either AMY or HPC to obtain the mean β -power spectra for either brain region.

Intrinsic coherence networks (ICNs)

ICNs were identified by deriving statistically significant, independent components of coherence matrices. First, principal component analysis (PCA) was carried out on each set of coherence matrices for dimensionality reduction and orthogonalization. PCA was performed on each $m \times n$ matrix M , with rows corresponding to m electrode-electrode coherence pairs and columns corresponding to n time points. To obtain an estimate of number of significant components, we used a statistical threshold set by the Marchenko-Pastur Law, which gives the distribution of eigenvalues expected by chance (Lopes-dos-Santos et al., 2013). Principal components (PCs) with eigenvalues above the Marchenko-Pastur eigenvalue threshold, λ_{MP} , represent statistically significant sources of variation within the dataset, where:

$$\lambda_{MP} = \sigma^2 \left(1 + \sqrt{1/q} \right)^2$$

for a matrix of size $q = n/m$ and variance of elements σ^2 . PCs with eigenvalues greater than λ_{MP} were extracted for each dataset. Next, independent components analysis (ICA) was carried out on the significant PCs to separate the signal mixtures into independent sources. We used Python's *fastICA* algorithm, which maximizes nongaussianity, measured by the approximation of negentropy, of prewhitened components using an iterative, fixed-point rotation scheme (Hyvärinen and Oja, 2000). Independent components (ICs) identified in this manner correspond to subnetworks of brain regions in which coherence tends to fluctuate in a coordinated manner, that capture significant sources of variation in the dataset, and that are maximally independent from one another at an instant in time. ICNs were visualized using chord diagrams with electrodes mapped onto nodes. Edges between nodes were found by thresholding the ICN loading above (below) the 98th (2nd) percentile. This threshold was chosen purely for illustrative purposes and has no bearing on any quantitative analyses or conclusions.

Common network analysis

A common network template was defined to consist of the four most medial electrodes in each of the ten following regions: anterior STC, middle STC, posterior STC, AMY, HPC, anterior OFC, posterior OFC, inferior CIN, superior CIN, and INS. I.e., if a subject had electrodes in any of these brain regions, the four most medial electrodes were selected for the common network analysis, while the remainder were excluded. If a subject did not have electrodes in a given brain region, the corresponding region was left empty (i.e., matrix elements corresponding to absent brain regions were set to have a value of zero). These absent electrodes are represented by hollow nodes in the ICN connectivity diagrams.

A similarity matrix to determine the spatial similarity of all ICN-ICN combinations (across all subjects and all frequency bands) was constructed by calculating the Pearson correlation coefficient (across electrode space) between each common network-projected ICN pairs. Network similarity was analyzed and visualized graphically using Python's *networkx* (Schult and Swart, 2008) and *pygraphviz* packages, in which nodes correspond to ICNs and edges link together similar ICNs. Briefly, an adjacency matrix was constructed by thresholding the absolute magnitude of the similarity matrix above 0.55. Core nodes were identified as those that were connected to at least 7 other nodes. We present a justification for this choice of cross-correlation threshold in Figure S3. These thresholds were chosen to be at a level that winnows the number of connections to yield a sparse graph that is suitable for

analysis. To avoid all nodes either being retained (red/orange in Figure S3A) or excluded (blue/purple in Figure S3A) requires a cross-correlation threshold between ~ 0.4 - 0.6 . The green band in Figure S3A corresponds to $\sim 40\%$ of nodes being included for analysis, thus balancing the inclusion/exclusion of nodes and presumably yielding a sparse graph that is appropriate for analysis. Choosing a cross-correlation threshold of 0.55 and core node threshold of 7 provided an appropriate trade-off between clusters that are highly overlapping and interconnected (Figure S3C, left) and clusters that start to break off and become independent islands (Figure S3C, right). Nonetheless, even though the topological graph structure changed as these thresholds are varied, the same overall clique structure was largely preserved over this range of thresholds (Figure S3C). Cliques of core nodes were identified by finding 5-clique clusters of nodes, in which each ICN node was connected to each other ICN node in that clique. To compare clique structure to random structure, we constructed a randomly shuffled adjacency matrix.

