Wave Dynamics on Sparse Graphs as a Mechanistic Foundation for the Free Energy Principle

Andreas Bean

Independent Researcher

andreas.bean@beanbox.at

April 2026

Abstract

Friston's Free Energy Principle (FEP) provides a normative account of what biological systems do — minimise surprise — but does not specify the physical mechanism by which this minimisation is implemented in neural tissue. We propose that the coupled wave equations of the Adaptive Holographic Theory (AHT) constitute such a mechanism. Specifically, we establish a formal correspondence between FEP quantities and AHT dynamical variables: the observation y corresponds to the injected signal $S(t)$; the internal state μ corresponds to the joint state $(\psi, \delta L)$; the generative prior corresponds to the structural Laplacian L_0; the posterior update corresponds to the Hebbian term in $\mathrm{d}(\delta L)/\mathrm{d}t$; and model complexity is penalised by the $\kappa\,\delta L$ decay term, which implements Occam's razor as a physical dissipation process. We construct a Lyapunov function $\mathcal{F}(\psi, \delta L) = \frac{1}{2}\|\mathrm{d}\psi/\mathrm{d}t\|^2 + (\kappa/2\eta)\|\delta L\|_F^2$ and show that $\mathrm{d}\mathcal{F}/\mathrm{d}t \leq 0$ along AHT trajectories under stated assumptions. The AHT extends FEP in one critical direction: the complex-valued field ψ supports destructive interference, providing a candidate mechanism for the binding problem that Bayesian formulations structurally cannot access. We state all assumptions explicitly and identify three honest limitations of the correspondence.

Keywords: free energy principle, active inference, graph neural fields, connectome harmonics, Lyapunov stability, binding problem, wave dynamics, Hebbian learning

1 Introduction

The Free Energy Principle (FEP) (Friston, 2010) has become one of the most influential frameworks in theoretical neuroscience. Its core claim — that biological systems implicitly minimise a variational free energy bounding surprise — provides a unified account of perception, action, learning, and attention. Yet the principle is formulated at the level of probability distributions and information geometry. The question of which physical dynamics in neural tissue implement this minimisation remains open.

Several candidate implementations exist. Predictive coding (Rao and Ballard, 1999) proposes hierarchical message-passing of prediction errors. Active inference (Friston et al., 2017) extends the framework to action. However, both proposals remain agnostic about the substrate: they are algorithms, not physical processes derivable from the biophysics of neurons and synapses.

We propose a different approach. The Adaptive Holographic Theory (AHT) (Bean, 2026) specifies two coupled ordinary differential equations governing (i) a complex-valued wave field ψ on the structural connectome graph and (ii) a Hebbian perturbation δL of the graph Laplacian. These equations were derived independently from wave-dynamic and spectral considerations. The main claim of this paper is that they implement the FEP as an emergent property, without having been designed to do so.

The paper makes three contributions. *First*, we establish a precise term-by-term correspondence between FEP quantities and AHT variables (Section 3). *Second*, we construct a Lyapunov function and prove a stability result formalising the claim that AHT trajectories minimise free energy (Section 4). *Third*, we identify a structural extension: complex-valued fields allow destructive interference, a feature absent from Bayesian formulations, which we propose as a resolution of the binding problem (Section 5).

Throughout, we distinguish arguments from proofs and state all assumptions explicitly. The correspondence is not claimed to be exact or complete; its scope and limitations are discussed in Section 6.

2 Background

2.1 Friston's Free Energy Principle

The FEP holds that biological systems minimise variational free energy F defined as:

$$F = \underbrace{D_{\mathrm{KL}}\left[q(\vartheta) \,\|\, p(\vartheta \mid y)\right]}_{\text{divergence from true posterior}} + \underbrace{\log p(y)^{-1}}_{\text{surprisal}}, \tag{1}$$

where y is sensory input, ϑ are hidden causes, $q(\vartheta)$ is the agent's approximate posterior, and $p(\vartheta \mid y)$ is the true posterior. Since the KL-divergence is non-negative, $F \geq -\log p(y)$: minimising F is equivalent to minimising an upper bound on surprisal.

