Dynamic Coupling Index-Based Network Reconstruction for Heterogeneous Time Series Systems

Jeshwanth Mohan[1,3], Bharath Ramsundar[2], and Sandya Subramanian[3]

[1]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
[2]Deep Forest Sciences, Inc.
[3]Department of Computational Precision Health, University of California, Berkeley and UCSF

Abstract

Learning networks from heterogeneous multivariate time series is a recurring problem across neuroscience, social science, and economics. Existing edge estimators sit at two extremes. Nonparametric correlation measures such as Pearson correlation, Spearman rank correlation, and mutual information are cheap but capture only simple coupling on near-stationary signals. Fully parametric approaches such as vector autoregression and DCC–GARCH model richer dynamics but need long records and degrade on large panels. The common middle, with moderate data, heterogeneous signals, and quasi-linear coupling, is served by neither. This paper introduces the Dynamic Coupling Index (DCI), a windowed correlation metric that fills this middle. DCI normalizes each signal against its own estimated period before taking a windowed inner product of the residuals. It reduces to Pearson correlation, adds a single per-signal hyperparameter, matches Pearson's computational cost, and has a closed-form null variance for significance testing. We validate DCI on three testbeds with external ground truth. On synthetic networks it recovers time-varying community structure where Pearson correlation and mutual information fail. On EEG it recovers occipital alpha synchronization during eyes-closed rest and ranks a default-mode-proxy region above the rest of cortex in more subjects than either baseline. On the FRED-MD macroeconomic panel evaluated against NBER recession dates it separates recessions from expansions more sharply than the Forbes–Rigobon adjustment, Spearman correlation, Pearson correlation, and DCC–GARCH. DCI improves on standard correlation estimators when signals differ in scale, periodicity, or volatility.

1 Introduction

Complex systems in neuroscience, ecology, climate, epidemiology, and macroeconomics are often modeled as dynamic networks. The components produce measurable signals, and the strength of the connections between them changes over time [1–6]. Systems and cognitive neuroscientists treat the brain as a network whose dynamic interregional communication underlies decision-making and sensory perception [7, 8]. Contact-network models drive infectious-disease forecasting in public health [9]. In finance and macroeconomics, networks are used to study market behavior and regime shifts [6, 10–12].

Any such framework depends on how the edges between nodes are defined. Defining an edge is hard when the nodes are different kinds of entity, because their signals run on different timescales, carry different units, and come from different generative processes. Existing edge measures handle this poorly at both ends of the complexity range. Simple nonparametric measures are cheap, interpretable, and work on short records. These include the Pearson correlation coefficient (PCC) [13], Spearman rank correlation [14], and binned mutual information (MI) [15]. They reconcile dissimilar signals poorly and capture only the simplest coupling [15, 16]. Richer models capture more structure but cost more. Vector autoregression [17], state-space system identification [18], multivariate volatility models such as DCC–GARCH [19, 20], and coupled-oscillator models [21] all need far more data and computation to fit their parameters, and they scale poorly to large networks. Few methods recover moderately complex coupling from short records at moderate cost.

This paper describes such a method, the Dynamic Coupling Index (DCI), introduced in preliminary form as the Windowed Variance Correlation [22]. DCI is a windowed correlation metric with one estimated parameter per signal. It handles differences in magnitude, periodicity, and variability, and it costs the same per window as PCC. Figure 1 summarizes the construction. Each channel is normalized against its own estimated period. A windowed inner product of the normalized signals then gives an edge weight for every ordered pair, and the resulting matrices define a graph that changes over time. DCI captures coupling carried by amplitude, phase, and volatility. It does not capture the nonlinear or conditional structure that a dynamical-systems model identifies. It is therefore a drop-in correlation measure that is more robust than the simple metrics and cheaper than the complex ones.

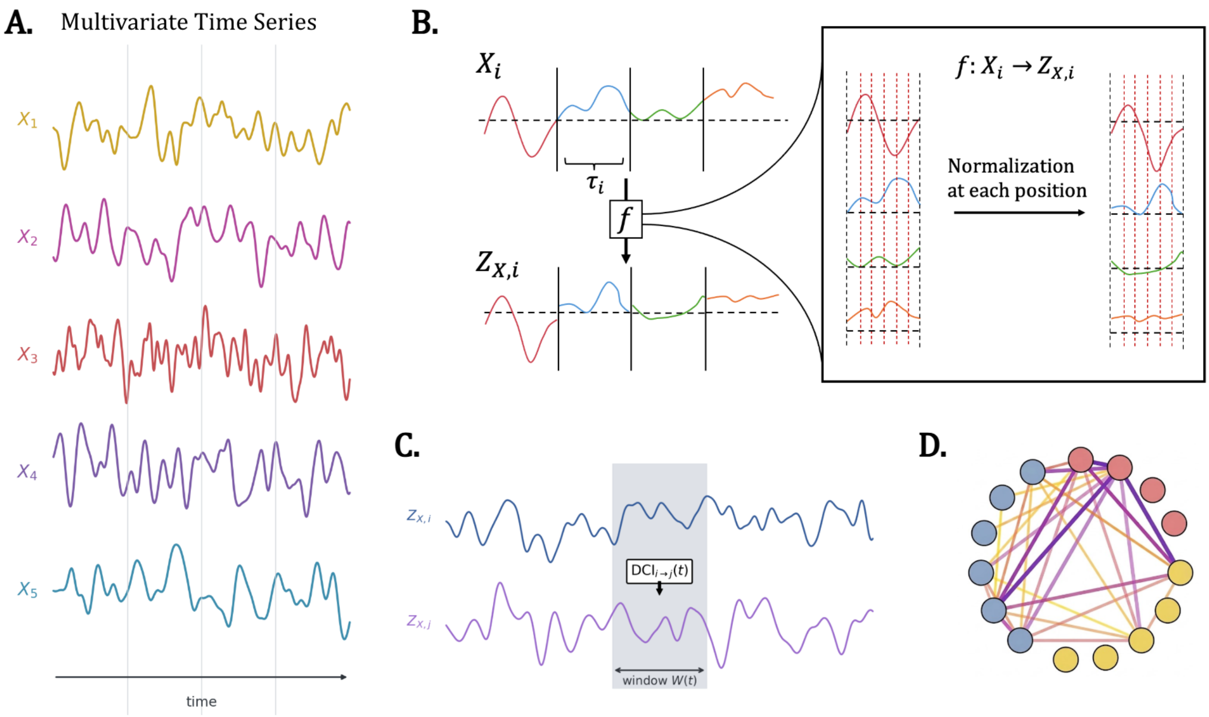

Figure 1. The Dynamic Coupling Index. (A) The input is a d-dimensional, length-L multivariate record; each channel becomes a graph node. **(B)** Each channel X_i is cut into β_i segments of its own estimated period τ_i and normalized position-wise (Eq. 2). This removes the shared periodic profile of the channel and leaves the deviations Z_i. **(C)** For every ordered pair (i, j) the coupling on a window is the inner product of the normalized signals over that window (Eq. 3). **(D)** Each pairwise calculation makes an edge in a dynamic graph

Several estimators sit near DCI. Amplitude envelope correlation and its leakage-corrected form AEC-c [23–26] correlate the Hilbert amplitude envelopes of two band-limited signals and dominate EEG and MEG practice. They capture amplitude coupling cleanly but discard phase-aligned linear coupling. Cyclostationary analysis handles signals whose correlations are themselves periodic in time [27], while DCI estimates one correlation after the periodic structure has been normalized away. DCC–GARCH conditions the correlation matrix on a parametric volatility model and is standard in finance [19, 28], but fitting it across many series drives it back toward the unadjusted correlation. Structure learning in dynamic Bayesian networks [29, 30] solves a different problem, recovering conditional-independence structure across time, and finds directional dependencies that DCI cannot at much greater cost. Figure 2 places DCI among these methods, and Appendix A reviews each one.

The periodic mean and variance construction used for normalization has a long history in orbital mechanics and astronomical time-series analysis, and to our knowledge it has not previously been applied to correlation-based graph inference.

Section 2 defines DCI and gives its basic properties. Sections 3 and 4 test it against established baselines on synthetic networks and on neurophysiological and macroeconomic data. Section 5 covers what the method does and does not recover, along with its cost and data requirements. Section 6 gives the experimental design. The appendices hold the related-work review, the derivations, the experimental details, and the supporting results, in that order.

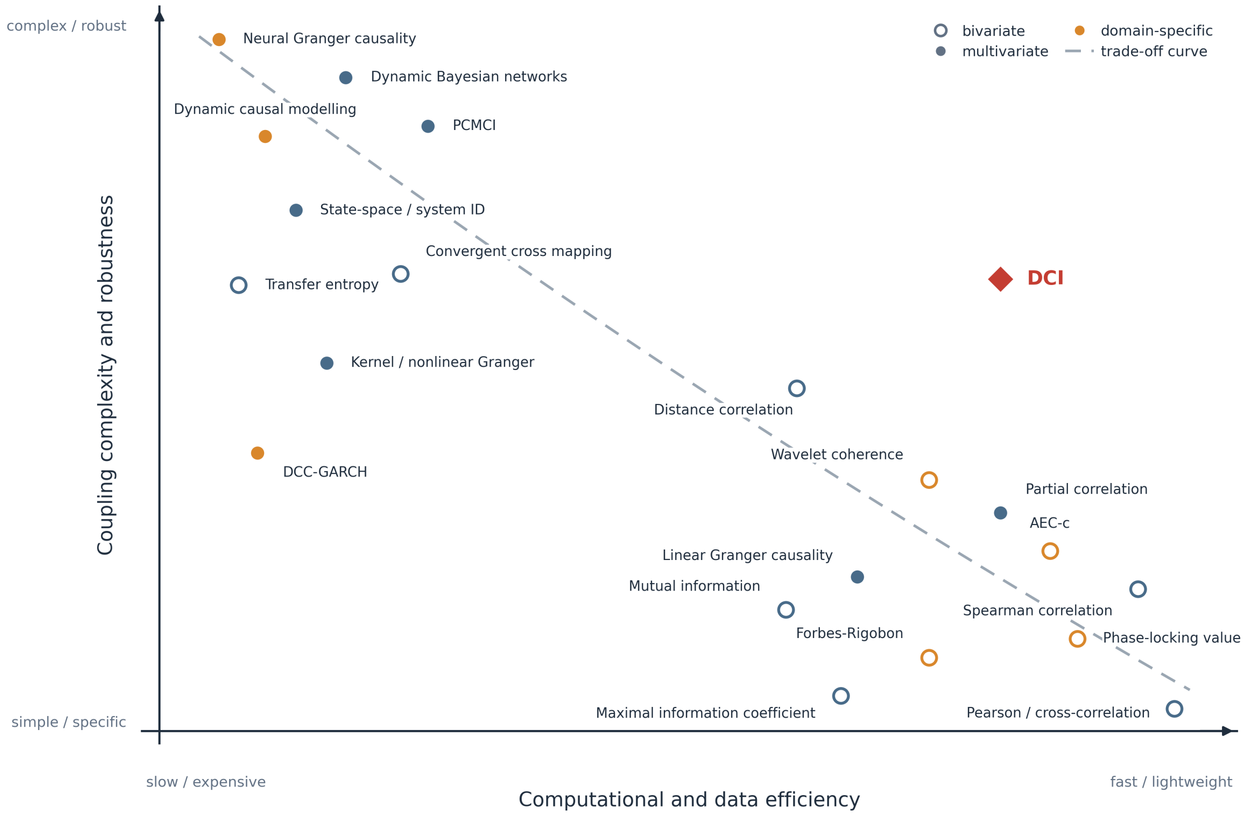

Figure 2. Coupling estimators arranged by computational and data efficiency against coupling complexity and robustness. Marker fill separates bivariate estimators, multivariate estimators, and estimators developed for a specific domain; the dashed line traces the empirical trade-off that most methods sit near. DCI is positioned off that trade-off curve because the periodic normalization adds robustness at the per-window cost of Pearson correlation. Appendix A reviews every method shown.

2 The Dynamic Coupling Index

2.1 Problem setup and notation

Let $X \in \mathbb{R}^{d \times L}$ be a d-dimensional time series of length L, with $X_i(t)$ the t-th sample of series i. For each ordered pair (i, j) and each window $\mathcal{W} = [t_1, t_2] \subseteq [1, L]$ of length $l = t_2 - t_1 + 1$, the goal is a coupling weight that can serve as an edge weight in a dynamic graph $G^{\mathcal{W}} = (V, E^{\mathcal{W}})$ with $V = \{X_1, \ldots, X_d\}$. Sliding $\mathcal{W}$ along the record yields a sequence of weighted adjacency matrices describing the time-evolving coupling topology.

The same symbols are used throughout the paper and the appendices. τ_i is the estimated period of channel i and $\beta_i = \lfloor L/\tau_i \rfloor$ is the number of whole periods in the record. Z_i is the periodicity-normalized form of X_i, and t_{off} is the lag between the two channels of a pair. In sliding-window analyses, w is the window length, Δ the stride, and m the number of windows.

2.2 Periodicity-aware normalization

DCI normalizes each channel against its own period, which decouples the normalization scale from the query scale and lets components with very different timescales coexist in one analysis. For each X_i, a characteristic period τ_i is estimated from the first statistically significant local maximum of the sample autocorrelation function, using the white-noise band of Bartlett [31]; Appendix B.1 gives the estimator. Given τ_i, the record is partitioned into β_i periods and position-wise statistics are computed at each within-period position $k \in \{1, \ldots, \tau_i\}$,

$$\mu_{X,i}(k) = \frac{1}{\beta_i} \sum_{u=0}^{\beta_i - 1} X_i\big(k + u\tau_i\big), \qquad \sigma_{X,i}(k) = \sqrt{\frac{1}{\beta_i} \sum_{u=0}^{\beta_i - 1} X_i\big(k + u\tau_i\big)^2 \;-\; \mu_{X,i}(k)^2}. \tag{1}$$

The periodicity-normalized signal is the phase-resolved z-score

$$Z_i(t) = \frac{X_i(t) - \mu_{X,i}(t \bmod \tau_i)}{\sigma_{X,i}(t \bmod \tau_i)}. \tag{2}$$

Removing each channel's intrinsic periodic structure before cross-channel comparison means that residual comovement reflects coupling rather than incidental phase alignment. Figure 1B illustrates the construction.

2.3 Definition

For a pair (i, j), window $\mathcal{W} = [t_1, t_2]$, and lag t_{off},

$$\boxed{\mathrm{DCI}^{\mathcal{W}}_{i,j}[t_{\text{off}}] \;=\; \sum_{t=t_1}^{t_2} Z_i(t)\, Z_j(t + t_{\text{off}})} \tag{3}$$

The lag t_{off} admits directed, time-delayed coupling. All analyses in this paper use $t_{\text{off}} = 0$; Appendix D.2 characterizes the behavior at nonzero lags as a sensitivity check.

2.4 Bounded variant and reduction to Pearson correlation

By the Cauchy–Schwarz inequality, with the combined volatility

$$D^{\mathcal{W}}_{i,j} \;=\; \sqrt{\Big(\textstyle\sum_{t=t_1}^{t_2} Z_i(t)^2\Big)\Big(\textstyle\sum_{t=t_1}^{t_2} Z_j(t + t_{\text{off}})^2\Big)}, \tag{4}$$

the bound $|\mathrm{DCI}^{\mathcal{W}}_{i,j}| \le D^{\mathcal{W}}_{i,j}$ holds, which gives the bounded variant

$$\boxed{\mathrm{Corr}^{\mathcal{W}}_{i,j} \;=\; \mathrm{DCI}^{\mathcal{W}}_{i,j} / D^{\mathcal{W}}_{i,j} \;\in\; [-1, 1]} \tag{5}$$

The pair (Corr, D) carries separable information, with Corr a structure-blind linear coupling and D the joint volatility over the window. Proposition B.1 in Appendix B.2 shows that when $\tau_i = \tau_j = 1$, $t_{\text{off}} = 0$, and $\mathcal{W} = [1, L]$, the construction reduces exactly to $\mathrm{DCI}^{\mathcal{W}}_{i,j} = L\rho_{X_i,X_j}$, $D^{\mathcal{W}}_{i,j} = L$, and $\mathrm{Corr}^{\mathcal{W}}_{i,j} = \rho_{X_i,X_j}$, so PCC is a degenerate special case of DCI.

2.5 Null variance and calibrated metric

Write $\mathrm{DCI}^{\mathcal{W}}_{i,j} = \sum_{t=t_1}^{t_2} Y_t$ with $Y_t = Z_i(t) Z_j(t + t_{\text{off}})$. Under cross-channel independence and weak stationarity, Theorem B.3 in Appendix B.3 gives

$$\boxed{\mathrm{Var}\big(\mathrm{DCI}^{\mathcal{W}}_{i,j}\big) \;=\; l \;+\; 2\sum_{\tau=1}^{l-1} (l - \tau)\, R_{Z_i Z_i}(\tau)\, R_{Z_j Z_j}(\tau)} \tag{6}$$

Under the same conditions DCI is approximately Gaussian, which gives the variance-normalized statistic

$$\boxed{\overline{\mathrm{DCI}}^{\mathcal{W}}_{i,j} \;=\; \mathrm{DCI}^{\mathcal{W}}_{i,j} \Big/ \sqrt{\mathrm{Var}(\mathrm{DCI}^{\mathcal{W}}_{i,j})}} \tag{7}$$

and, after mapping through the standard normal distribution function Φ, a signed and null-calibrated score

$$\boxed{\widetilde{\mathrm{DCI}}^{\mathcal{W}}_{i,j} \;=\; 2\,\Phi\big(\overline{\mathrm{DCI}}^{\mathcal{W}}_{i,j}\big) - 1 \;\in\; [-1, 1]} \tag{8}$$

The magnitude of $\widetilde{\mathrm{DCI}}$ equals $1 - 2p$, where p is the two-tailed Gaussian-null p-value, so the score encodes sign and statistical strength together. Equation (6) is validated by a 10^5-trial Monte Carlo in which empirical and theoretical standard deviations agree to within 0.1% (Appendix D.1). Sweeps over l, τ_i, and t_{off} confirm robustness to moderate misspecification (Appendix D.2).