Regression analysis

We used an elastic net approach (Zou and Hastie, 2005) to identify linear models that predict IMS score based on activity within multiple ICNs. This is a regularized linear regression model that determines the coefficients, β , which map a feature vector X onto observations y , by minimizing the following loss function:

$$\hat{\beta} = \arg \min \left(|y - X\beta|^2 + \lambda_2 |\beta|^2 + \lambda_1 |\beta| \right)$$

where, λ_1 and λ_2 are L_1 and L_2 regularization parameters, respectively. In this study, y corresponds to IMS score and X corresponds to a representation of ICN activity. The latter was determined by first projecting the time series of coherence matrices onto each ICN, separately for each frequency band and subject. The variance of the projection was then measured using a 1-minute sliding window, in order to capture transient peaks observed in the temporal structure of ICN projections. Feature vectors X were then constructed using all frequency band ICNs for each subject, by taking the mean ICN projection variance over a given time window around each IMS time point in vector y . An elastic net model was fit using Python's *ElasticNetCV* class in the scikit-learn library. The best model (the ratio of regularization parameters λ_1 and λ_2 that minimizes the loss function) was selected using 3-fold cross-validation. A learning rate of $\alpha = 0.1$ was used. After testing 6 different time windows (5, 10, 15, 20, 25, or 30 min), we found that only when we used a time window of 20 min did the elastic net yield a model that could predict IMS (i.e., non-zero R^2 regression score) in all 6 subjects for whom we performed this analysis. Importantly, we did not choose the time window separately for each subject based on the 'best' analysis results (e.g., highest R^2). Rather, we selected a single time window that yields a non-zero R^2 for all subjects. Furthermore, the choice of this time window is completely agnostic to the details of the model, i.e., which networks are used to predict mood. Thus the temporal window optimization does not increase the likelihood of observing convergence onto a specific mood-predictive network across subjects. As a result no correction for multiple comparisons is warranted. In other words, even though different windows may have been optimal (i.e., yielding higher R^2) for individual subjects, this 20-minute time window was optimal at the population level and used for all subsequent analyses. R^2 regression scores were compared to those from models fit to randomly shuffled ICN time series, computed using Python's *random.shuffle* function on ICN projections, and repeated 100 times for each subject.

Clique-convergence probability

We estimated the probability that 6/6 subjects would have the most IMS-predictive ICN in a single ICN clique, with 5/6 of these in the same frequency band by summing over all 9 cliques, c , and all 4 frequency bands, f , as follows:

$$p \approx \sum_c \sum_f C_6^6 \left(\frac{n_c}{n_{tot}} \right)^6 C_5^6 \left(\frac{n_{c,f}}{n_c} \right)^5$$

where n_{tot} is the total number of ICNs for the 6 subjects, n_c is the number of ICNs for the 6 subjects in clique c , $n_{c,f}$ is the number of ICNs for the 6 subjects in clique c frequency band f , and C_k^n corresponds to the binomial coefficient for choosing k elements from a set of n elements.

Interictal activity analysis

The iEEG recordings in this study consist of large-scale, semi-chronic recordings (~ 40 - 70 electrodes recorded continuously over the course of ~ 7 - 10 days). Due to the large scale of these recordings, rather than excluding electrodes or epochs with epileptiform activity, we opted to retain all data and to perform the appropriate post hoc analyses to control for possible mood effects of interictal activity. iEEG signals on AMY and HPC electrodes were extracted in the 20-minutes around each IMS score time point and normalized by z-score. Interictal discharges were identified as transient spikes that crossed a threshold of -4 standard deviations from the mean, going away from zero, using Python's *quickspikes* package (Meliza and Margoliash, 2012). Discharge rate was estimated by calculating the number of spikes per unit time.

Medication analysis

Potential effects of anti-epileptic drugs (AEDs) on neural activity and mood were assessed in the three subjects with the strongest mood-subnetwork relationship: EC79, EC82 and EC108. EC79 received levetiracetam (LEV) for 4/6 days at 3000mg/day, and was

off AEDs for 2/6 days; EC82 was off AEDs for 6/6 days; EC108 received LEV for 5/5 days at 2000mg/day. Medication was administered twice daily at 9AM and 9PM. The mean temporal variance of AMY-HPC β -coherence and mean IMS score were calculated over each 4-hour period following medication administration (or equivalent 4-hour periods when subjects did not receive medication). The mean temporal variance values were z-scored within each subject, and the mean IMS score was z-scored based on the full distribution of IMS points for each subject.