2.6 Behavior under noise

Before moving to networks, DCI was characterized at the level of a single pair on a controlled amplitude-coupling benchmark. Three heterogeneous base waveforms receive injected amplitude-coupling events spanning the four sign combinations of joint scaling, and the signal-to-noise ratio (SNR) is swept while the recovered coupling trace is compared against the noiseless ground truth. The waveform definitions and sweep settings are in Appendix C.1.

Both DCI variants recover all four injected events down to SNR ≈ 5 dB (Fig. 3). The bounded form Corr of Eq. (5) and the variance-normalized form $\overline{\mathrm{DCI}}$ of Eq. (7) give the most stable root-mean-square error across the sweep, while MI and PCC computed on the raw signals report spurious coupling whenever the unnormalized waveforms happen to align.

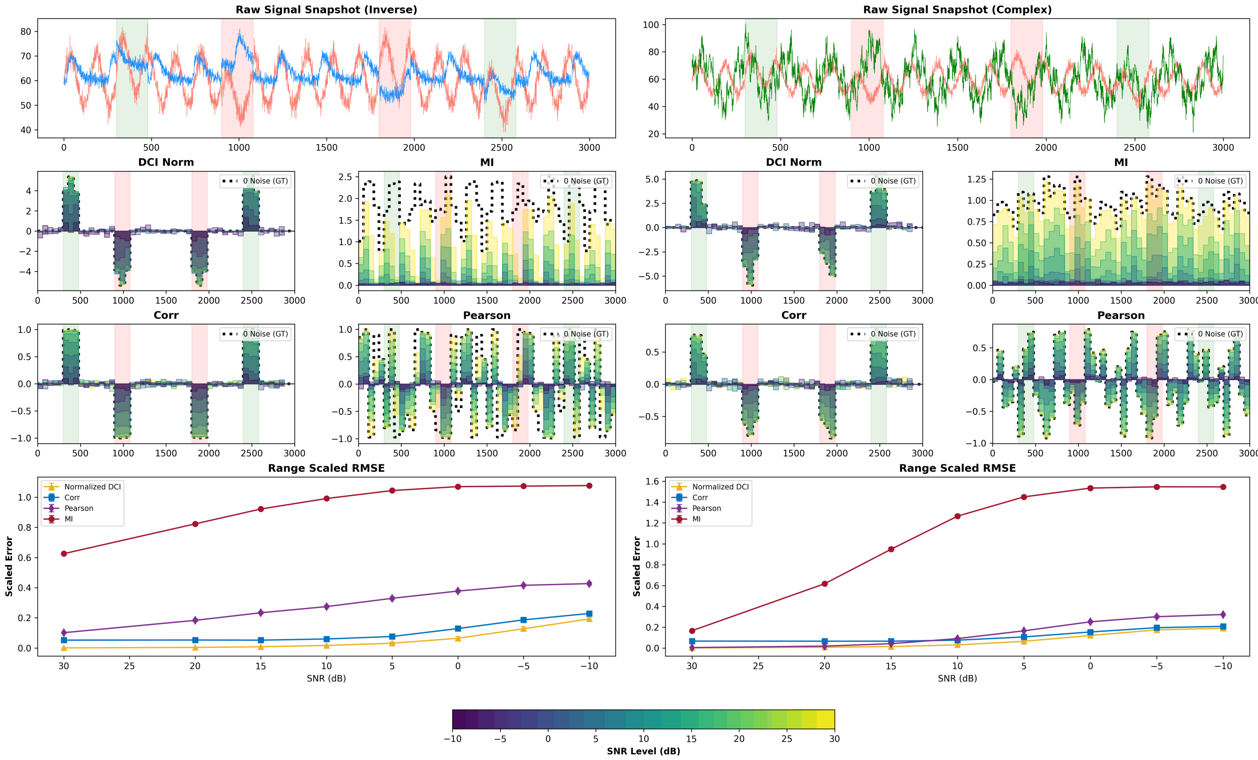

Figure 3. Noise robustness. Both DCI variants recover the four injected amplitude-coupling events down to SNR ≈ 5 dB. Bottom row: root-mean-square error against SNR. PCC and MI fail to recover the injected coupling and report spurious edges.

3 Synthetic Network Experiments

Two complementary synthetic experiments evaluate DCI in a full network setting. The first recovers time-varying community structure under several injected coupling mechanisms and is the central synthetic result. The second recovers a fixed structural graph from a stable linear state-space model and is summarized here, with the setup in Appendix C.3 and the results in Appendix D.3.

3.1 Dynamic community recovery

3.1.1 Generative model

Each simulation draws d base channels, where channel i is a random waveform $X_i^{(0)}$ from a family of sines, multi-frequency sines, and square waves (Appendix C.2.1, Eq. (23)). The ground truth is a set of K coupling modules. Module c has a node set $M_c \subseteq \{1, \ldots, d\}$, an activity window $\mathcal{W}_c = [a_c, b_c]$, and one of five coupling mechanisms g_c. Inside its window, each member of a module is rewritten as a function of its own value, the module mean, and a strength parameter s,

$$X_i(t) \leftarrow g_c\Big(X_i^{(0)}(t),\ \bar{X}_{M_c}(t);\ s\Big), \qquad i \in M_c,\ t \in \mathcal{W}_c, \qquad \bar{X}_{M_c}(t) = \frac{1}{|M_c|} \sum_{u \in M_c} X_u^{(0)}(t). \tag{9}$$

The five mechanisms apply Eq. (9) to different coupling channels. BLEND mixes each member toward the module mean, $g(x, \bar{x}; s) = (1 - s)x + s\bar{x}$, which aligns phase and frequency. ELEVATE and DEPRESS apply a common amplitude gain $g(x, \bar{x}; s) = (1 \pm s)x$. OPPOSEAMP applies opposite gains to the two halves of the module, which gives anti-correlated amplitude modulation at the same coupling magnitude. PHASELOCK rebuilds each member from its own Hilbert envelope and a phase shifted toward the module-mean phase, leaving amplitude untouched. Appendix C.2.3 gives each mechanism in full, and Appendix C.2 gives the simulator parameters and the merge rule applied when amplitude-driven modules overlap in time (Appendix C.2.4).

3.1.2 Evaluation and results

A window of length w is slid across each record, adjacency matrices are built with $\widetilde{\mathrm{DCI}}$, sliding-window PCC, and binned MI, and each window is partitioned by Leiden modularity maximization [32]. Recovered communities are scored against the injected ground truth. The headline score is Significant Community Recall (SCR), which weights each statistically significant window by its topological recall and is defined in Appendix B.6, Eq. (22); it is reported alongside confusion-based metrics and active-window accuracy. A matched injection-free control run of each estimator on the same base signals verifies null behavior. Window and scoring settings are in Section 6.

Across 100 simulations (Table 1), DCI outperforms PCC and MI on every metric, with SCR $= 0.713 \pm 0.029$ against 0.211 ± 0.030 for PCC and 0.077 ± 0.014 for MI, and it is the only estimator whose scores separate cleanly from its own injection-free control. Figure 4 shows a representative recovery. The advantage holds across injection strengths and graph sizes: sweeps over $s \in [0, 3]$ and $d \in \{10, \ldots, 100\}$ appear in Supplement Figs. S1 and S2, together with a worst-case example (Supplement Fig. S3) in which one dominant PHASELOCK module inflates the clustering baseline and masks smaller events even though the Leiden partition recovers all three modules correctly.

Table 1. Dynamic community recovery on multimodal synthetic networks.

Method	SCR	Active Acc.	Recall	Jaccard	F_2	Sig. Dur.
DCI	$\mathbf{0.713 \pm 0.029}$	$\mathbf{0.818 \pm 0.010}$	$\mathbf{0.878 \pm 0.013}$	$\mathbf{0.588 \pm 0.021}$	$\mathbf{0.797 \pm 0.015}$	$\mathbf{0.799 \pm 0.030}$
Pearson	0.211 ± 0.030	0.672 ± 0.010	0.612 ± 0.018	0.349 ± 0.014	0.556 ± 0.015	0.279 ± 0.033
MI	0.077 ± 0.014	0.587 ± 0.009	0.650 ± 0.013	0.313 ± 0.013	0.550 ± 0.011	0.107 ± 0.015
DCI (Control)	0.059 ± 0.004	0.625 ± 0.004	0.580 ± 0.008	0.301 ± 0.008	0.515 ± 0.007	0.098 ± 0.006
Pearson (Control)	0.062 ± 0.004	0.636 ± 0.007	0.538 ± 0.011	0.294 ± 0.010	0.490 ± 0.009	0.111 ± 0.006
MI (Control)	0.066 ± 0.004	0.573 ± 0.006	0.641 ± 0.012	0.297 ± 0.010	0.537 ± 0.010	0.103 ± 0.007

Mean ±95% CI; SCR = Significant Community Recall; Active Acc. = active-window accuracy; Sig. Dur. = Duration of Significant Clustering. DCI median SCR = 0.7345.

Figure 4. Representative module detection. SCR = 0.992. Three sequential coupling modules (Depress, OpposeAmp, Elevate) occupy the latter half of the record. DCI resolves all three with distinct clustering-coefficient peaks and threshold crossings; PCC and MI miss the events.

3.2 Structural connectivity and the indirect-coupling limitation

The second experiment asks whether DCI recovers a fixed structural graph from a linear state-space model. Stochastic block model weight matrices W are drawn, rescaled to a fixed stable spectral radius, and used to drive a purely stochastic linear time-invariant recursion with white Gaussian innovations,

$$x(t) \;=\; Wx(t-q) \;+\; \epsilon(t), \qquad \epsilon(t) \sim \mathcal{N}(0, \sigma^2 I), \tag{10}$$

with lag q. Time-averaged DCI and PCC adjacency matrices are compared against W using matrix correlation and cosine similarity. Appendix C.3 gives the setup and the evaluation measures, and Appendix D.3 gives the results.

DCI and PCC recover the block structure comparably, with matrix correlation 0.917 ± 0.003 and cosine similarity 0.916 ± 0.004 for DCI against 0.929 ± 0.004 and 0.929 ± 0.004 for PCC. Both estimate correlation rather than the state-transition matrix W, and because a stable recursion of the form in Eq. (10) propagates influence along every path of the graph, indirect couplings bleed into the recovered weights for both methods.

4 Real-World Data Experiments

DCI was evaluated on two heterogeneous testbeds whose ground truth is established independently of any connectivity method: the eyes-closed and eyes-open alpha response in EEG, and NBER recession labels for the FRED-MD macroeconomic panel.

4.1 EEG: alpha-band connectivity under the Berger effect

4.1.1 Background and design

Closing the eyes raises occipital alpha-band (8–12 Hz) activity sharply and synchronizes it across regions, and opening the eyes breaks that synchrony [33–35]. This Berger effect is one of the most reproducible phenomena in electrophysiology, which makes it a useful external check. Any method that measures coupling should report more coupling among posterior channels with the eyes closed than with the eyes open. A second, weaker expectation concerns the default mode network (DMN), a midline fronto-parietal system engaged during rest [36, 37], which should also strengthen during eyes-closed rest while remaining anatomically distinct from visual cortex.

Recordings come from the EEG Alpha Waves dataset [38], in which subjects alternate ten-second eyes-closed and eyes-open blocks. Thirteen of twenty subjects passed amplitude and alpha-SNR screening. After ocular-artifact removal and band-pass filtering to the alpha band, a window is slid across the continuous record and adjacency matrices are built in three parallel pipelines using $\widetilde{\mathrm{DCI}}$, absolute PCC, and leakage-corrected amplitude envelope correlation (AEC-c), the latter being the standard connectivity metric in the field [23–26]. Three regions of interest were fixed in advance: occipital channels, a four-channel DMN proxy over posterior-midline hubs, and the remaining nine channels as a rest-of-brain comparison. Preprocessing and statistical procedures are in Section 6.

4.1.2 Results

DCI reports the expected eyes-closed increase and the baselines do not. Node strength in the occipital region roughly doubles when the eyes close under DCI, and every one of the thirteen subjects shows the increase individually. PCC reports no occipital difference at all, and AEC-c reports a smaller increase. The same eyes-closed increase appears in the DMN proxy under DCI in twelve of thirteen subjects, where PCC reverses direction. Clustering coefficients follow node strength for DCI in both regions. Figure 5 shows the effect in a representative subject, and Table 2 gives the full region statistics for all three estimators. Per-subject values (Appendix D.4.5) and the within-subject percent-change analysis (Appendix D.4.2) are in Appendix D.4.

Table 2. EEG ROI statistics: DCI, PCC, and corrected AEC, eyes-closed vs. eyes-open ($n = 13$ subjects). Shaded p-value cells mark $p < 0.05$; asterisks indicate level ($^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, ns otherwise). Paired two-tailed t-tests across subjects.

		Node strength			*Clustering coefficient*		
ROI / Metric	Method	EC	EO	p-value	EC	EO	p-value
Occipital	DCI	11.41 ± 1.38	6.18 ± 2.05	$5.45 \times 10^{-5***}$	0.777 ± 0.074	0.523 ± 0.123	$2.76 \times 10^{-4***}$
Occipital	PCC	7.91 ± 0.96	8.31 ± 1.23	0.195^{ns}	0.514 ± 0.052	0.559 ± 0.075	$2.48 \times 10^{-2*}$
Occipital	AEC	2.30 ± 0.40	1.89 ± 0.30	$9.52 \times 10^{-3**}$	0.635 ± 0.055	0.647 ± 0.063	0.446^{ns}
DMN-proxy	DCI	11.42 ± 1.53	6.84 ± 2.13	$1.09 \times 10^{-4***}$	0.778 ± 0.081	0.559 ± 0.119	$5.37 \times 10^{-4***}$
DMN-proxy	PCC	7.80 ± 0.89	8.48 ± 1.03	$9.31 \times 10^{-3**}$	0.513 ± 0.050	0.572 ± 0.066	$2.52 \times 10^{-3**}$
DMN-proxy	AEC	2.38 ± 0.43	1.93 ± 0.25	$6.53 \times 10^{-3**}$	0.647 ± 0.059	0.655 ± 0.051	0.653^{ns}
REST	DCI	11.10 ± 1.97	6.77 ± 1.72	$1.40 \times 10^{-4***}$	0.761 ± 0.106	0.552 ± 0.094	$5.73 \times 10^{-4***}$
REST	PCC	7.45 ± 0.68	7.46 ± 1.02	0.948^{ns}	0.497 ± 0.038	0.525 ± 0.066	0.085^{ns}
REST	AEC	2.26 ± 0.37	1.89 ± 0.23	$9.41 \times 10^{-3**}$	0.628 ± 0.048	0.648 ± 0.050	0.266^{ns}

The alpha response is scalp-wide, so coupling also rises in the rest-of-brain region and a specificity test is needed. Two such tests agree. Within subjects, the percent increase in coupling is largest in the occipital region and smallest in the rest of the brain, and the occipital-over-rest ordering is significant under DCI but not under AEC-c, while PCC runs in the opposite direction. Across channel subsets, an exact permutation test that excludes visual channels from the null pool places the DMN proxy in the top 2.5% of all four-channel subsets of non-visual cortex during eyes-closed rest, and returns it to the middle of the null during eyes-open. Occipital clustering moves in the direction expected from stimulus-locked alpha desynchronization [39], and the DMN proxy behaves as expected from its anticorrelation with goal-directed attention [40]. Neither shift reaches significance in the expected direction under PCC. Appendices D.4.2 and D.4.3 give both tests in full, and Appendix D.4.4 adds an edge-level network-based statistic [41] and a broadband reanalysis in which DCI still resolves the occipital effect, PCC still fails, and AEC-c reverses sign because broadband envelope energy is dominated by frequencies unrelated to the Berger phenomenon.

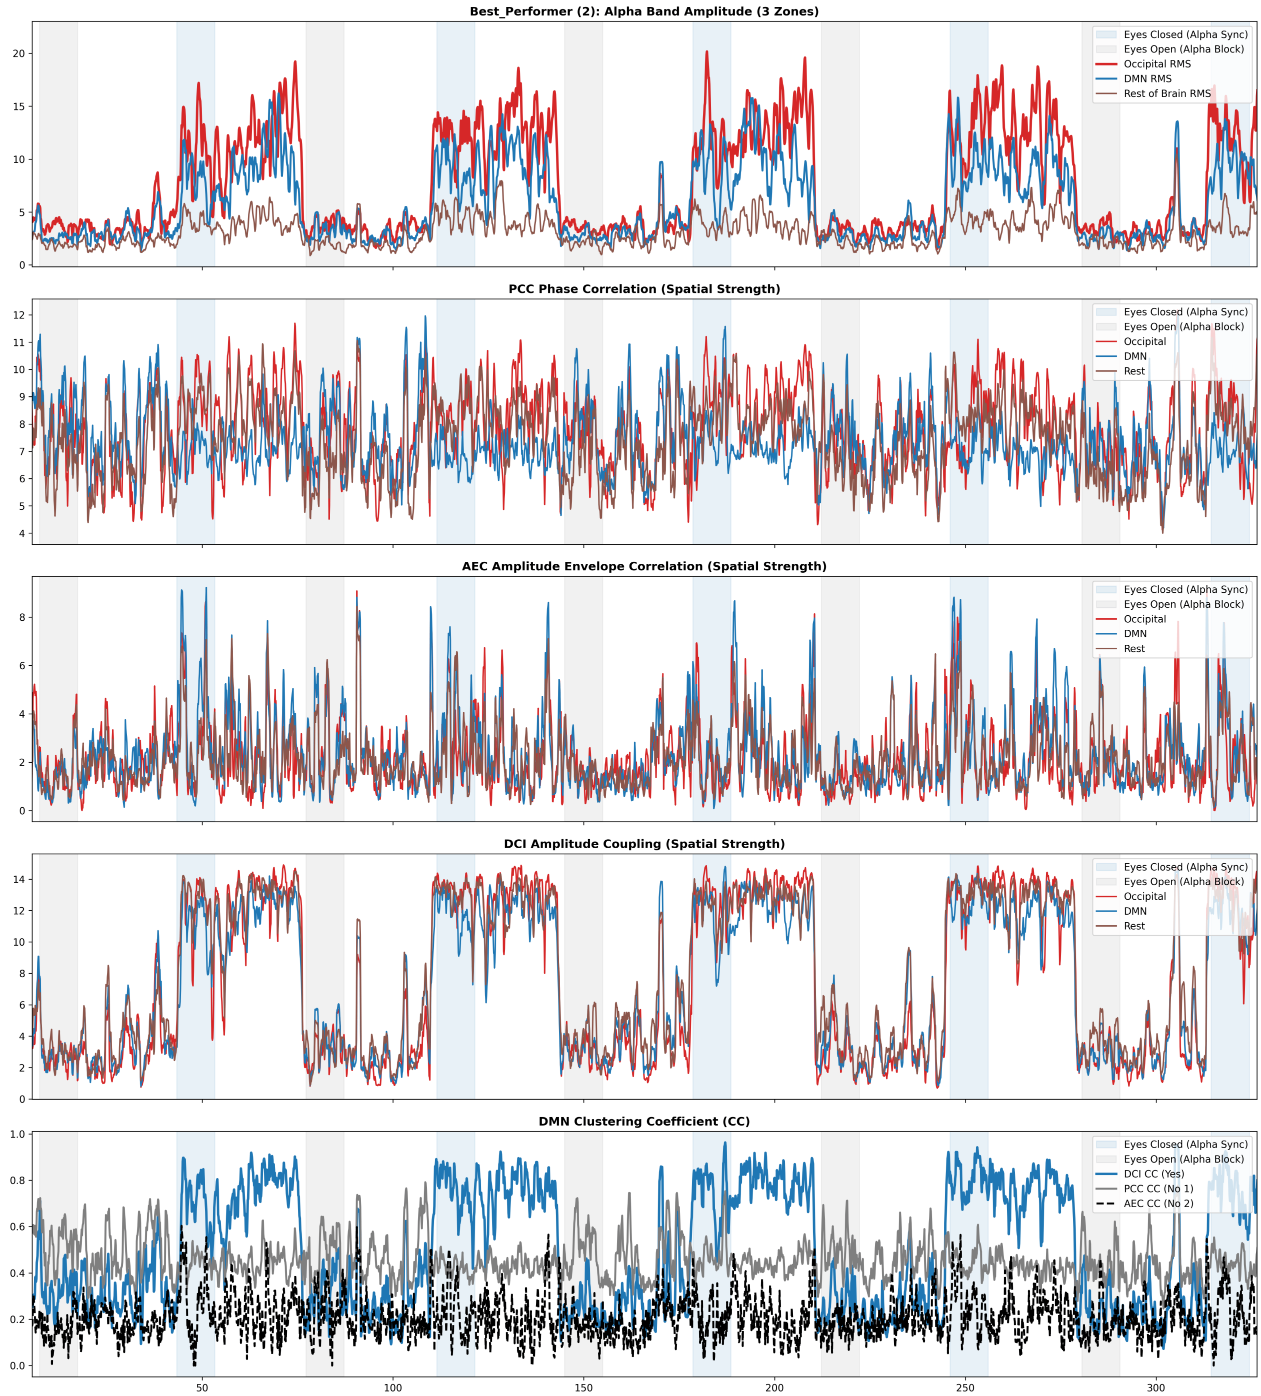

Figure 5. Alpha-band connectivity under the Berger effect in a representative subject. DCI tracks the eyes-closed and eyes-open alternation with amplitude-coupling excursions roughly an order of magnitude larger than the PCC response.

4.2 Macroeconomics: FRED-MD and NBER recessions

4.2.1 Background and design

Macroeconomic indicators spanning output, employment, prices, credit, housing, and financial markets are expected to couple more tightly during recessions than during expansions, a prediction from the financial-contagion literature [6, 28, 42]. Ordinary correlation estimators separate genuine coupling from a coincident rise in volatility poorly, which is the problem the Forbes–Rigobon (FR) adjustment was designed to address [20, 28].

Networks were built from FRED-MD [43], a panel of 134 monthly U.S. macroeconomic indicators running from 1959 to the present, and evaluated against the NBER recession indicator. A one-year window is slid across the panel, and each window yields five adjacency matrices from windowed PCC, Spearman rank correlation, the FR adjustment, DCC–GARCH [19], and $\widetilde{\mathrm{DCI}}$. Each matrix is summarized by five standard global graph features, and each window is aligned to the NBER label to test retrospectively how separable recession windows are from expansion windows. DCC–GARCH serves as the standard parametric heteroskedasticity-aware baseline; joint re-estimation is infeasible at this panel size, so canonical scalar defaults are used, and the resulting estimator tracks Pearson closely. Specifications are in Section 6.

4.2.2 Results

All five DCI features separate recession windows from expansion windows highly significantly, with density, transitivity, efficiency, and entropy rising during recessions and modularity falling (Fig. 6). Welch $|t|$ for DCI ranges from 13 to 17, against $|t| \in [4.7, 7.5]$ for PCC, Spearman, and DCC–GARCH; the full test table is Table 7 in Appendix D.5. Over the 65-year record, DCI produces localized excursions at almost every recession, with the 2007–09 and 2020 episodes driving the two largest movements, whereas PCC registers the same shocks at smaller magnitude against a noisier baseline (Fig. 7).

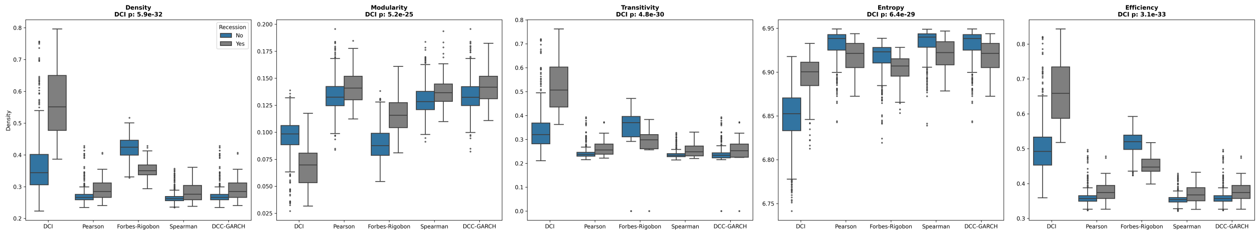

Figure 6. Distributional separation of graph features across NBER regimes. All five DCI features shift significantly between expansion (blue) and recession (grey) windows, with effect sizes larger than PCC, Spearman, or DCC–GARCH. DCC–GARCH tracks Pearson at this panel size because the parametric correlation step regularizes toward the sample correlation when $d = 134$.

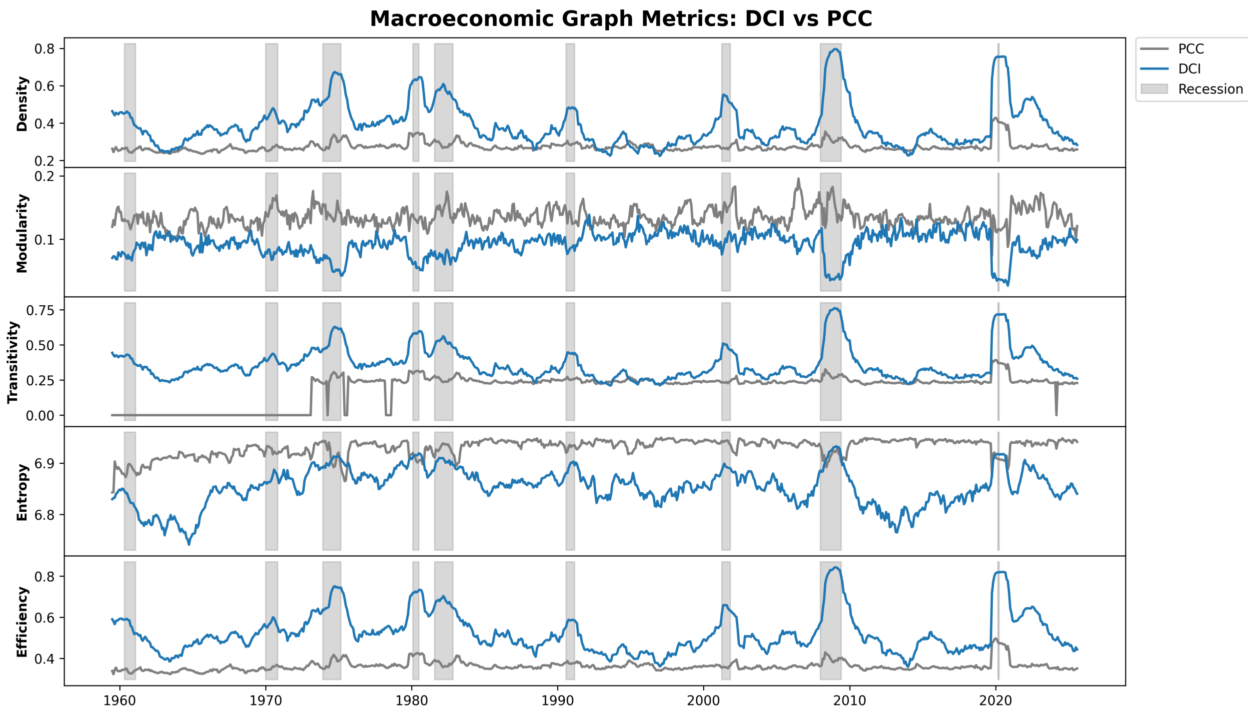

Figure 7. DCI and PCC graph features over U.S. business-cycle history. DCI (blue) localizes recession events (grey shading) as sharp topological excursions, with density, transitivity, efficiency, and entropy rising and modularity falling. PCC (grey) registers the same shocks at smaller magnitude against a noisier baseline.

Treating each feature as a single-feature recession score and sweeping its threshold gives the benchmarks in Table 3. DCI dominates Pearson, Spearman, and DCC–GARCH on every feature, and the aggregate curves in Fig. 8 place it highest overall at AUROC 0.90 and AUPRC 0.57. The comparison against FR is less clear-cut.

FR matches or exceeds DCI on density and efficiency, but four of its five features carry the opposite sign from DCI and need to be sign-flipped before they track recessions. The reason is structural. The FR adjustment divides by a denominator that grows with conditional variance, so when comovement and individual variance rise together, as they do in recessions, the volatility penalty outweighs the increase in the numerator and the adjusted features fall. DCI keeps the raw coupling direction in all five features. Its inverted modularity direction is itself meaningful, because recession networks become less modular and more systemically integrated [42]. DCC–GARCH and Pearson are nearly identical on a 134-variable panel, which is the parametric collapse noted above: an estimator that is sharp on small graphs regularizes toward simple correlation at scale.

Table 3. Retrospective separability of NBER recession windows.

Method	Metric	Density	Modularity	Transitivity	Entropy	Efficiency
DCI	AUROC	**0.93**	**0.87**†	**0.93**	**0.86**	0.93
	AUPRC	0.59	0.45	**0.60**	**0.59**	0.59
	Max F_1	0.60	**0.55**	**0.61**	**0.61**	0.62
Pearson	AUROC	0.70	0.66	0.72	0.75†	0.70
	AUPRC	0.28	0.20	0.28	0.24	0.28
	Max F_1	0.45	0.30	0.44	0.36	0.44
FR	AUROC	0.93†	0.87	0.77†	0.76†	**0.93**†
	AUPRC	**0.68**	**0.52**	0.31	0.24	**0.64**
	Max F_1	**0.65**	0.54	0.42	0.39	**0.65**
Spearman	AUROC	0.68	0.66	0.71	0.72†	0.68
	AUPRC	0.31	0.20	0.32	0.24	0.31
	Max F_1	0.44	0.30	0.43	0.37	0.44
DCC–GARCH	AUROC	0.70	0.66	0.66	0.75†	0.70
	AUPRC	0.28	0.20	0.28	0.24	0.28
	Max F_1	0.45	0.29	0.44	0.36	0.44

For each (feature, method) pair, AUROC, AUPRC, and Max F_1 are reported with the score sign-flipped if needed. **Bold**: best across methods within each feature on each metric. † inverted relationship (score sign-flipped); the flag applies to all three metrics in that cell. Light blue rows: DCI; light orange rows: FR. Recession base rate = 0.12.

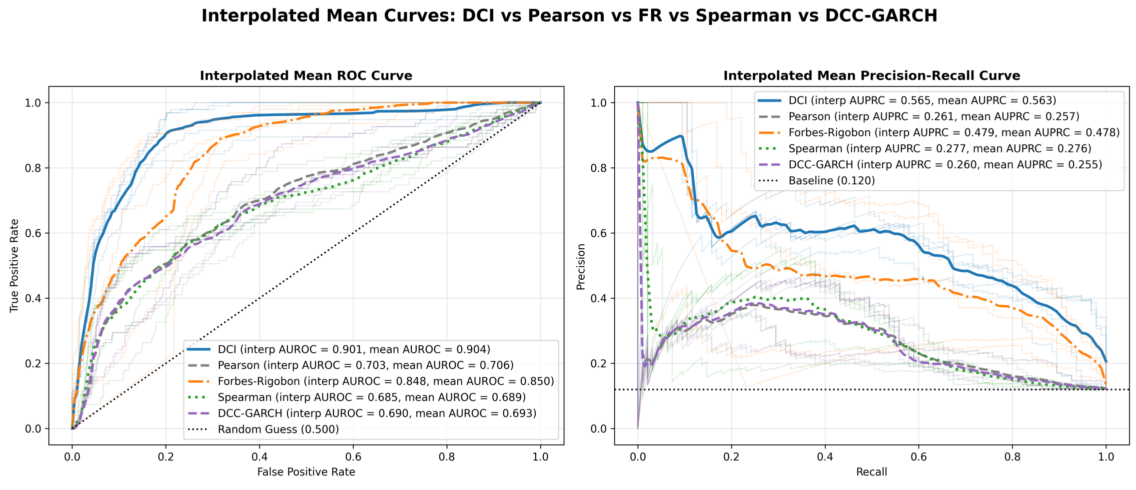

Figure 8. Aggregate ROC and precision–recall performance across macroeconomic graph features. DCI achieves the highest AUROC (0.90) and AUPRC (0.57), exceeding Forbes–Rigobon (0.85, 0.48) and substantially exceeding Pearson, Spearman, and DCC–GARCH, all of which fall in the 0.69–0.70 AUROC and 0.26–0.28 AUPRC range. Recession windows are about 12% of the monthly record, so AUPRC is the more informative aggregate benchmark.

5 Discussion

DCI is a periodicity-aware windowed correlation metric for heterogeneous time series. It sits in a middle ground among existing estimators. It stays nonparametric and costs the same per window as PCC, while the per-signal periodic normalization makes it more robust than plain correlation. Across synthetic networks, EEG under the Berger effect, and the FRED-MD macroeconomic panel, DCI improved on PCC and matched or exceeded domain-specific baselines without tuning.

The three testbeds mark out what the method does. The dynamic-community experiment shows that when coupling is carried by amplitude, phase, or volatility on heterogeneous signals, the periodic normalization lets DCI separate injected modules from the null. The two real-world testbeds show that this survives messy data. On EEG, DCI recovers the Berger response that PCC misses and that the field-standard AEC-c registers with less confidence. On FRED-MD, DCI separates NBER recessions from expansions at AUROC 0.90 and keeps the coupling direction that macroeconomic theory predicts. In the same setting the parametric DCC–GARCH baseline collapses toward sample correlation, and the Forbes–Rigobon adjustment reports decreased coupling during recessions. The structural experiment marks the limit. On a stable linear state-space model, DCI and PCC recover the block structure equally well, and neither recovers the generating operator, because both are correlation estimators.

Many functional-connectivity and connectedness studies build adjacency matrices from windowed PCC, including work in systems neuroscience, climate networks, and systemic-risk monitoring. DCI can be substituted in those pipelines at the same per-window cost, and it recovers amplitude, phase, and volatility coupling that plain correlation misses. The change is local to edge construction, so downstream graph statistics such as community structure, clustering, efficiency, and centrality inherit the improvement unchanged.

5.1 Computational cost

Let d be the panel size, L the record length, w the window length, Δ the stride, and $m = \lfloor (L-w)/\Delta \rfloor + 1$ the number of windows. Sliding-window PCC costs $O(d^2 m w)$, one inner product of length w per pair per window. DCI matches this. Period detection (Appendix B.1) and the phase-resolved normalization of Eq. (2) cost $O(dL)$ once, and the null variance of Eq. (6) adds $O(w)$ per pair. The total is

$$T_{\mathrm{DCI}} \;=\; O(dL) \;+\; O(d^2 m w) \;=\; O\big(d^2 \cdot \tfrac{L}{\Delta} \cdot w\big)\,, \tag{11}$$

with the preprocessing term dominated whenever $d\,m\,w \gtrsim L$. Specialized estimators such as MI, AEC-c, and DCC–GARCH carry higher per-pair costs, heavier preprocessing, or additional parametric assumptions.

5.2 Data requirement

DCI estimates a period and a phase-resolved normalization from the record itself, so it carries a minimum data requirement that plain correlation does not. Appendix B.4 derives this requirement. Two quantities control it: the number of whole periods $\beta_i = \lfloor L/\tau_i \rfloor$ available for estimating the phase moments, and the window length l available for accumulating evidence of coupling. For a coupling event of effect size $\delta^\star = \mathbb{E}[Z_i Z_j \mid H_1]$, Theorem B.6 gives the detection signal-to-noise ratio over a window of length l as

$$\mathrm{SNR}(\widetilde{\mathrm{DCI}}) \;\approx\; \frac{\sqrt{l/v_0}\;\delta^\star}{\sqrt{1 + c_1/\beta_i}}, \tag{12}$$

where v_0 is the per-sample null variance and c_1 is a constant determined by the phase-conditional fourth moment, derived in Lemma B.4 and equal to 3 for Gaussian data.