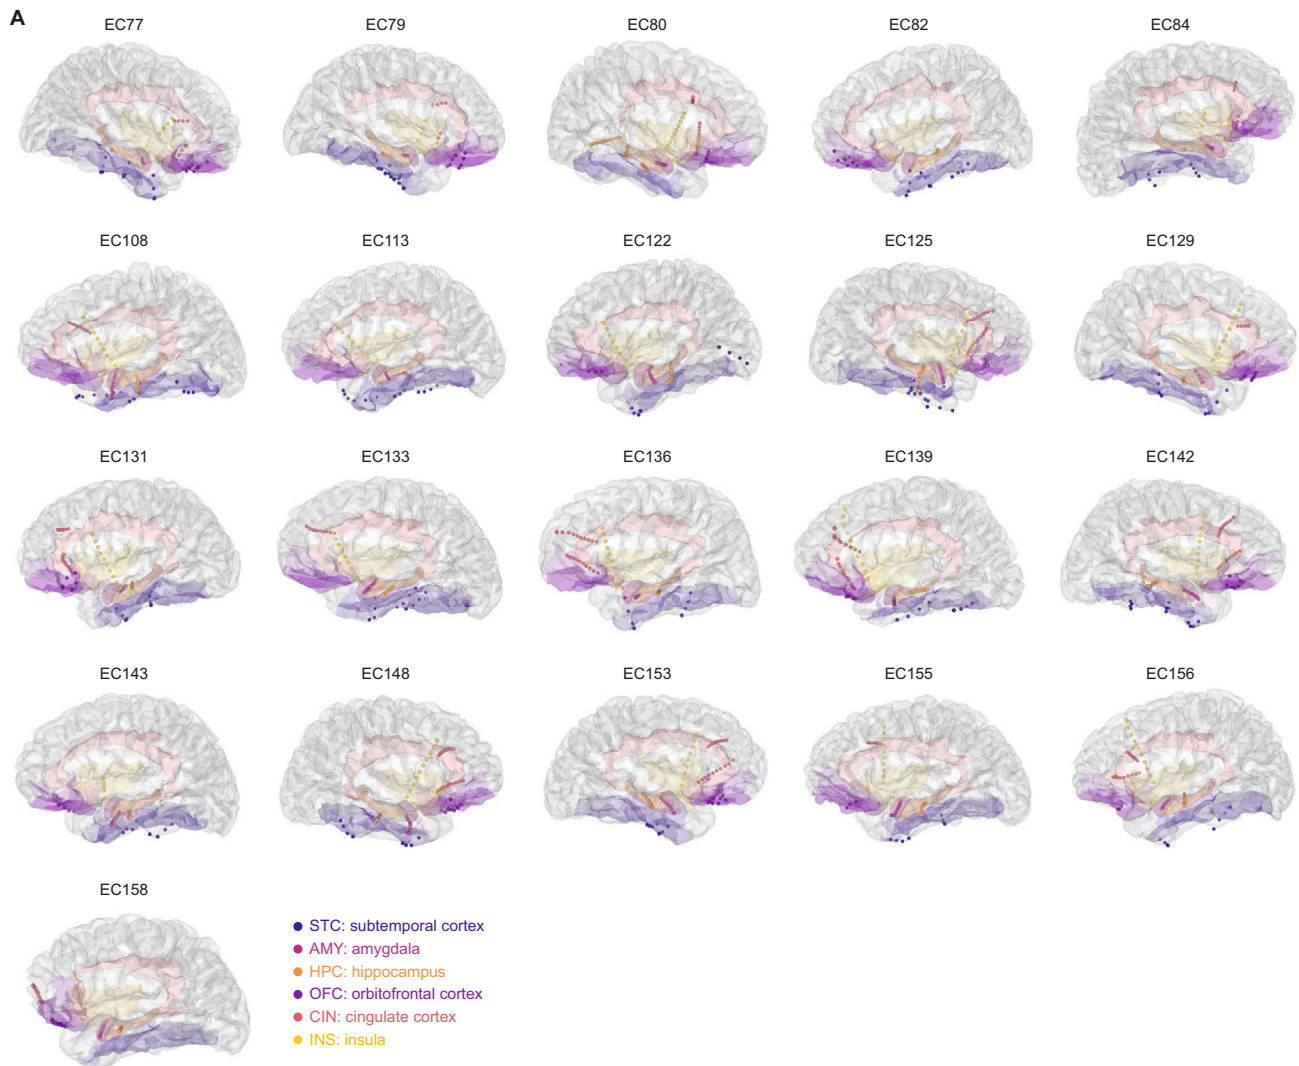
DATA AND SOFTWARE AVAILABILITY

Data availability

Data used to generate the findings of this study will be freely available upon request (subject to participant consent) to the Lead Contact.