The two regimes behave differently. Under strong coupling ($\delta^\star \gtrsim 1$) the numerator clears the detection threshold easily, and the binding constraint is that the variance inflation $1 + c_1/\beta_i$ stay small. Keeping half the ideal signal-to-noise ratio then requires $L \gtrsim 6\tau_i$. Under weak coupling ($\delta^\star \ll 1$) the window-length requirement grows as $(\delta^\star)^{-2}$, and the tighter $L \gtrsim 30\tau_i$ is needed (Corollary B.7). Two hard floors sit below both cases (Proposition B.8): $\beta_i \geq 2$, because with a single period the phase-wise variance is zero and Z_i is undefined, and $l \geq 1$. The overall requirement sits between that of the stationary per-pair correlation measures DCI generalizes and that of a fully identified system-identification model.

5.3 Scope and future directions

DCI is an inner product of phase-normalized residuals, so it covers amplitude modulation, phase locking, frequency modulation, and joint volatility, none of which are nonlinear relations. Information-theoretic estimators, conditional estimators, and fully parametric models remain preferable when the coupling is genuinely nonlinear, and extending the phase-normalization step to a nonlinear similarity measure is a direction for future work. DCI is also built for continuous-valued signals and does not yet generalize to discrete-valued series.

DCI reports marginal rather than conditional dependence. It estimates symmetric coupling at zero lag and directed coupling at a fixed lag t_{off}. Correlation is transitive enough that indirect paths still carry edge weight, which is what the structural experiment of Section 3.2 exposes. In a stable linear state-space model the recovered adjacency reflects the full correlation structure, including the downstream influence of every connected node. All correlation-based estimators behave this way, PCC included. Dynamic Bayesian networks [29] and conditional-mutual-information methods [44] are the appropriate tools when conditional-independence structure is the target.

Estimating the normalization from each signal's own structure has one further consequence. A coupling event that occupies a large fraction of the record enters the normalization and dilutes its own apparent strength, so DCI measures coupling relative to the rest of the record. This is why the experiments of Section 3.1 limit injection duration to 10–25% of the record. When a signal is effectively aperiodic, as in chaotic, random-walk, or trend-dominated cases, the period search returns $\tau_i = 1$, the normalization becomes an ordinary z-score, and DCI gains nothing over PCC. These boundaries place DCI as a method for moderate-data, heterogeneous-signal systems whose coupling is linear or close to it.

6 Methods

This section states the experimental design for each result. Derivations, parameter tables, and the definitions of the coupling mechanisms and graph features are deferred to the appendices, as noted in each subsection.

6.1 Signal-to-noise characterization

Three heterogeneous base waveforms of length $L = 3000$ received four injected amplitude-coupling events spanning the four sign combinations of joint scaling. Independent Gaussian noise scaled to the signal variance was added at target SNR values from 30 dB down to -10 dB in 5 dB steps, with 20 Monte Carlo runs per level, and sliding-window DCI, Corr, and $\overline{\mathrm{DCI}}$ were computed with window $w = 60$. Performance is the root-mean-square error between the empirical and noiseless traces. Waveform definitions are in Appendix C.1.

6.2 Dynamic community recovery

Simulations generate $d \in [8, 25]$ nodes over 500 s at 6.0 Hz, giving $L = 3000$ samples. Each base channel is drawn from the heterogeneous waveform family of Eq. (23), and each simulation injects $K \in [2, 6]$ disjoint modules through Eq. (9), each spanning 10–25% of the record. Because the amplitude-driven mechanisms share absolute magnitude as their coupling channel, temporally overlapping amplitude modules are merged into a single ground-truth module before scoring (Appendix C.2.4).

Adjacency matrices were built for every window of length $w = 30$ at stride $\Delta = 6$ using $\widetilde{\mathrm{DCI}}$, sliding-window PCC, and binned MI. Entrywise absolute values were taken and each window was partitioned by Leiden modularity maximization [32]. For each ground-truth module and each active window, the detected cluster with the largest overlap with the ground-truth membership was selected, and confusion-based topological metrics were computed against that membership alongside SCR (Appendix B.6). An injection-free control variant of each method on the same base signals verifies null behavior. The sweeps in Supplement S1 repeat this pipeline while varying $s \in [0, 3]$ and $d \in \{10, \ldots, 100\}$.

6.3 EEG

Dataset and regions. The AlphaWaves dataset [38] provides 16-channel scalp recordings at 512 Hz from 20 subjects alternating ten-second eyes-closed and eyes-open blocks. The occipital region is O1, Oz, and O2; the DMN proxy is Fz, P3, Pz, and P4, overlying posterior-midline DMN hubs; the rest-of-brain group is the remaining nine channels.

Screening and artifact removal. Recordings whose channel RMS amplitudes gave $|Z_{\mathrm{MAD}}| > 3.5$ were excluded, as were subjects whose occipital alpha SNR, defined as 8–12 Hz power divided by 35–45 Hz power under a Welch estimate [45], fell below 2.0. The high-frequency reference band flags muscle artifact, which dominates broadband EEG power above roughly 30 Hz [46, 47]. Thirteen subjects passed both checks. Each retained recording was then decomposed by independent component analysis on PCA-whitened data and components dominated by ocular artifact were removed [48, 49]; Appendix D.4 records the per-subject component counts.

Connectivity and statistics. Records were band-pass filtered to 8–12 Hz with a zero-phase two-pass FIR filter using a Hamming window and MNE-Python default transition bands. Adjacency matrices were computed over 512-sample windows at 51-sample stride using $\widetilde{\mathrm{DCI}}$, absolute PCC, and AEC-c. Per-region node strength and clustering coefficient (Appendix B.5) were compared by paired t-tests across subjects. A within-subject percent-change analysis tested whether the occipital and DMN-proxy increases exceeded the rest-of-brain increase, and an exact conditional permutation test on within-region clustering, with visual channels excluded from the null pool, tested whether the DMN proxy is topologically distinguished from arbitrary four-channel subsets of non-visual cortex. Edge-level network-based statistics [41] provided a complementary global comparison.

6.4 Macroeconomics

Preprocessing. Each indicator was made stationary using the transformation code assigned to it in the FRED-MD documentation [43], which specifies for each series one of no transformation, a first or second difference, a logarithm, a first or second difference of the logarithm, or a first difference of the percentage change. Following the same source, an observation deviating from its series median by more than ten interquartile ranges was set to missing before windowing.

Network construction. A window of $w = 12$ months at stride $\Delta = 1$ month was slid across the panel, and five adjacency matrices were built per window. Pearson and Spearman are the standard windowed estimators. FR is the heteroskedasticity-adjusted Pearson correlation of Forbes and Rigobon [28], dividing by $\sqrt{1 + \delta_i(1 - \rho^2)}$ with δ_i the within-window variance ratio. DCC–GARCH follows the two-step procedure of Engle [19], fitting a univariate GARCH(1,1) per node with the `arch` package and then applying the DCC recursion with canonical scalar defaults, as stated in Appendix C.4, Eq. (24). The fifth is $\widetilde{\mathrm{DCI}}$. Each window midpoint was aligned to the NBER USREC indicator, self-loops were zeroed, entrywise absolute values were taken, and edges below 0.05 were removed.

Features and statistics. Five global features were extracted from each thresholded matrix: density, Louvain modularity [50], transitivity, Shannon entropy of normalized node strengths, and weighted global efficiency, all defined in Appendix B.5. Welch's t-test compared recession against expansion windows, and retrospective separability used each feature as a single-feature recession score with the sign chosen automatically.

References

[1] Proulx, S. R., Promislow, D. E. L., and Phillips, P. C. (2005). Network thinking in ecology and evolution. *Trends in Ecology & Evolution*, 20(6):345–353.

[2] Rubinov, M. and Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage*, 52(3):1059–1069.

[3] Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45(2):167–256.

[4] Tsonis, A. A. and Roebber, P. J. (2004). The architecture of the climate network. *Physica A*, 333:497–504.

[5] Donges, J. F., Zou, Y., Marwan, N., and Kurths, J. (2009). Complex networks in climate dynamics. *European Physical Journal Special Topics*, 174(1):157–179.

[6] Diebold, F. X. and Yilmaz, K. (2015). *Financial and Macroeconomic Connectedness: A Network Approach to Measurement and Monitoring.* Oxford University Press.

[7] Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198.

[8] Hevey, D. (2018). Network analysis: a brief overview and tutorial. *Health Psychology and Behavioral Medicine*, 6(1):301–328.

[9] Keeling, M. J. and Eames, K. T. D. (2005). Networks and epidemic models. *Journal of the Royal Society Interface*, 2(4):295–307.

[10] Diebold, F. X. and Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal*, 119(534):158–171.

[11] Diebold, F. X. and Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1):57–66.

[12] Diebold, F. X. and Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134.

[13] Pearson, K. (1895). Note on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58:240–242.

[14] Spearman, C. (1904). The proof and measurement of association between two things. *American Journal of Psychology*, 15(1):72–101.

[15] Kraskov, A., Stögbauer, H., and Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, 69:066138.

[16] Li, F. and Zhu, H. (2014). Testing for financial contagion based on a nonparametric measure of the cross-market correlation. *Review of Financial Economics*, 23(3):141–147.

[17] Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 48(1):1–48.

[18] Ljung, L. (1999). *System Identification: Theory for the User.* Prentice Hall, 2nd edition.

[19] Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350.

[20] Starkey, C. M. and Tsafack, G. (2023). Measuring financial contagion: Dealing with the volatility bias in the correlation dynamics. *International Review of Financial Analysis*, 90:102863.

[21] Kuramoto, Y. (1975). Self-entrainment of a population of coupled non-linear oscillators. In *International Symposium on Mathematical Problems in Theoretical Physics*, pages 420–422.

[22] Mohan, J., Ramsundar, B., and Subramanian, S. (2025). Inferring dynamic hidden graph structure in heterogeneous correlated time series. In *BayLearn Machine Learning Symposium*. arXiv:2512.01301.

[23] Brookes, M. J., Woolrich, M. W., and Barnes, G. R. (2012). Measuring functional connectivity in MEG: A multivariate approach insensitive to linear source leakage. *NeuroImage*, 63(2):910–920.

[24] Hipp, J. F., Hawellek, D. J., Corbetta, M., Siegel, M., and Engel, A. K. (2012). Large-scale cortical correlation structure of spontaneous oscillatory activity. *Nature Neuroscience*, 15(6):884–890.

[25] Colclough, G. L., Brookes, M. J., Smith, S. M., and Woolrich, M. W. (2015). A symmetric multivariate leakage correction for MEG connectomes. *NeuroImage*, 117:439–448.

[26] O'Mahony, J., Niso, G., and Tadel, F. (2024). Amplitude envelope correlation as a dominant connectivity metric for large-scale brain networks: review and recommendations. *NeuroImage*, in press.

[27] Gardner, W. A., Napolitano, A., and Paura, L. (2006). Cyclostationarity: Half a century of research. *Signal Processing*, 86(4):639–697.

[28] Forbes, K. J. and Rigobon, R. (2002). No contagion, only interdependence: Measuring stock market comovements. *Journal of Finance*, 57(5):2223–2261.

[29] Murphy, K. P. (2002). *Dynamic Bayesian Networks: Representation, Inference and Learning.* PhD thesis, University of California, Berkeley.

[30] Friedman, N., Murphy, K., and Russell, S. (1998). Learning the structure of dynamic probabilistic networks. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 139–147.

[31] Bartlett, M. S. (1946). On the theoretical specification and sampling properties of autocorrelated time-series. *Supplement to the Journal of the Royal Statistical Society*, 8(1):27–41.

[32] Traag, V. A., Waltman, L., and van Eck, N. J. (2019). From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports*, 9:5233.

[33] Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 87(1):527–570.

[34] Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: A review and analysis. *Brain Research Reviews*, 29(2–3):169–195.

[35] Palva, S. and Palva, J. M. (2007). New vistas for alpha-frequency band oscillations. *Trends in Neurosciences*, 30(4):150–158.

[36] 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. *Proceedings of the National Academy of Sciences*, 98(2):676–682.

[37] Buckner, R. L., Andrews-Hanna, J. R., and Schacter, D. L. (2008). The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, 1124(1):1–38.

[38] Cattan, G., Rodrigues, P. L. C., and Congedo, M. (2018). EEG alpha waves dataset. *Zenodo*, doi:10.5281/zenodo.2348892.

[39] Pfurtscheller, G. and Lopes da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: Basic principles. *Clinical Neurophysiology*, 110(11):1842–1857.

[40] Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., and Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, 102(27):9673–9678.

[41] Zalesky, A., Fornito, A., and Bullmore, E. T. (2010). Network-based statistic: Identifying differences in brain networks. *NeuroImage*, 53(4):1197–1207.

[42] James, N., Menzies, M., and Chin, K. (2022). On financial market correlation structures and diversification benefits across and within equity sectors. *Physica A: Statistical Mechanics and its Applications*, 604:127682.

[43] McCracken, M. W. and Ng, S. (2016). FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589.

[44] Runge, J., Nowack, P., Kretschmer, M., Flaxman, S., and Sejdinovic, D. (2019). Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5(11):eaau4996.

[45] Welch, P. D. (1967). The use of fast Fourier transform for the estimation of power spectra. *IEEE Transactions on Audio and Electroacoustics*, 15(2):70–73.

[46] Goncharova, I. I., McFarland, D. J., Vaughan, T. M., and Wolpaw, J. R. (2003). EMG contamination of EEG: spectral and topographical characteristics. *Clinical Neurophysiology*, 114(9):1580–1593.

[47] Muthukumaraswamy, S. D. (2013). High-frequency brain activity and muscle artifacts in MEG/EEG: a review and recommendations. *Frontiers in Human Neuroscience*, 7:138.

[48] Hyvärinen, A. and Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural Networks*, 13(4–5):411–430.

[49] Pion-Tonachini, L., Kreutz-Delgado, K., and Makeig, S. (2019). ICLabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, 198:181–197.

[50] Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008.

[51] Reshef, D. N., Reshef, Y. A., Finucane, H. K., Grossman, S. R., McVean, G., Turnbaugh, P. J., Lander, E. S., Mitzenmacher, M., and Sabeti, P. C. (2011). Detecting novel associations in large data sets. *Science*, 334(6062):1518–1524.

[52] Székely, G. J., Rizzo, M. L., and Bakirov, N. K. (2007). Measuring and testing dependence by correlation of distances. *Annals of Statistics*, 35(6):2769–2794.

[53] Marrelec, G., Krainik, A., Duffau, H., Pélégrini-Issac, M., Lehéricy, S., Doyon, J., and Benali, H. (2006). Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage*, 32(1):228–237.

[54] Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438.

[55] Marinazzo, D., Pellicoro, M., and Stramaglia, S. (2008). Kernel method for nonlinear Granger causality. *Physical Review Letters*, 100:144103.

[56] Tank, A., Covert, I., Foti, N., Shojaie, A., and Fox, E. B. (2021). Neural Granger causality. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4267–4279.

[57] Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85:461–464.

[58] Sugihara, G., May, R., Ye, H., Hsieh, C.-h., Deyle, E., Fogarty, M., and Munch, S. (2012). Detecting causality in complex ecosystems. *Science*, 338(6106):496–500.

[59] Lachaux, J.-P., Rodriguez, E., Martinerie, J., and Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8(4):194–208.

[60] Bastos, A. M. and Schoffelen, J.-M. (2016). A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*, 9:175.

[61] Torrence, C. and Compo, G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*, 79(1):61–78.

[62] Grinsted, A., Moore, J. C., and Jevrejeva, S. (2004). Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Processes in Geophysics*, 11(5/6):561–566.

[63] Friston, K. J., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, 19(4):1273–1302.

[64] Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *Review of Economics and Statistics*, 72(3):498–505.

[65] Ibragimov, I. A. (1962). Some limit theorems for stationary processes. *Theory of Probability and Its Applications*, 7(4):349–382.

[66] Newman, M. E. J. (2004). Analysis of weighted networks. *Physical Review E*, 70:056131.

[67] Newman, M. E. J. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582.

[68] Onnela, J.-P., Saramäki, J., Kertész, J., and Kaski, K. (2005). Intensity and coherence of motifs in weighted complex networks. *Physical Review E*, 71:065103.

A Related Work and the Method Landscape

This appendix reviews the estimators shown in Fig. 2 and says what each one recovers and what it costs. The horizontal axis of that figure is computational and data efficiency and the vertical axis is coupling complexity and robustness. Most methods sit near a trade-off curve on which robustness is bought with data and computation. DCI is designed to sit slightly off that curve, because its added robustness comes from one normalization step whose cost is linear in the record length.

A.1 Bivariate nonparametric correlation

The Pearson correlation coefficient [13], in windowed and lagged form, is the default edge estimator in every field cited in Section 1. It is the cheapest option and the least robust. It measures linear covariation on the raw scale, so it is sensitive to differences in units, trends, periodic structure, and volatility between the two signals. Spearman rank correlation [14] replaces values with ranks. This buys invariance to monotone transformation for the cost of a sort and captures monotone nonlinear association. Neither method addresses differences in periodicity, which is what DCI targets.

Binned or nearest-neighbor mutual information [15] captures dependence of any functional form, but it needs a density estimate whose variance grows quickly as the window shortens. Short windows are exactly the regime used for dynamic graphs, and this variance is the dominant error there, as the results in Section 3.1 show. The maximal information coefficient [51] searches over grid resolutions to produce an equitable dependence score, and inherits the same small-sample problem at a higher per-pair cost. Distance correlation [52] is zero if and only if the two variables are independent and needs no density estimate, but its per-pair cost is quadratic in the window length, which places it well to the left of the simple correlations.

Partial correlation [53] conditions each pair on the remaining channels through the inverse covariance matrix. This removes the indirect-path bleed-through discussed in Section 3.2, but it needs the window to be long relative to the panel size for the inverse to be well conditioned, which is restrictive at the panel sizes used here.

A.2 Model-based and directed estimators

Linear Granger causality [54] tests whether the past of one series improves prediction of another under a vector autoregression [17]. It gives a directed edge, and it costs a fitted lag-order model. Kernel and other nonlinear Granger variants [55] run the same test in a reproducing-kernel space and capture nonlinear predictive structure with correspondingly larger data requirements. Neural Granger causality [56] replaces the predictive model with a sparse-input neural network and is the most expressive and most data-hungry method in Fig. 2.

Transfer entropy [57] measures directed information flow without assuming a functional form. It carries the density-estimation cost of mutual information plus the cost of conditioning on the target's own past. PCMCI [44] combines a condition-selection stage with a momentary conditional-independence test and recovers time-lagged causal graphs at panel scale, which makes it the standard choice when the target is a causal graph rather than a coupling weight. Dynamic Bayesian networks [29, 30] recover conditional-independence structure across time by structure search, and state-space system identification [18] recovers a latent linear operator. Both need long records relative to the number of parameters.

Convergent cross mapping [58] tests for causal coupling in deterministic dynamical systems by asking whether the attractor reconstructed from one series predicts the states of another. It suits short, noisy ecological records from low-dimensional deterministic systems and does not target the linear coupling weight that DCI estimates.

A.3 Domain-specific estimators

Neuroscience contributes several estimators built around a specific signal model. The phase-locking value [59] discards amplitude and measures how concentrated the instantaneous phase difference between two band-limited signals is. Amplitude envelope correlation and its leakage-corrected form AEC-c [23–26] do the opposite, correlating Hilbert amplitude envelopes after orthogonalization to suppress volume-conduction leakage, and they are the dominant connectivity metric in EEG and MEG. Bastos and Schoffelen [60] review both alongside the rest of the functional-connectivity toolkit. Wavelet coherence [61, 62] localizes coherence in time and frequency together and is standard in geophysics and climate science. Dynamic causal modelling [63] fits a generative biophysical model of neuronal populations and is the most parameter-heavy method in the figure.

In finance and macroeconomics, the Forbes–Rigobon adjustment [28] corrects the correlation coefficient for the heteroskedasticity that inflates apparent comovement during crises. Its behavior on the FRED-MD panel is analyzed in Section 4.2. Conditional correlation GARCH models [19, 64] instead model the conditional covariance directly, and DCC–GARCH is the standard heteroskedasticity-aware baseline. Its two-step estimator scales poorly in the panel dimension, and at the panel sizes used here the correlation step regularizes toward the sample correlation, which is the failure mode documented in Section 4.2. Li and Zhu [16] and Starkey and Tsafack [20] review the volatility bias in contagion measurement.

A.4 Position of DCI

DCI is bivariate and nonparametric, and it is directed only through the explicit lag t_{off}. Its per-window cost matches Pearson correlation (Section 5.1). Its one extra requirement is enough record length to estimate the phase-resolved moments of Eq. (1), which Appendix B.4 quantifies. The robustness it gains over Pearson correlation comes from removing each signal's own periodic profile before comparison, which is a preprocessing step rather than a richer model of the coupling. DCI therefore does not recover conditional independence, nonlinear functional dependence, or a latent generating operator, and the methods above remain the right tools when those are the target.

Cyclostationary signal analysis [27] uses a similar periodic-moment construction but applies it differently. It estimates correlations that are themselves periodic functions of time, while DCI estimates one correlation after the periodic structure has been normalized away.

B Mathematical Details

This appendix contains the derivations underlying Section 2. Notation is that of Section 2.1: X_i denotes a channel of length L, τ_i its estimated period, $\beta_i = \lfloor L/\tau_i \rfloor$ the number of whole periods contained in the record, Z_i the normalized channel defined by Eq. (2), $\mathcal{W} = [t_1, t_2]$ a window of length l, and t_{off} the lag between the two channels of a pair.

B.1 Estimating the period

Let $\hat{\mu}_i$ be the sample mean of X_i and

$$\hat{\rho}_i(\tau) \;=\; \frac{\sum_{t=1}^{L-\tau} \big(X_i(t) - \hat{\mu}_i\big)\big(X_i(t+\tau) - \hat{\mu}_i\big)}{\sum_{t=1}^{L} \big(X_i(t) - \hat{\mu}_i\big)^2} \tag{13}$$

denote its sample autocorrelation at lag τ. Under a white-noise null, $|\hat{\rho}_i(\tau)| \le 1.96/\sqrt{L}$ with probability 0.95 [31]. The period estimate τ_i is defined as the smallest lag $\tau \ge 2$ that exceeds this band and is a local maximum of $\hat{\rho}_i$,

$$\tau_i \;=\; \min\big\{\tau \ge 2 \;:\; |\hat{\rho}_i(\tau)| > 1.96/\sqrt{L} \;\text{ and }\; \hat{\rho}_i(\tau-1) < \hat{\rho}_i(\tau) > \hat{\rho}_i(\tau+1)\big\}, \tag{14}$$

with $\tau_i = 1$ if no such lag exists. The case $\tau_i = 1$ is well defined: by Proposition B.1 it reduces Eq. (2) to a global z-score, so an aperiodic channel contributes to DCI as it would to a Pearson-based estimator, with the normalization conferring no additional robustness.

B.2 Reduction to Pearson correlation

Proposition B.1. *If $\tau_i = \tau_j = 1$, $t_{\text{off}} = 0$, and $\mathcal{W} = [1, L]$, then $\mathrm{DCI}^{\mathcal{W}}_{i,j} = L\rho_{X_i,X_j}$, $D^{\mathcal{W}}_{i,j} = L$, and $\mathrm{Corr}^{\mathcal{W}}_{i,j} = \rho_{X_i,X_j}$, where ρ_{X_i,X_j} is the sample Pearson correlation over the whole record.*

Proof. With $\tau_i = 1$ there is one within-period position, so Eq. (1) reduces to the sample mean $\hat{\mu}_i$ and sample standard deviation $\hat{\sigma}_i$ of the whole channel, and Eq. (2) becomes $Z_i(t) = (X_i(t) - \hat{\mu}_i)/\hat{\sigma}_i$. Substituting into Eq. (3),

$$\mathrm{DCI}^{\mathcal{W}}_{i,j} = \sum_{t=1}^{L} \frac{(X_i(t) - \hat{\mu}_i)(X_j(t) - \hat{\mu}_j)}{\hat{\sigma}_i \hat{\sigma}_j} = L\,\rho_{X_i,X_j}.$$

Because $\hat{\sigma}_i$ is the sample standard deviation, $\sum_t Z_i(t)^2 = L$ and likewise for j, so Eq. (4) gives $D^{\mathcal{W}}_{i,j} = L$ and the ratio is ρ_{X_i,X_j}. □

B.3 The null variance

Write $\mathrm{DCI}_{i,j}^{\mathcal{W}} = \sum_{t=t_1}^{t_2} Y_t$ with $Y_t = Z_i(t)Z_j(t + t_{\mathrm{off}})$. The autocorrelation of a normalized channel at lag τ is

$$R_{Z_iZ_i}(\tau) \;=\; \mathbb{E}\big[Z_i(t)\, Z_i(t+\tau)\big]. \tag{15}$$

For weakly stationary Z_i this quantity is independent of t and satisfies $R_{Z_iZ_i}(0) = 1$. Since Z_i has zero mean and unit variance, $R_{Z_iZ_i}(\cdot)$ is simultaneously the autocovariance and the autocorrelation function of Z_i. The quantity $R_{Z_jZ_j}(\cdot)$ is defined analogously.

The null hypothesis H_0 states that channels i and j are uncoupled. Under H_0 it is assumed that (i) Z_i and Z_j are weakly stationary, (ii) Z_i and Z_j are mutually independent, and (iii) $\sum_{\tau=1}^{\infty} |R_{Z_iZ_i}(\tau)| < \infty$, and likewise for Z_j.

Lemma B.2. *For every channel i and every within-period position k, the β_i samples of Z_i at position k have sample mean 0 and sample second moment 1. Consequently $\mathbb{E}[Z_i(t)] = 0$ and $\mathbb{E}[Z_i(t)^2] = 1$.*

Proof. Fix k and set $x_u = X_i(k + u\tau_i)$ for $u = 0, \ldots, \beta_i - 1$. The scale $\sigma_{X,i}(k)$ is constant over these samples, so by the definition of $\mu_{X,i}(k)$ in Eq. (1),

$$\frac{1}{\beta_i}\sum_u Z_i(k+u\tau_i) = \frac{1}{\sigma_{X,i}(k)}\left(\frac{1}{\beta_i}\sum_u x_u - \mu_{X,i}(k)\right) = 0.$$

Expanding the square and applying the definition of $\sigma_{X,i}(k)$ in the same equation,

$$\frac{1}{\beta_i}\sum_u Z_i(k+u\tau_i)^2 = \frac{\frac{1}{\beta_i}\sum_u x_u^2 - \mu_{X,i}(k)^2}{\sigma_{X,i}(k)^2} = 1.$$

Averaging both identities over the τ_i within-period positions yields the stated expectations. □

Theorem B.3. *Under H_0, for a window of length l,*

$$\mathrm{Var}\big(\mathrm{DCI}_{i,j}^{\mathcal{W}}\big) \;=\; l \;+\; 2\sum_{\tau=1}^{l-1}(l-\tau)\, R_{Z_iZ_i}(\tau)\, R_{Z_jZ_j}(\tau).$$

Proof. By assumption (ii) and Lemma B.2, $\mathbb{E}[Y_t] = \mathbb{E}[Z_i(t)]\,\mathbb{E}[Z_j(t+t_{\mathrm{off}})] = 0$ and $\mathrm{Var}(Y_t) = \mathbb{E}[Z_i(t)^2]\,\mathbb{E}[Z_j(t+t_{\mathrm{off}})^2] = 1$. For $s < t$ with $\tau = t - s$, the same independence factorizes the covariance,

$$\mathrm{Cov}(Y_s, Y_t) = \mathbb{E}\big[Z_i(s)Z_i(t)\big]\,\mathbb{E}\big[Z_j(s+t_{\mathrm{off}})Z_j(t+t_{\mathrm{off}})\big] = R_{Z_iZ_i}(\tau)\, R_{Z_jZ_j}(\tau).$$

Expanding the variance of the sum gives $\sum_t \mathrm{Var}(Y_t) + 2\sum_{s<t}\mathrm{Cov}(Y_s, Y_t)$. The first term equals l. In the second, the pairs are grouped by their separation $\tau = t - s$: there are exactly $l - \tau$ pairs at separation τ, each contributing $R_{Z_iZ_i}(\tau)R_{Z_jZ_j}(\tau)$. Summing over $\tau = 1, \ldots, l-1$ yields the stated expression. □

Dividing the expression in Theorem B.3 by l and letting $l \to \infty$ defines the per-sample null variance

$$v_0 \;=\; 1 \;+\; 2\sum_{\tau=1}^{\infty} R_{Z_iZ_i}(\tau)\, R_{Z_jZ_j}(\tau). \tag{16}$$

The series converges by assumption (iii). If the samples are uncorrelated then $v_0 = 1$; autocorrelation in either channel increases v_0, so a longer window is required to attain a given detection power. The sequence $\{Y_t\}$ is stationary and weakly dependent, so by the central limit theorem for such sequences [65], $\mathrm{DCI}_{i,j}^{\mathcal{W}}$ is asymptotically Gaussian under H_0 with the variance given in Theorem B.3. Equations (7) and (8) follow.

B.4 Data requirements

Since the normalization in Eq. (2) is estimated from the record, DCI imposes a minimum data requirement that the Pearson correlation does not. The requirement is governed by two quantities: the number of periods β_i available for estimating the phase moments in Eq. (1), and the window length l over which evidence of coupling accumulates. A bound on each is derived below.

Let $\mu_i^\star(k)$ and $\sigma_i^\star(k)$ denote the true mean and standard deviation of X_i at within-period position k, and let $Z_i^\star(t) = (X_i(t) - \mu_i^\star(k))/\sigma_i^\star(k)$ denote the ideal normalized channel obtained when these are known. Assume that, for each position k, the β_i samples of X_i at that position are independent and identically distributed across periods; the variance calculation in Lemma B.4 assumes further that they are Gaussian. Under a coupling alternative H_1 the effect size is $\delta^\star = \mathbb{E}[Z_i^\star(t)\, Z_j^\star(t + t_{\mathrm{off}}) \mid H_1]$.

Lemma B.4. *Let $a_i = \beta_i^{-1}\sum_u Z_i^\star(k+u\tau_i)$ denote the error in the estimated phase mean, and let $1+b_i$ denote the ratio of the estimated to the true phase variance. Then*

$$Z_i(t) \;=\; \frac{Z_i^\star(t)-a_i}{\sqrt{1+b_i}}, \tag{17}$$

and, for Gaussian data,

$$\mathbb{E}\big[Z_i(t)Z_j(t+t_{\text{off}}) \mid H_1\big] = \delta^\star + O(\beta^{-1}), \qquad \operatorname{Var}\big(Y_t \mid H_0\big) = \Big(1+\tfrac{c_1}{\beta_i}\Big)\Big(1+\tfrac{c_1}{\beta_j}\Big) + O(\beta^{-2}), \tag{18}$$

with $c_1 = 3$ and $\beta = \min(\beta_i, \beta_j)$. Estimation of the phase moments therefore preserves the coupling signal to first order and inflates the null variance by a term of order β^{-1}.

Proof. By Eq. (1), $\mu_{X,i}(k)$ and $\sigma_{X,i}(k)^2$ are the sample mean and sample variance of the β_i values at position k, whence $\mu_{X,i}(k) = \mu_i^\star(k) + \sigma_i^\star(k)\,a_i$ and $\sigma_{X,i}(k)^2 = \sigma_i^\star(k)^2(1+b_i)$. Substituting both into Eq. (2) and cancelling $\sigma_i^\star(k)$ gives Eq. (17).

For the first claim, expand $(1+b_i)^{-1/2} = 1 - \frac{1}{2}b_i + O(b_i^2)$ and take the product of the two decompositions. The leading term is $\mathbb{E}[Z_i^\star Z_j^\star] = \delta^\star$. Each remaining term contains at least one of a_i, a_j, b_i, b_j. Each of these is an average over β samples and hence of magnitude $O(\beta^{-1/2})$, so every term in which one appears together with a mean-zero factor is $O(\beta^{-1})$.

For the second claim, Lemma B.2 states that the empirical second moment of Z_i at each within-period position is exactly one, so a_i does not contribute to $\mathbb{E}[Z_i(t)^2]$ and only the scale error b_i remains. For Gaussian samples at position k the sample variance satisfies $\beta_i(1+b_i) \sim \chi^2_{\beta_i-1}$, whence

$$\mathbb{E}\Big[\frac{1}{1+b_i}\Big] = \beta_i\,\mathbb{E}\Big[\frac{1}{\chi^2_{\beta_i-1}}\Big] = \frac{\beta_i}{\beta_i-3} = 1 + \frac{3}{\beta_i} + O(\beta_i^{-2}).$$

By assumption (ii), $\operatorname{Var}(Y_t) = \mathbb{E}[Z_i(t)^2]\,\mathbb{E}[Z_j(t+t_{\text{off}})^2]$, and each factor equals the corresponding expectation above. This establishes Eq. (18) with $c_1 = 3$. □

Lemma B.5. *Consider the one-sided test that rejects H_0 when $\mathrm{DCI}_{i,j}^{\mathcal{W}} > z_{1-\alpha}\sqrt{l\,v_0}$. This test has level α and power at least $1-\gamma$ against an alternative of effect size $\delta^\star > 0$ provided*

$$l \;\geq\; \frac{\big(z_{1-\alpha}+z_{1-\gamma}\big)^2 v_0}{(\delta^\star)^2}. \tag{19}$$

Proof. Under H_0 the statistic is asymptotically $\mathcal{N}(0, lv_0)$, so the stated critical value attains level α. Under H_1 its mean is $l\delta^\star$ by Lemma B.4, and its standard deviation is $\sqrt{lv_0}$ to leading order. Power at least $1-\gamma$ requires the mean to exceed the critical value by $z_{1-\gamma}$ standard deviations,

$$l\,\delta^\star - z_{1-\alpha}\sqrt{lv_0} \;\geq\; z_{1-\gamma}\sqrt{lv_0}.$$

Dividing by $\sqrt{l}$ and squaring yields Eq. (19). □

Theorem B.6. *Let $r \in (0,1)$ be a tolerance on the fractional inflation of the null variance contributed by each channel. Then DCI detects a coupling event of effect size $\delta^\star$ at level α and power $1-\gamma$ provided*

$$\beta_i \;\geq\; \frac{c_1}{r} \quad\iff\quad L \;\geq\; \frac{c_1}{r}\,\tau_i, \qquad \text{and} \qquad l \;\geq\; \frac{(1+r)^2\big(z_{1-\alpha}+z_{1-\gamma}\big)^2 v_0}{(\delta^\star)^2}, \tag{20}$$

with the same condition on β_j. The corresponding detection signal-to-noise ratio is

$$\mathrm{SNR}(\widetilde{\mathrm{DCI}}) \;=\; \frac{\mathbb{E}[\mathrm{DCI} \mid H_1]}{\sqrt{\operatorname{Var}(\mathrm{DCI} \mid H_0)}} \;\approx\; \frac{\sqrt{l/v_0}\;\delta^\star}{\sqrt{1+c_1/\beta_i}}. \tag{21}$$

Proof. By Lemma B.4 channel i inflates the per-sample null variance by the factor $1 + c_1/\beta_i$, which is at most $1+r$ if and only if $\beta_i \geq c_1/r$. Since $\beta_i = \lfloor L/\tau_i \rfloor$, this is equivalent to the stated condition on L. With both channels at tolerance r, the effective per-sample variance is at most $v_0(1+r)^2$; substituting this into Lemma B.5 yields the window condition. Equation (21) follows on dividing the mean $l\delta^\star$ of Lemma B.4 by the standard deviation $\sqrt{lv_0(1+c_1/\beta_i)}$. □

Corollary B.7. *Let* $c_1 = 3$*. Retention of half the ideal signal-to-noise ratio corresponds to* $r = 0.5$ *and yields* $\beta_i \geq 6$*, equivalently* $L \geq 6\tau_i$*. This is the binding condition in the strong-coupling regime* $\delta^\star \gtrsim 1$*, in which the window condition of Eq.* (20) *is satisfied by a short window. In the weak-coupling regime* $\delta^\star \ll 1$ *the window condition grows as* $(\delta^\star)^{-2}$ *and the tolerance must be tightened;* $r = 0.1$ *yields* $\beta_i \geq 30$*, equivalently* $L \geq 30\tau_i$*.*

Proposition B.8. *For every effect size, DCI requires* $\beta_i \geq 2$ *and* $l \geq 1$*. If* $\beta_i = 1$ *then* $\sigma_{X,i}(k) = 0$ *for every within-period position* k *and* Z_i *is undefined.*

Proof. If $\beta_i = 1$, each within-period position contains a single sample x, so Eq. (1) gives $\mu_{X,i}(k) = x$ and $\sigma_{X,i}(k) = \sqrt{x^2 - x^2} = 0$, and the denominator of Eq. (2) vanishes. The bound $l \geq 1$ is immediate from Eq. (3). □

The benchmark of Appendix C.1 lies in the strong-coupling regime: the injected events have magnitude near unity and occupy a fifth of the record, and the normalization is formed from $\beta = 5$ periods. This exceeds the floor of Proposition B.8 and approaches the condition $\beta_i \geq 6$ of Corollary B.7, and Eq. (21) predicts a peak normalized DCI of the observed magnitude. Theorem B.6 implies three failure modes: $\beta_i = 1$, which is structural; $\delta^\star \to 0$, in which case no effect is present; and noise sufficient to reduce the effective $\delta^\star$ below the threshold of Lemma B.5, which corresponds to the low-SNR end of Fig. 3.