Code availability

Custom computer code used to generate the findings of this study will be made available upon request to the Lead Contact.



B AMY and HPC electrode centroids (MNI coordinates)


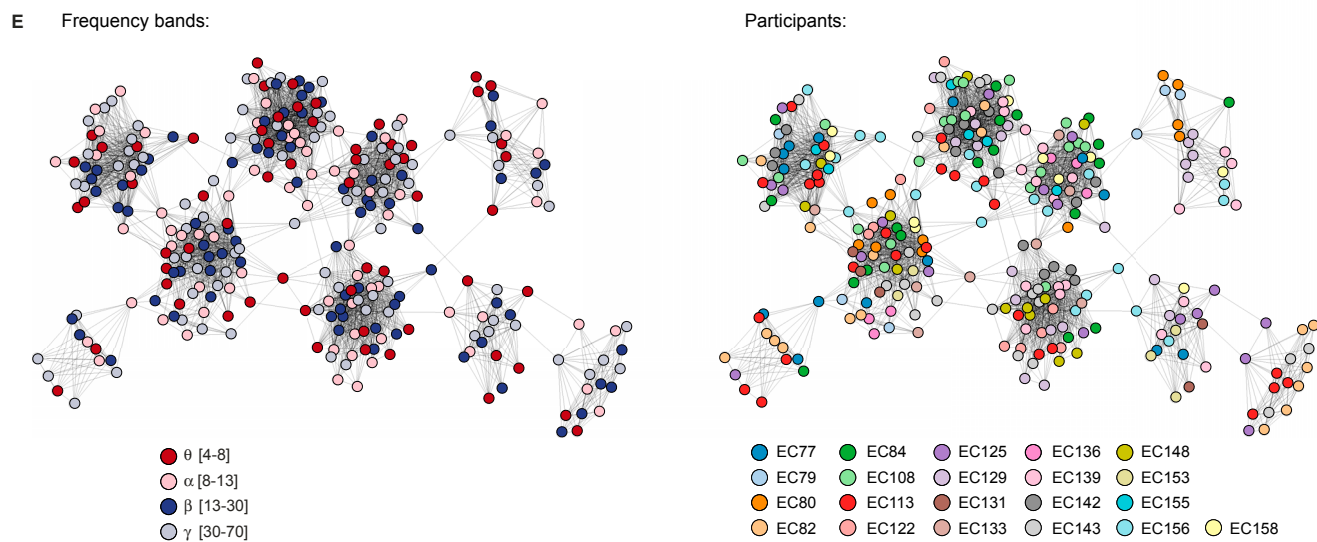

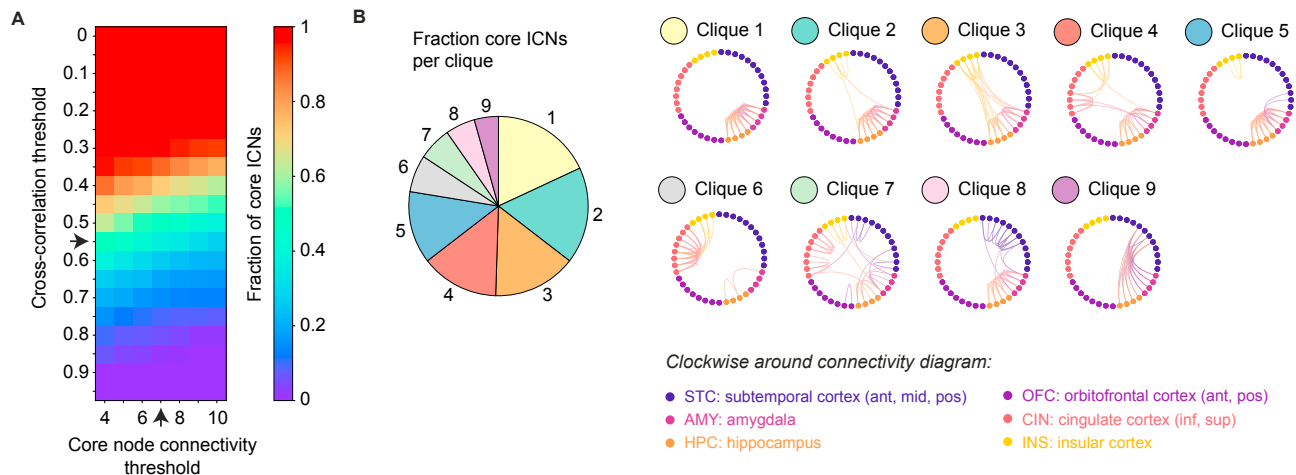


Figure S1. Electrode Locations, Related to Figure 1

(A) Sagittal views of electrode locations for each individual subject in native coordinates. STC: subtemporal cortex (blue); AMY: amygdala (magenta); HPC: hippocampus (orange); OFC: orbitofrontal cortex (purple); CIN: cingulate cortex (red); INS: insular cortex (yellow).

(B) Sagittal view of the centroid locations of AMY (magenta) and HPC (orange) electrodes superimposed onto reconstructed brain in MNI coordinate space for subjects with and without β -AH ICN. The overlapping centroids within each brain region for each group indicate that no marked differences in electrode locations between these two groups exists.



(legend on next page)

Figure S3. Clique Structure Is Preserved for Different Threshold Values, Related to Figure 2

(A) Fraction of “core ICNs” (#core ICNs/#all ICNs) over range of cross-correlation and core node connectivity thresholds.

(B) Left: Proportion of ICNs per clique (#ICNs in clique x /#core ICNs, where $x = 1-9$) for topological ICN graph in Figure 2 (cross-correlation threshold = 0.55, core node connectivity threshold = 7, clique connectivity threshold = 5). Right: Corresponding chord diagrams of average spatial connectivity patterns of all ICNs represented in a given clique (as in Figure 1D).

(C) Topological graphs (as in Figure 2) for two different combinations of threshold values: cross-correlation threshold = 0.5 and core node connectivity threshold = 8 (left); cross-correlation threshold = 0.6 and core node connectivity threshold = 6 (right). Nodes are color-coded according to clique assignment as in (B); white nodes correspond to unassigned ICNs.

(D) Topological graph for randomly-shuffled similarity matrix (cross-correlation threshold = 0.55 and core node connectivity threshold = 7).

(E) Topological graphs (as in Figure 2) with ICN nodes color-coded by frequency band (left) and subject (right).