B.5 Graph-theoretic definitions

All definitions follow the Brain Connectivity Toolbox [2]. Adjacency matrices are weighted, symmetric, and non-negative, with self-loops set to zero and entries satisfying $|A_{ij}| < 0.05$ removed. Throughout, d denotes the number of nodes, $k_i = \sum_j A_{ij}$ the strength of node i, k_i^{deg} its number of nonzero neighbors, and $\omega = \frac{1}{2}\sum_{ij} A_{ij}$ the total edge weight.

Density. $\text{Density}(A) = \dfrac{1}{d(d-1)} \sum_{i \neq j} A_{ij}$.

Weighted modularity. For a partition assigning node i to community c_i [66, 67],

$$Q \;=\; \frac{1}{2\omega} \sum_{i,j} \Big[A_{ij} - \frac{k_i k_j}{2\omega} \Big] \delta(c_i, c_j),$$

with partitions from Louvain optimization [50] for the macroeconomic features and Leiden optimization [32] for the community-recovery experiments.

Weighted clustering and transitivity. With $\widehat{A}_{ij} = A_{ij} / \max_{ab} A_{ab}$ [68],

$$C_i = \frac{1}{k_i^{\text{deg}}\big(k_i^{\text{deg}} - 1\big)} \sum_{j,h} \big(\widehat{A}_{ij}\widehat{A}_{jh}\widehat{A}_{hi}\big)^{1/3}, \qquad T = \frac{\sum_{i,j,h} \big(\widehat{A}_{ij}\widehat{A}_{jh}\widehat{A}_{hi}\big)^{1/3}}{\sum_i k_i^{\text{deg}}\big(k_i^{\text{deg}} - 1\big)}.$$

Entropy of node strengths. $H = -\sum_i s_i \log_2 s_i$ with $s_i = k_i / \sum_j k_j$ and $0 \log_2 0 = 0$.

Weighted global efficiency. With edge lengths $L_{ij} = 1/A_{ij}$, infinite where $A_{ij} = 0$, and shortest-path distances d_{ij} from Dijkstra's algorithm, $E_{\text{glob}} = \dfrac{1}{d(d-1)} \sum_{i \neq j} 1/d_{ij}$.

B.6 Significant Community Recall

Let $C_c(t)$ denote the mean weighted clustering coefficient over the node set M_c of ground-truth module c in window t, computed on the native linear edge scale and evaluated on the merged node set when amplitude-driven modules overlap (Appendix C.2.4). Let $\mathcal{B}_c$ denote the set of windows in which module c is inactive, with mean μ_c and standard deviation $\hat{\sigma}_c$ of C_c over that set, and let $\mathcal{A}_c$ denote the set of windows in which it is active. For $t \in \mathcal{A}_c$, the standardized shift $\big(C_c(t) - \mu_c\big)/\hat{\sigma}_c$ yields a right-tailed normal p-value $p_c(t)$, and window t is declared significant when $p_c(t) < 0.10$. The duration of significant clustering is the fraction of active windows declared significant.

Let $r_c(t) \in [0,1]$ denote the recall, in window t, of the detected community against the ground-truth membership of module c, the detected community being the Leiden cluster of largest overlap with M_c. Significant Community Recall weights each significant active window by this recall,

$$\mathrm{SCR}_c \;=\; \frac{1}{|\mathcal{A}_c|} \sum_{t \in \mathcal{A}_c} r_c(t)\, \mathbf{1}\big[p_c(t) < 0.10\big], \tag{22}$$

and the reported value of SCR is the mean of SCR_c over the ground-truth modules.

C Experimental Details

This appendix gives the signal models and estimator settings used by the experiments in Sections 2–4.

C.1 Amplitude-coupling benchmark

The pairwise benchmark of Section 2.6 uses three base waveforms of length $L = 3000$: S_1 a sine of period 150, S_2 a repeating inverse Gaussian of period 240, and S_3 a three-component sine. Each pair receives four injected amplitude-coupling events covering the sign combinations $(+,+)$, $(+,-)$, $(-,+)$, and $(-,-)$ of joint scaling, with each event occupying a fifth of the record. Noise is added at target SNR from 30 dB to -10 dB in 5 dB steps, with 20 Monte Carlo runs per level and edge window $w = 60$. The same benchmark is used for the parameter sweeps of Appendix D.2.

C.2 Synthetic network simulator

C.2.1 Base waveforms

Each base channel is drawn from a heterogeneous family. With phase $\phi \sim \mathcal{U}[0, 2\pi)$ and frequency $\omega \sim \mathcal{U}[0.1, 3.5]$ Hz,

$$X_i^{(0)}(t) \;\sim\; \begin{cases} \sin(2\pi\omega t + \phi), & \text{simple sine,} \\ \sum_{p=1}^{3} \alpha_p \sin(2\pi\omega_p t + \phi_p), & \text{multi-frequency sine,} \\ \operatorname{sgn}\big(\sin(2\pi\omega t + \phi)\big), & \text{square wave,} \end{cases} \tag{23}$$

with the three families equally likely. With probability 0.3 the drawn channel is replaced by the sum of two independent draws from Eq. (23). The heterogeneity is deliberate and matches the setting DCI targets.

C.2.2 Coupling channels

Four kinds of coupling are relevant to the injections below. Two channels show *amplitude coupling* when their analytic-signal magnitudes $a_i(t) = |\mathcal{H}\{X_i\}(t)|$ covary after marginal normalization. They show *phase coupling* when the difference of their instantaneous phases $\phi_i(t) = \arg \mathcal{H}\{X_i\}(t)$ concentrates around a preferred value, which is what the phase-locking value measures [59]. They show *frequency coupling* when their instantaneous frequencies $\mathrm{d}\phi_i/\mathrm{d}t$ covary. They show *joint volatility coupling* when both deviate from their normalized baselines at the same time regardless of sign, which is what the combined volatility $D_{i,j}^{\mathcal{W}}$ of Eq. (4) measures and what the Forbes–Rigobon and DCC–GARCH baselines target.

C.2.3 Injection mechanisms

Each mechanism instantiates the map g_c of Eq. (9) for a module $M = \{i_1, \ldots, i_K\}$ active on $[t_1, t_2]$ with strength $s \in [0,1]$ (default $s = 0.5$), where $\bar{X}(t) = K^{-1} \sum_k X_{i_k}(t)$ is the module mean.

Blend. $X_{i_k}(t) \leftarrow (1-s)X_{i_k}(t) + s\bar{X}(t)$, which aligns phase and frequency across module members.

Elevate. $X_{i_k}(t) \leftarrow (1+s)X_{i_k}(t)$, a synchronous amplitude increase.

Depress. $X_{i_k}(t) \leftarrow (1-s)X_{i_k}(t)$, a synchronous amplitude decrease.

OpposeAmp. The first $\lceil K/2 \rceil$ members are scaled by $1+s$ and the rest by $1-s$, giving anti-correlated amplitude modulation at the same coupling magnitude as Elevate and Depress.

PhaseLock. With a_{i_k} and ϕ_{i_k} the Hilbert amplitude and phase of X_{i_k}, and $\phi_{\bar{X}}$ the phase of the module mean,

$$X_{i_k}(t) \leftarrow a_{i_k}(t)\cos\bigl((1-s)\phi_{i_k}(t) + s\,\phi_{\bar{X}}(t)\bigr),$$

which locks phase to a shared reference while leaving each member's amplitude envelope untouched.

C.2.4 Overlapping modules

Blend, Elevate, Depress, and OpposeAmp all act on amplitude. When two of them are active at the same time, no amplitude statistic computed within a window can tell which module a node belongs to, so the ground truth merges amplitude-driven modules that overlap in time into one module before scoring. PhaseLock acts only on phase and is never merged with an amplitude module.

C.3 Structural connectivity model

Stochastic block model graphs of $d = 18$ nodes in three modules of six were generated with within-module edge probability 0.70, between-module probability 0.05, and edge weights drawn from $\mathcal{U}[0.6, 1.0]$. Every draw is rescaled to spectral radius 0.9 by $W \leftarrow W \cdot 0.9/\rho_{\text{raw}}$, where ρ_{raw} is the spectral radius of the raw draw, so the recursion of Eq. (10) is stable on every trial and no draws are discarded. The recursion uses lag $q = 2$ and runs for $L = 3000$ samples per node over 100 trials.

Deterministic intrinsic sinusoids used in earlier versions of this experiment were removed. They are autocorrelated and appear in both $x(t)$ and the lagged regressor $x(t-q)$, which makes the lagged regression endogenous and biased, whereas white innovations make the identification problem well posed.

For each method the time-averaged signed adjacency matrix is compared against W on the off-diagonal entries under two scale-invariant measures. Matrix correlation is the Pearson correlation between the flattened predicted and true edge vectors and shows whether high-weight edges are ranked above low-weight ones. Cosine similarity is $\langle \hat{A}, W\rangle / (\|\hat{A}\|\|W\|)$ and measures angular alignment independently of overall scale. Frobenius distance is not reported. DCI recovers the correlation structure rather than the literal entries of W, so its correctly larger weights are penalized in absolute matrix space even when the topology is right, while a method returning uniformly small values appears artificially close to W.

C.4 DCC–GARCH specification

The DCC–GARCH baseline follows the two-step procedure of Engle [19]. A univariate GARCH(1,1) is fitted per node with the `arch` Python package, and the standardized residuals ϵ_t drive

$$Q_t = (1-a-b)\bar{Q} + a\,\epsilon_{t-1}\epsilon_{t-1}^{\top} + b\,Q_{t-1}, \qquad R_t = \operatorname{diag}(Q_t)^{-1/2} Q_t \operatorname{diag}(Q_t)^{-1/2}, \tag{24}$$

with R_t averaged within each window. Joint re-estimation of (a, b) across all $d = 134$ series is computationally infeasible at this panel size, so the canonical scalar defaults $a = 0.05$ and $b = 0.93$ are used. The result tracks Pearson correlation closely on a panel this size, which is the parametric collapse discussed in Section 4.2.

D Supporting Results

D.1 Validation of the null distribution

Two independent synthetic signals, a sine wave and a repeating inverse Gaussian, were generated at length $L = 10^6$ and 10^5 Monte Carlo trials were run. Each trial drew independent start indices a and b and computed $\text{DCI} = \sum_{u=0}^{l-1} Z_i(a+u) Z_j(b+u)$ on the normalized signals. The empirical distribution matches the Gaussian predicted by Theorem B.3, with empirical and theoretical standard deviations agreeing to three significant figures (Fig. 9).

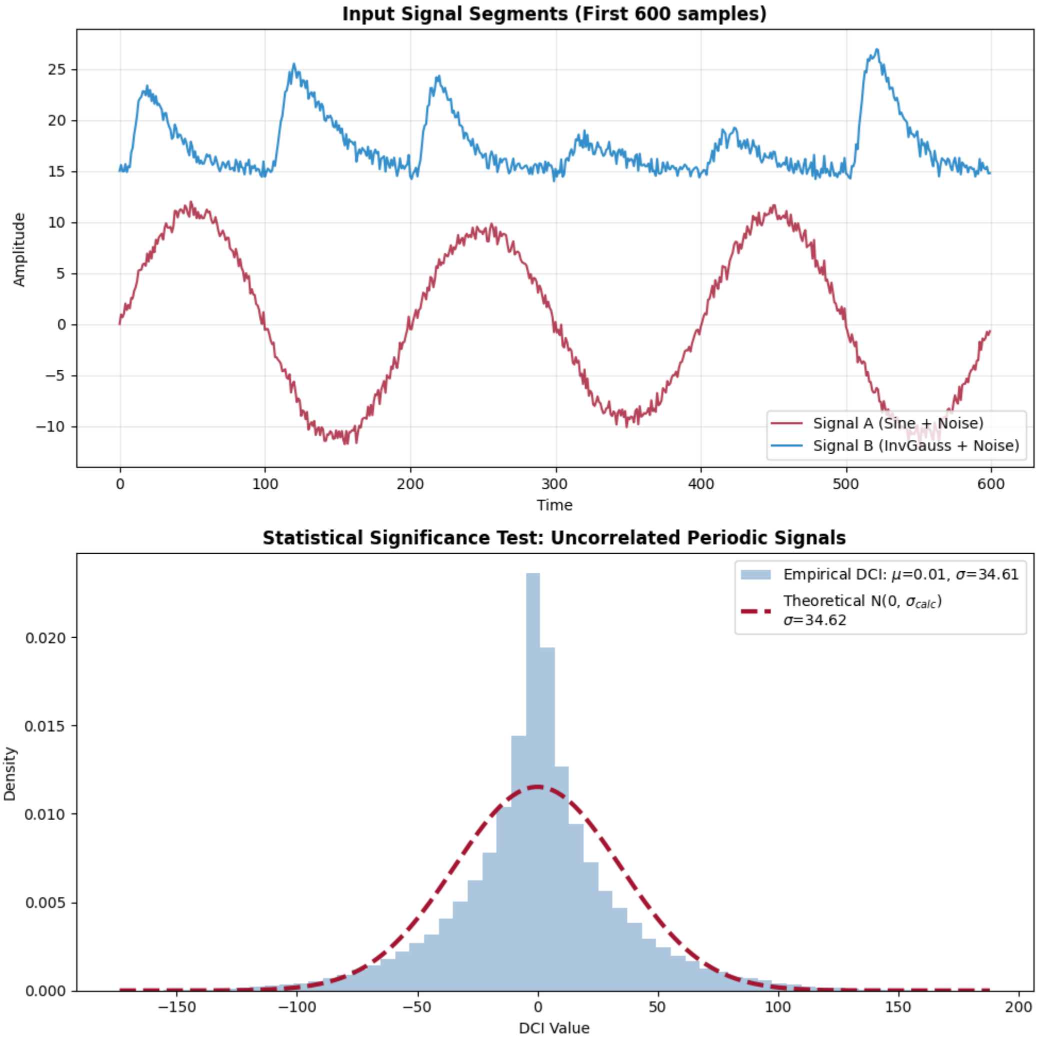

Figure 9. Validation of the DCI null distribution. The empirical distribution of 10^5 DCI values matches the Gaussian predicted by Theorem B.3.

D.2 Parameter sensitivity

Each of the three free parameters was swept independently on the benchmark of Appendix C.1. The edge window $l \in \{1, 20, 50, 100, 150, 200, 300, 500\}$ shows the usual bias–variance behavior, with sharpest recovery when l is near the event duration. The normalization period τ_i, swept around the true periods 150 and 240, tolerates moderate misspecification, and at $\tau_i = 1$ the metric reduces to an ordinary z-score by Proposition B.1. The lag offset $t_{\text{off}} \in \{0, 10, \ldots, 500\}$ peaks at $t_{\text{off}} = 0$ on synchronous events and falls off as $|t_{\text{off}}|$ grows. Figure 10 shows all three sweeps.

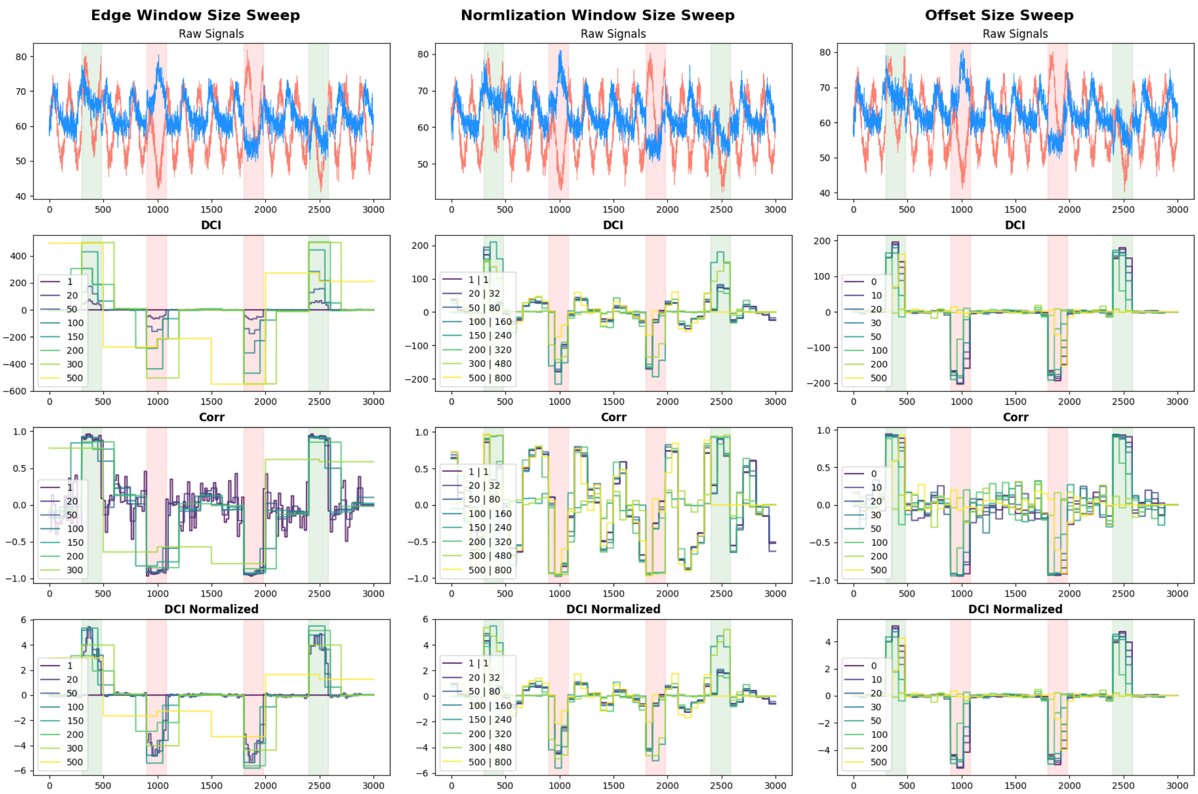

Figure 10. Parameter sensitivity. DCI tolerates moderate misspecification in the edge window, the normalization period, and the lag offset.

D.3 Structural connectivity recovery

Table 4 reports the two measures across trials and Fig. 11 shows a best-case and a worst-case trial. The setup is in Appendix C.3.

Table 4. Structural connectivity recovery (SBM, $d = 18$, 100 trials). Mean ±95% CI. DCI and PCC achieve comparable scale-invariant scores on a stable white-noise-driven LTI model, with DCI marginally below PCC. Neither recovers the literal generating operator (see text).

Method	Matrix corr.	Cosine sim.
DCI	0.917 ± 0.003	0.916 ± 0.004
PCC	0.929 ± 0.004	0.929 ± 0.004

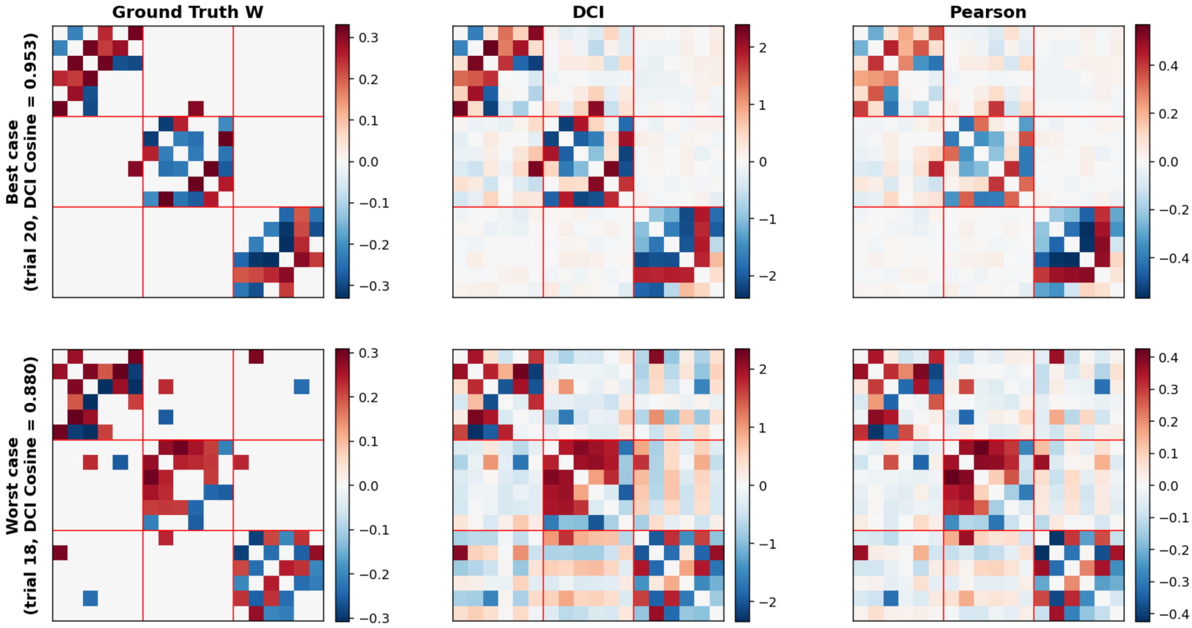

Figure 11. Structural connectivity recovery: best-case and worst-case trials. Ground truth $|W|$ shows three blocks. DCI and PCC both place positive weight on the within-block edges in the right positions, and both also place weight between blocks where indirect paths bleed through the stable dynamics of Eq. (10).

D.4 EEG results in detail

D.4.1 Artifact removal

Eleven of the thirteen retained subjects had exactly one independent component removed. In ten of these it was the first component, the highest-variance component after whitening, which captured ocular activity; in subject 12 it was the second. Subjects 1 and 19 needed no removal.

D.4.2 Within-subject percent change

DCI reports its largest increase in the occipital region and its smallest in the rest of the brain: +112.5% occipital, +87.7% DMN proxy, +77.4% rest of brain. Paired tests give $\Delta_{\mathrm{OCC}} > \Delta_{\mathrm{REST}}$ at $p = 6.0 \times 10^{-3}$, and the $\Delta_{\mathrm{DMN}} > \Delta_{\mathrm{REST}}$ contrast runs in the same direction without reaching significance ($p = 6.7 \times 10^{-2}$). Ten of thirteen subjects satisfy both orderings. Under PCC the changes run the other way (−3.9%, −7.7%, +0.8%) and 2/13 subjects satisfy the ordering. AEC-c gives positive shifts in all three regions (+24.0%, +24.8%, +21.0%), but neither contrast is significant ($p = 0.34$ and $p = 0.20$) and 4/13 subjects satisfy both orderings.

D.4.3 Permutation test on region clustering

With the eyes closed, DMN-proxy clustering sits in the top 2.5% of the 715 possible four-channel subsets of non-visual cortex, so the DMN proxy is more clustered than most other non-occipital channel groups. The occipital region sits in the top 11.1% of its three-channel null, and the rest-of-brain group sits at $p = 0.994$, so the remaining nine channels are correspondingly depleted. With the eyes open the pattern reverses in an interpretable way. Occipital clustering drops to $p = 0.916$ during active visual processing, the expected signature of stimulus-locked alpha desynchronization [39], and the DMN proxy returns to the middle of its null ($p = 0.400$), consistent with anticorrelation between the default mode network and goal-directed attention [40]. Under PCC none of these shifts reaches significance in the expected direction. Table 5 gives the full test.

Table 5. Permutation test on ROI clustering coefficient.

State	Module (k)	Null size	DCI CC	DCI p	PCC p	AEC CC	AEC p
EC	Occipital (3)	560	0.777	0.1107^{ns}	0.2250^{ns}	0.635	0.4893^{ns}
EC	DMN-proxy (4)	715	0.778	$\mathbf{0.0252}^{*}$	0.1399^{ns}	0.647	0.0615^{ns}
EC	Rest of Brain (9)	11440	0.761	0.9943^{ns}	0.9285^{ns}	0.628	0.9326^{ns}
EO	Occipital (3)	560	0.523	0.9161^{ns}	0.2232^{ns}	0.636	0.7214^{ns}
EO	DMN-proxy (4)	715	0.559	0.4000^{ns}	0.0308^{*}	0.643	0.3944^{ns}
EO	Rest of Brain (9)	11440	0.552	0.3238^{ns}	0.9916^{ns}	0.641	0.4000^{ns}

D.4.4 Edge-level and broadband analyses

Network-based statistics [41] find that all 120 edges of the 16-channel graph differ significantly between eyes-closed and eyes-open under DCI ($p < 0.001$). No significant edge component appears under PCC, and 58 edges differ under AEC-c. The alpha response is a scalp-wide phenomenon, and the amplitude sensitivity of DCI captures its full reach. A broadband reanalysis over 0.5–45 Hz without the alpha filter still resolves occipital coupling under DCI ($p = 0.036$), still fails under PCC ($p = 0.140$), and reverses sign under AEC-c ($p = 3.7 \times 10^{-8}$), because broadband envelope energy is dominated by frequencies unrelated to the Berger effect.

D.4.5 Per-subject node strength

Table 6 provides the full per-subject breakdown. DCI shows a within-subject eyes-closed increase in every subject for the occipital ROI and in 12/13 subjects for the DMN-proxy ROI, with Sub_19 reversing, consistent with its low-reactivity status. PCC shows no consistent asymmetry.

Table 6. Per-subject node strength, eyes-closed vs. eyes-open ($n = 13$ subjects).

	DCI						**PCC**						**AEC**					
	OCC		DMN-proxy		REST		OCC		DMN-proxy		REST		OCC		DMN-proxy		REST	
Subj.	EC	EO	EC	EO	EC	EO	EC	EO	EC	EO	EC	EO	EC	EO	EC	EO	EC	EO
Sub_1	9.37	8.02	9.74	8.76	9.03	8.56	6.39	6.61	6.17	6.67	5.71	5.79	1.99	1.50	1.88	1.61	1.89	1.67
Sub_2	12.73	2.91	11.84	3.17	12.63	3.55	8.47	7.47	7.23	7.86	8.39	6.74	2.88	1.76	3.09	1.79	2.68	1.80
Sub_3	10.88	4.53	10.90	7.05	10.72	7.06	8.21	7.54	7.70	7.77	7.14	6.46	2.67	2.12	2.51	1.97	2.40	1.86
Sub_5	12.24	7.13	12.19	7.14	11.90	7.40	8.58	10.69	8.34	10.28	7.96	9.33	2.34	2.67	2.40	2.61	2.42	2.59
Sub_6	11.15	10.05	12.48	11.01	11.21	10.14	7.23	8.09	8.68	8.96	7.60	7.61	1.83	1.67	1.93	1.73	1.85	1.72
Sub_8	11.99	4.34	11.89	5.09	11.80	5.53	7.26	6.54	6.98	7.52	6.99	6.24	2.62	1.69	2.83	1.71	2.56	1.84
Sub_12	13.08	4.55	12.93	4.77	12.32	5.98	9.61	8.49	9.32	8.70	8.17	7.63	1.69	2.13	1.68	2.15	1.68	2.13
Sub_14	11.36	6.91	11.65	6.52	11.35	7.00	7.91	7.89	8.17	7.36	7.58	6.57	2.34	1.81	2.68	1.82	2.45	1.71
Sub_15	10.41	7.03	10.98	7.50	10.98	6.92	6.45	7.14	6.84	7.61	6.97	6.48	2.07	1.85	2.19	2.01	2.14	1.87
Sub_16	12.13	8.13	11.91	9.60	11.81	8.60	7.36	8.42	7.25	8.48	7.06	7.05	2.32	1.53	2.19	1.76	2.22	1.84
Sub_17	12.26	4.31	12.43	4.56	12.35	4.79	7.52	8.99	7.89	8.85	7.90	7.88	2.98	2.09	2.94	2.11	2.79	1.99
Sub_18	12.21	4.12	11.73	5.71	11.76	6.31	8.03	7.78	7.08	7.59	7.16	6.83	2.55	1.87	2.49	1.80	2.30	1.85
Sub_19	7.97	7.85	6.97	7.48	5.06	6.02	9.58	10.28	8.82	9.72	6.98	8.15	1.73	1.85	1.95	1.83	1.70	1.95
Mean	**11.41**	**6.18**	**11.42**	**6.84**	**11.10**	**6.77**	**7.91**	**8.31**	**7.80**	**8.48**	**7.45**	**7.46**	**2.30**	**1.89**	**2.38**	**1.93**	**2.26**	**1.89**

D.5 Macroeconomic regime separation

Table 7. Welch's t-test on graph features, recession vs. expansion windows.

Feature	DCI t	DCI p	PCC t	PCC p	FR t	FR p	Spear. t	Spear. p	DCC t	DCC p
Density	$+16.64$	5.9×10^{-32}	$+5.64$	1.3×10^{-7}	-21.46	3.3×10^{-45}	$+6.06$	2.2×10^{-8}	$+5.64$	1.3×10^{-7}
Modularity	-13.34	5.2×10^{-25}	$+5.16$	1.0×10^{-6}	$+14.22$	2.9×10^{-27}	$+4.70$	7.1×10^{-6}	$+5.05$	1.6×10^{-6}
Transitivity	$+15.84$	4.8×10^{-30}	$+6.00$	2.6×10^{-8}	-3.83	2.0×10^{-4}	$+6.56$	2.2×10^{-9}	$+1.27$	2.1×10^{-1}
Entropy	$+14.41$	6.4×10^{-29}	-6.93	2.5×10^{-10}	-7.10	1.0×10^{-10}	-6.63	1.2×10^{-9}	-6.93	2.5×10^{-10}
Efficiency	$+17.06$	3.1×10^{-33}	$+5.54$	2.2×10^{-7}	-21.05	1.2×10^{-43}	$+5.76$	8.8×10^{-8}	$+5.54$	2.2×10^{-7}

Supplement

S1 Additional Synthetic Sweeps

These runs repeat the dynamic-community pipeline of Section 3.1 while varying one factor at a time.

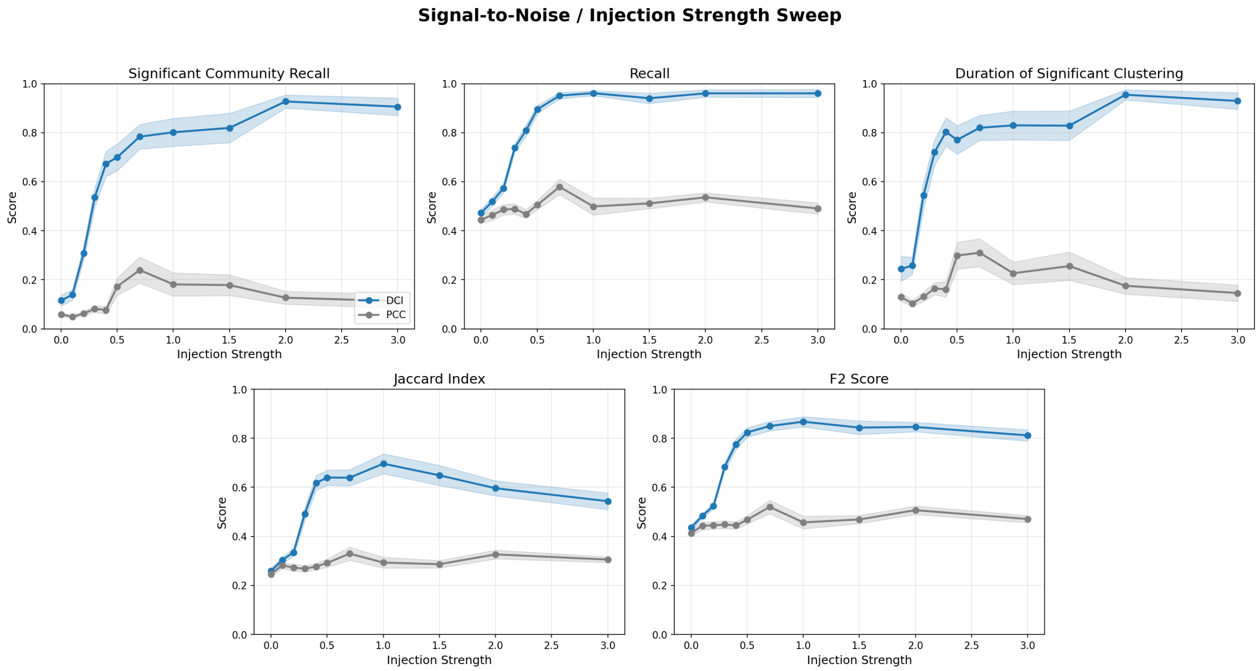

Figure S1. Module detection against injection strength. DCI rises sharply with strength on every metric and saturates near ceiling for $s \geq 0.5$ on the recall-weighted criteria. Pearson stays near baseline across the whole range $s \in [0, 3]$, with the null calibrated at $s = 0$.

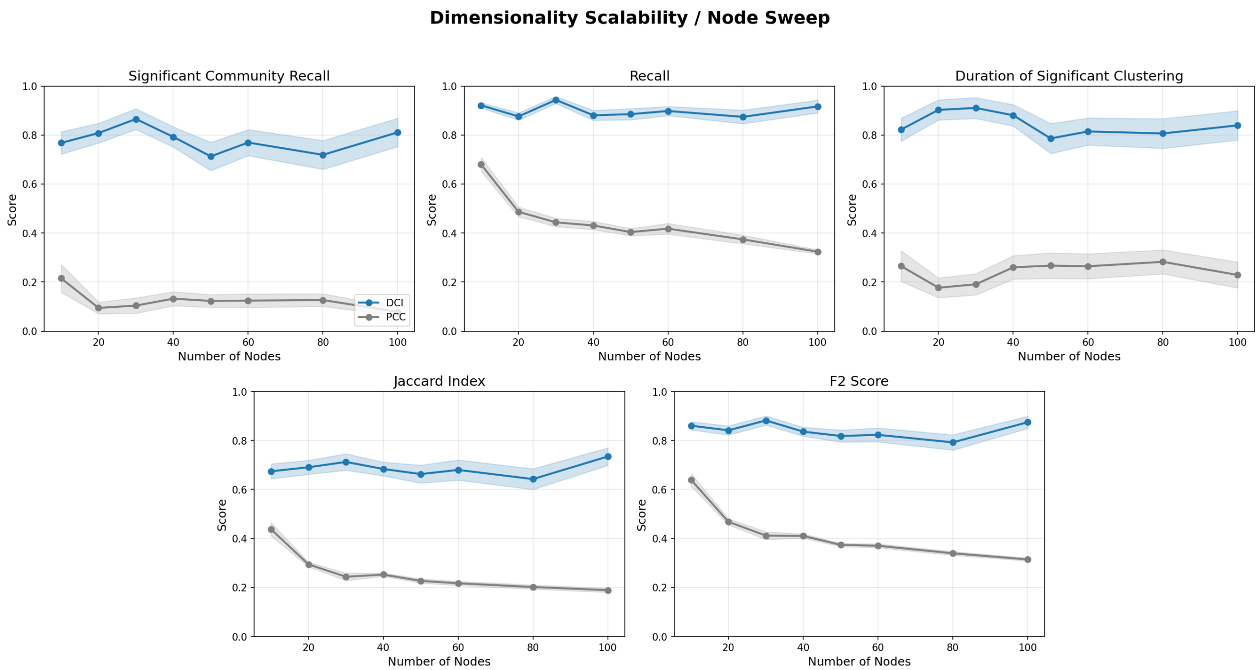

Figure S2. Module detection against node count. The advantage of DCI over Pearson holds across the full range $d \in \{10, \ldots, 100\}$ on every metric.