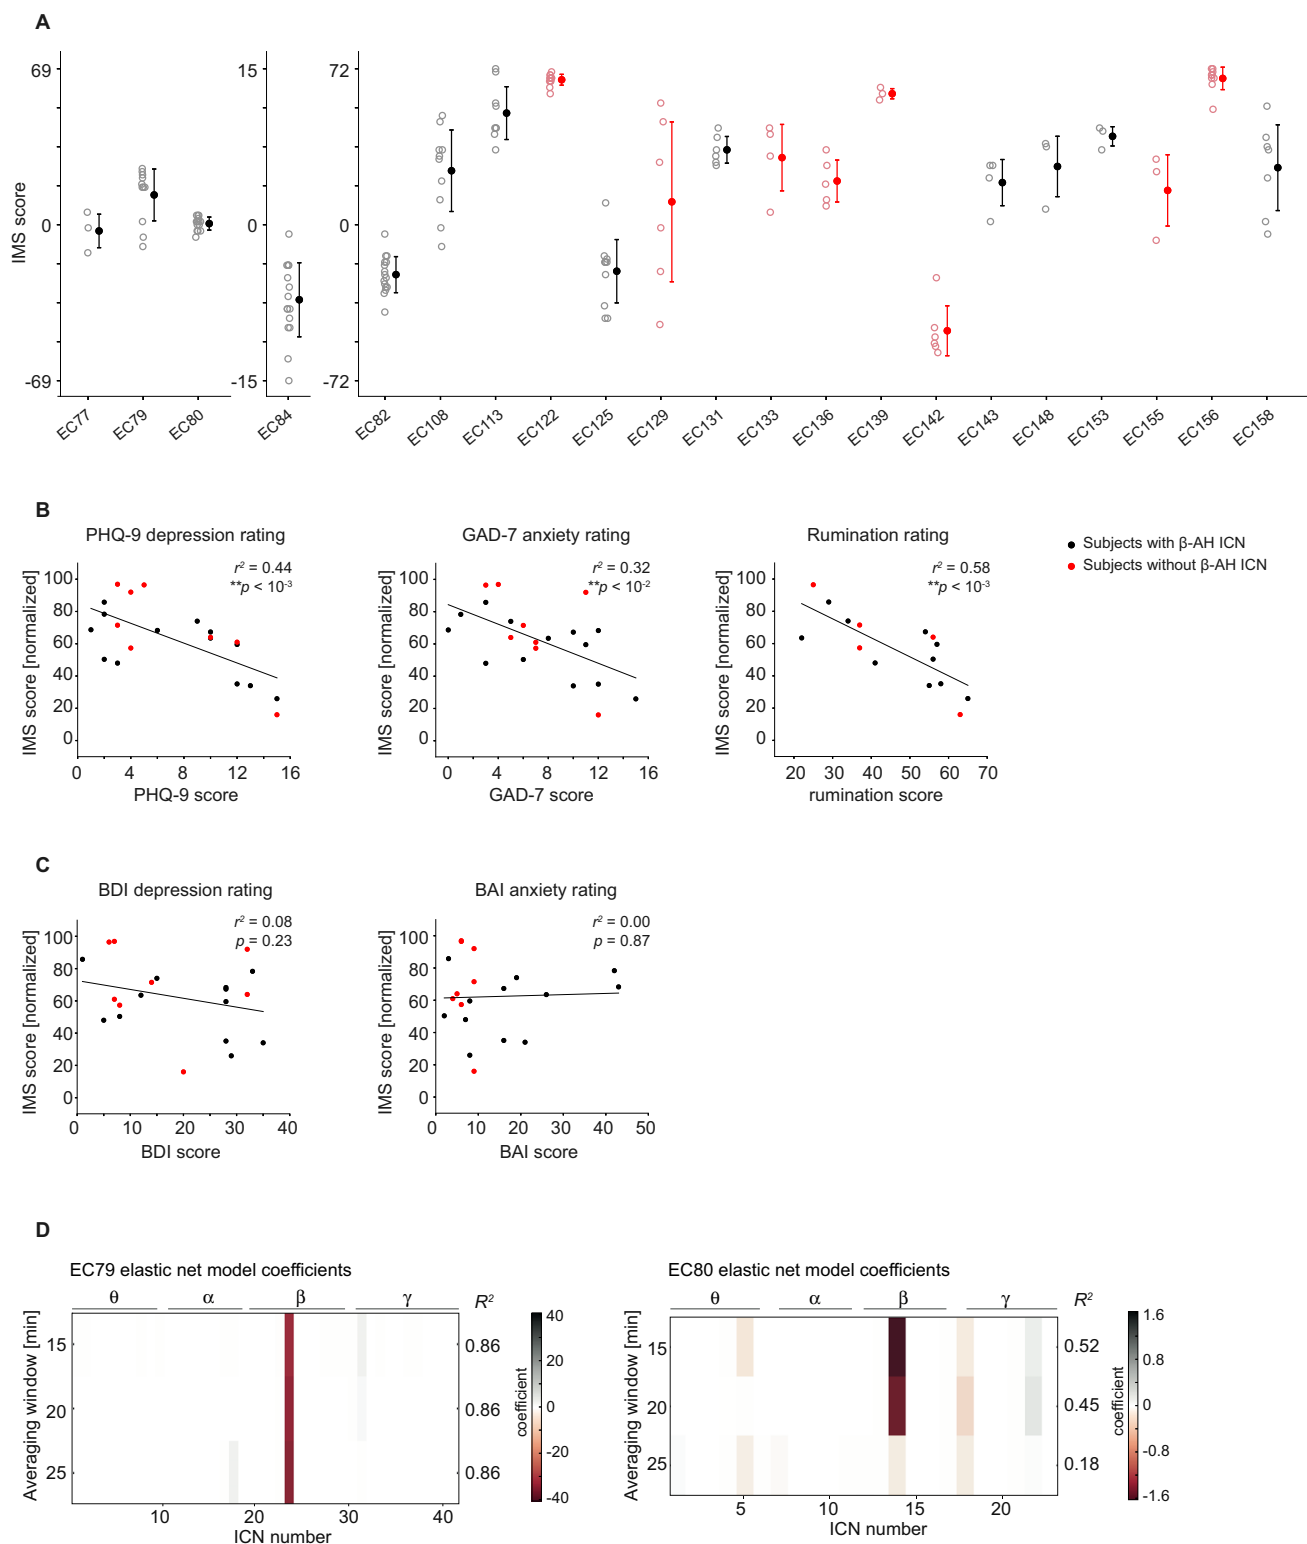


Figure S4. IMS Scores across Subjects, Related to Figure 3

(A) IMS scores for each subject, color-coded by presence (black) or absence (red) of mood-related β -AH ICN. IMS score range depends on total number of questions in IMS-questionnaire: 23-questions (EC77, EC79, EC80), 5-questions (EC84), or 24-questions (rest).