S2 Worst-Case Examples

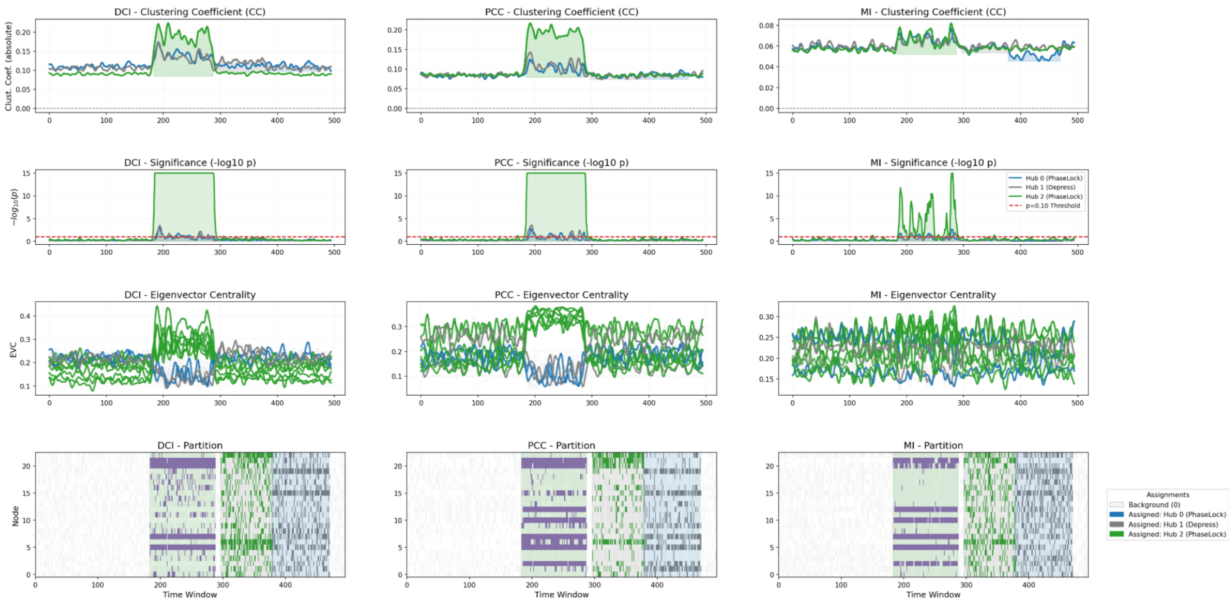

Figure S3. Worst-case module detection. SCR = 0.224. Three modules, two PhaseLock and one Depress, with one dominant PhaseLock module raising the clustering-coefficient baseline. The DCI partition heatmap still places the smaller modules' nodes correctly, so the loss appears in event localization while community recovery stays intact.

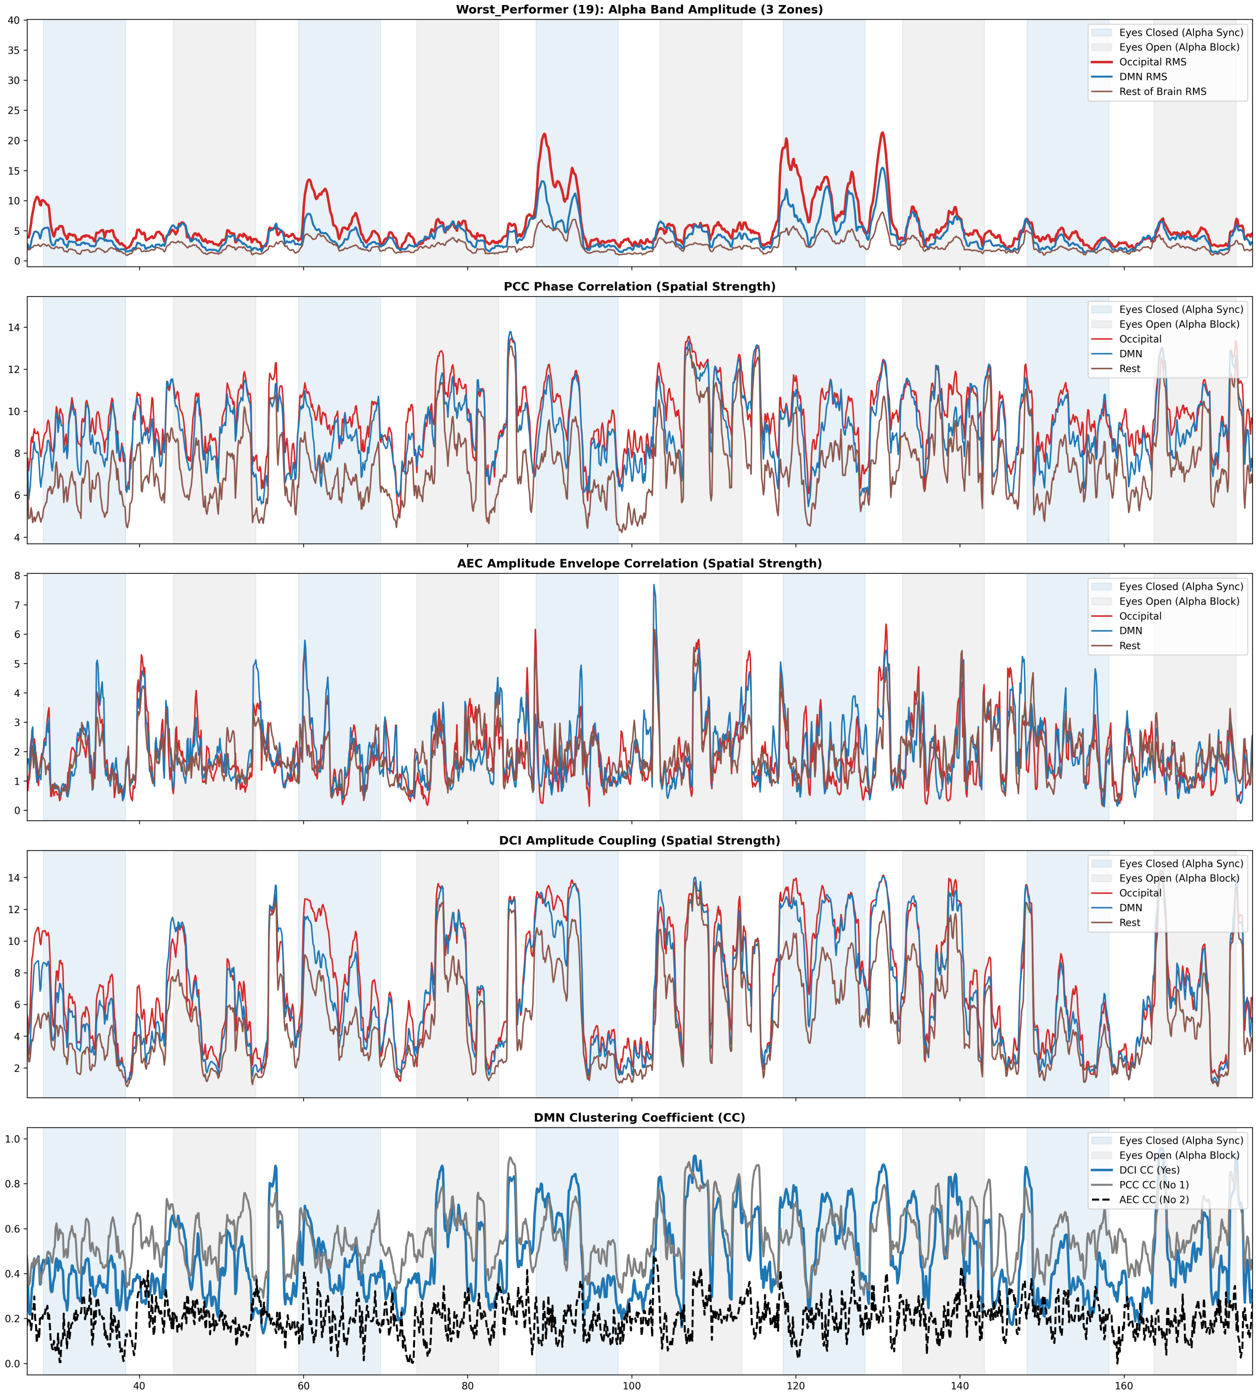

Figure S4. Worst-case alpha-band connectivity under the Berger effect. DCI and PCC both fail to capture state transitions in this low-reactivity subject.

S3 Phase–Amplitude and Phase–Frequency Coupling

To show that DCI generalizes beyond amplitude modulation, two synthetic experiments coupled a slow carrier to a fast companion through either its amplitude envelope or its instantaneous frequency. In both cases the Hilbert-extracted envelope or instantaneous frequency was supplied as the second input to DCI, following the definitions in Appendix C.2.2.

In the first experiment a slow 0.05 Hz sine modulated the amplitude of a fast 2 Hz sine during a central window. All three DCI variants gave a sharp, well-localized response during the coupling window (Fig. S5). In the second the instantaneous frequency of the fast signal was modulated by the slow signal. DCI and its derived forms isolate the coupling window, and the bounded form Corr of Eq. (5) saturates near unity inside the window and vanishes outside (Fig. S6). Given an appropriate upstream decomposition, whether an envelope, an instantaneous frequency, or another derived signal, DCI captures the coupling in the chosen representation.

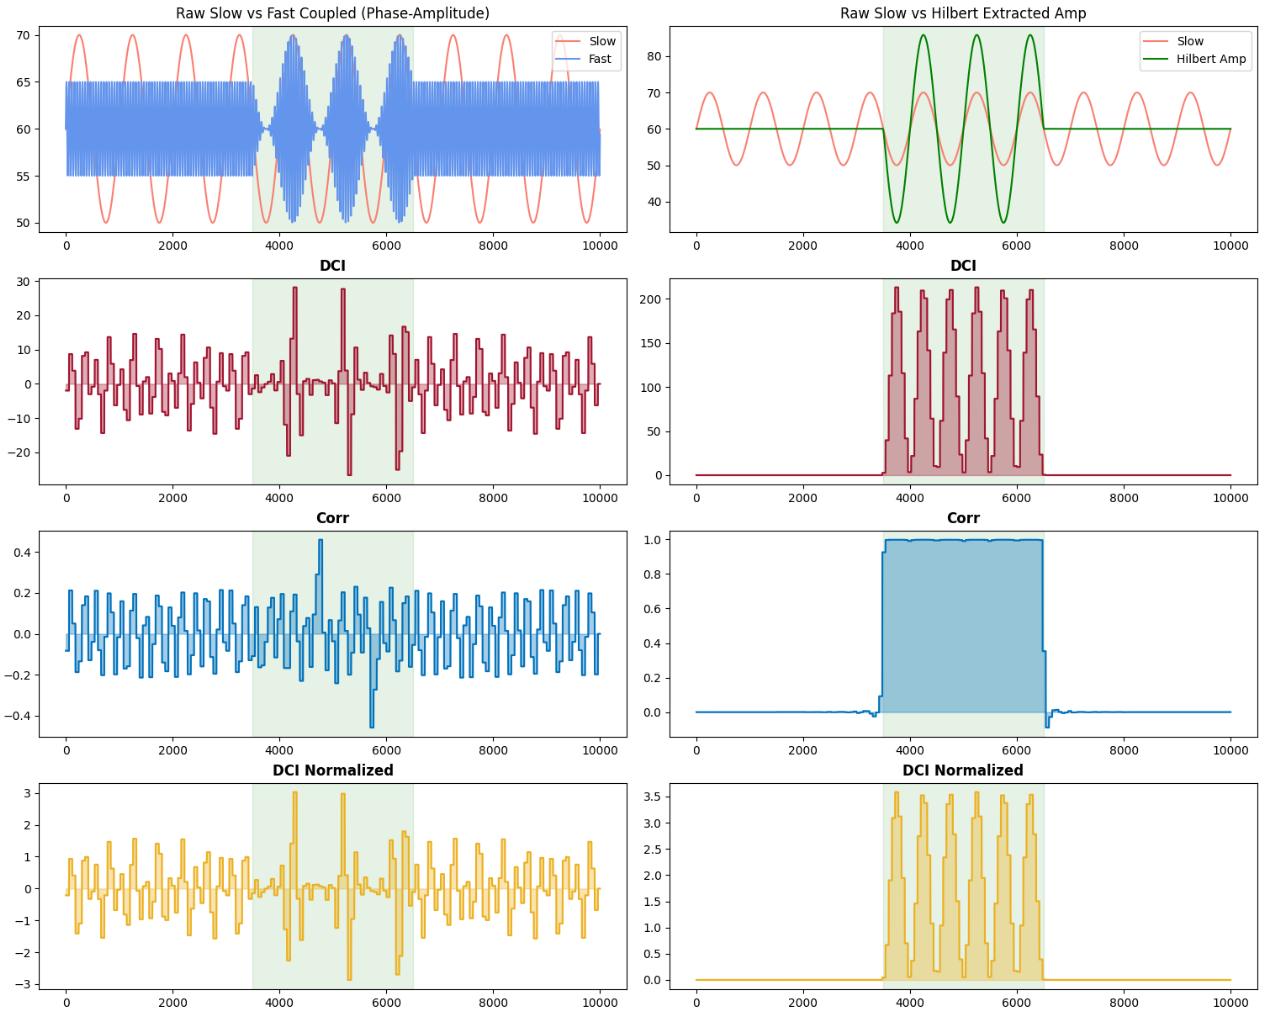

Figure S5. Phase–amplitude coupling detection. All three DCI variants produce a localized response marking the coupling epoch.

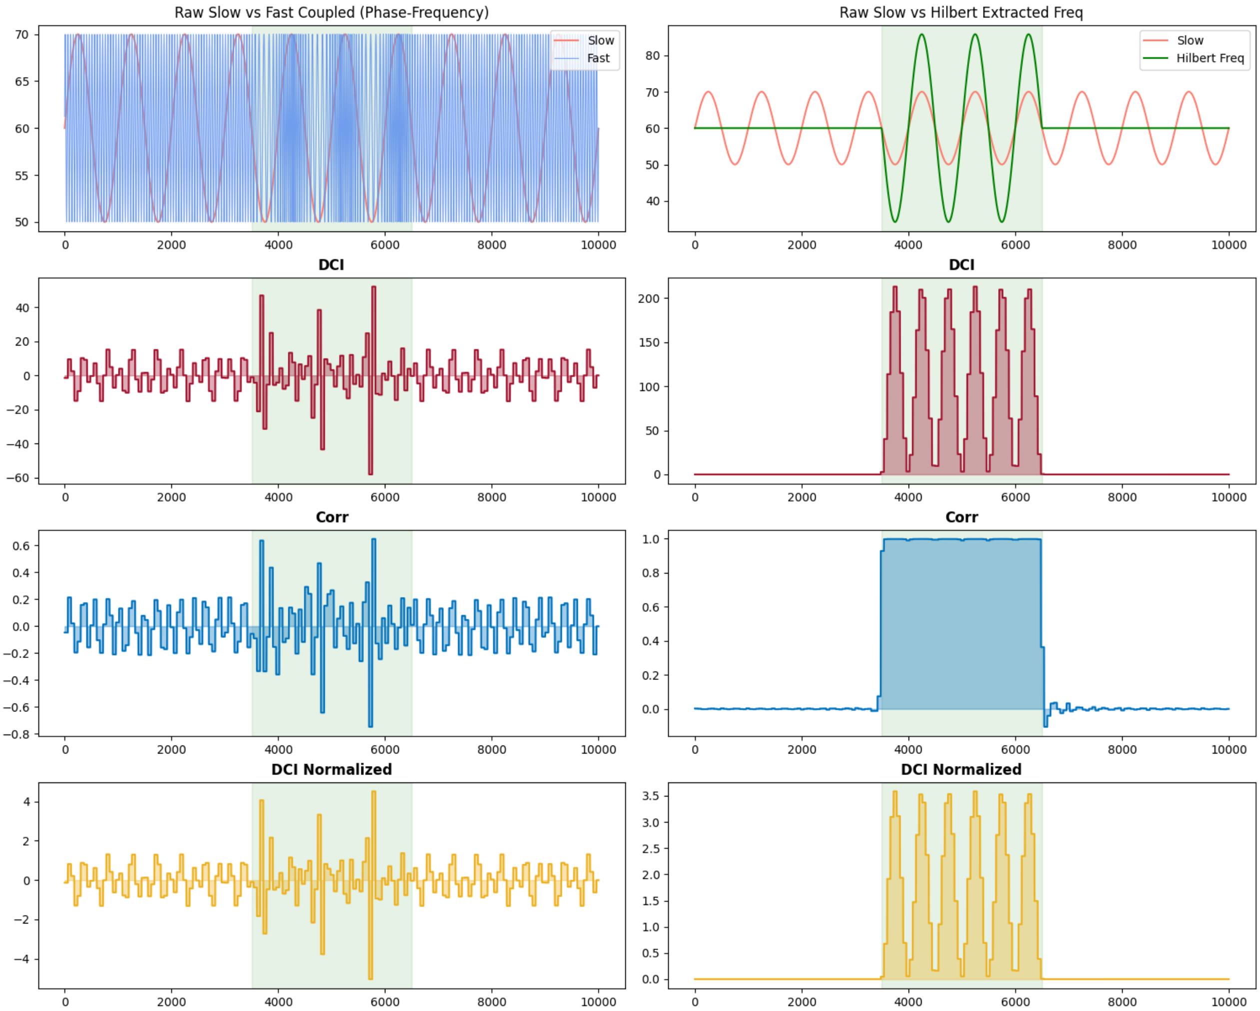

Figure S6. Phase–frequency coupling detection. DCI and its bounded form isolate the coupling window; Corr saturates near unity inside it.