(B) Mean IMS score versus PHQ-9 (left), GAD-7 (middle) or Rumination (right) scores for subjects with (black) or without (red) mood-related β -AH ICN.

(legend continued on next page)

(C) Mean IMS score versus BDI (left) or BAI (right) scores for subjects with (black) or without (red) mood-related β -AH ICN. IMS: immediate mood scaler; PHQ-9: Patient-Health Questionnaire-9; GAD-7: Generalized Anxiety Disorder-7; BDI: Beck Depression Inventory; BAI: Beck Anxiety Inventory.

(D) Elastic net model coefficients for 15-, 20- and 25-minute averaging windows for two example subjects (EC79 and EC80). The model coefficients are very similar over all three averaging windows for EC79, and over 15- and 20-minute windows for EC80. This illustrates that the model output is robust against a range of window parameters.






Figure S5. The Same Activity-Mood Relationship Holds across All Subjects with a β -AH ICN, Related to Figure 5

Left: regression of IMS score against temporal variance of AMY-HPC β -coherence over 20-minute averaging window centered around each IMS score time point and z-score normalized within each individual for the 7 subjects with the β -AH ICN who were not used in the elastic net analysis. Grey confidence interval represents 95% confidence on regression estimate. A linear model was fit to these data and used for mood predictions for remaining subjects. Center/right: IMS score predictions using linear model for 6 subjects used in elastic net analysis and with β -AH ICN (center), and 8 subjects not used in elastic net analysis and without β -AH ICN (right). R^2 regression scores describe how well the model fits the data, where a negative score corresponds to a worse prediction than a horizontal line. p -values assess significance using 10^6 -fold permutation test for residuals against the line $y = x$ (dashed line).


A Variance of AMY-HPC coherence




B AMY-AMY coherence



C AMY power



HPC-HPC coherence



HPC power




Figure S6. Activity within AMY or HPC Alone Does Not Correlate with IMS Score across Subjects, Related to Figure 6

(A) Heatmaps of regression *p*-value between IMS score and temporal variance of AMY-HPC coherence in θ , α , β or γ frequency bands, across range of averaging windows, t_{ave} , and start times from each IMS data point, Δt (as in Figure 6B).

(B) Left: Regression of IMS score against mean AMY-AMY β -coherence (upper) and HPC-HPC β -coherence (lower) over 20-minute averaging window centered around each IMS score time point and z-score normalized within each individual, for 13/21 subjects with β -AH ICN (color-coded as in Figure 6A). Right: heatmap of regression *p* values across range of averaging windows, t_{ave} , and start times from each IMS data point, Δt (as in Figure 6B). Regression plots on left correspond to $t_{ave} = 20$ min and $\Delta t = -10$ min.

(C) Regression and heatmaps (as in (B)) for mean AMY β -power (upper) and HPC β -power (lower).




Figure S7. Interictal Discharge Rate Does Not Correlate with IMS Score across Subjects, Related to Figure 6

(A) Examples of interictal discharges (red dots) detected from AMY electrodes in two example subjects.

(B-D) IMS score (B), β -coherence variance (C), and β -coherence mean (D) versus interictal discharge rates in AMY (left) and HPC (right) electrodes, using 20-minute averaging window centered around each IMS data point for 13/21 subjects with β -AH ICN (color-coded as in Figure 6A).

(E) IMS score versus β -coherence variance (left) or β -coherence mean (right) after correcting for relationship with AMY discharge rate (color-coded as in Figure 6A).

(F) The mean temporal variance of β -coherence (left) and mean IMS score (right) during each 4-hour period following administration of the anti-epileptic drug levetiracetam (LEV), or during equivalent periods when no medication was received. Data are shown for three subjects, one who consistently received LEV (EC108, 5 days on medication, green), one who received LEV on some days (EC79, 4 days on and 2 days off medication, black), and one who did not receive any anti-epileptic medication (EC82, 6 days off medication, red).