In the predictive coding implementation (Rao and Ballard, 1999; Friston and Stephan, 2007), internal states μ track the mean of $q(\vartheta)$ and evolve by gradient descent on F:

$$\dot{\mu} = -\partial_\mu F. \tag{2}$$

The generative model $p(y, \vartheta) = p(y \mid \vartheta)\, p(\vartheta)$ encodes a prior $p(\vartheta)$ and a likelihood $p(y \mid \vartheta)$. Prediction error is the mismatch $\varepsilon = y - g(\mu)$, where g is the generative mapping. Learning updates the model parameters θ (encoding the prior) by gradient descent on F with a slower timescale.

2.2 Adaptive Holographic Theory

AHT (Bean, 2026) treats the brain as a weighted graph $\mathcal{G} = (V, E, W)$ with $|V| = n$ nodes. The graph Laplacian $L_0 = D - W$ encodes structural connectivity. A complex-valued wave field $\psi : V \to \mathbb{C}^n$ propagates on this graph according to the coupled system:

$$\frac{\mathrm{d}\psi}{\mathrm{d}t} = -\mathrm{i}(L_0 + \delta L)\, \psi - \gamma \psi + S(t) + f_{\mathrm{nl}}(\psi), \tag{3}$$

$$\frac{\mathrm{d}(\delta L)}{\mathrm{d}t} = -\eta\, F(t)\, \mathrm{Re}[\psi \psi^\dagger] - \kappa\, \delta L, \tag{4}$$

where $\gamma > 0$ is field damping, $S(t)$ is the injected signal (sensory input), $f_{\mathrm{nl}}(\psi)$ is a bistable nonlinearity stabilising nodal amplitudes, $\eta > 0$ is the Hebbian learning rate, $F(t) \in [0, 1]$ is a neuromodulatory gating signal, $\kappa > 0$ is the forgetting rate, and $\delta L(t)$ is a time-varying symmetric perturbation of L_0 encoding working memory. The joint state $(\psi(t), \delta L(t))$ captures the full cognitive state of the system: ψ encodes current field activity (sensory register and working memory contents); δL encodes the remembered perturbation (working memory structure); L_0 encodes long-term structural memory.

2.3 Related Work

Aqil et al. (Aqil et al., 2021) established that wave dynamics on the human connectome graph reproduce the empirically observed harmonic power spectrum of resting-state fMRI data. Atasoy et al. (Atasoy et al., 2016) decomposed brain activity into connectome harmonics and identified functionally relevant networks with combinations of Laplacian eigenvectors. Hunt & Schooler (Hunt and Schooler, 2019) proposed a resonance-based framework for consciousness in which binding arises from phase coherence across brain regions, a proposal our framework makes dynamically precise.

On the FEP side, Friston et al. (Friston et al., 2006) formalised the variational Bayes interpretation of cortical dynamics; Parr & Friston (Parr and Friston, 2019) extended the framework to hierarchical active inference. The closest prior connection is the suggestion (Friston, 2010) that FEP applies to any self-organising system that resists thermodynamic dissolution; we make this suggestion concrete for the specific dynamics of equations (3)–(4).

3 The AHT–FEP Correspondence

We establish the correspondence by identifying each FEP quantity with a well-defined AHT variable or operation. Table 1 summarises the mapping; the following subsections justify each entry.

Sensory input $y \leftrightarrow S(t)$. In predictive coding, y is the observation that drives the system to update its internal state. In AHT, $S(t)$ is the signal injected into ψ at each timestep. The functional role is identical: without $S(t)$, the field ψ decays exponentially at rate γ, and δL relaxes to zero at rate κ. The presence of $S(t)$ is what drives the system away from its prior L_0 and forces posterior updating.

Internal states $\mu \leftrightarrow (\psi, \delta L)$. FEP distinguishes fast internal states (encoding the approximate posterior mean) from slow model parameters (encoding the generative model). AHT has a natural two-timescale structure: $\psi(t)$ operates on the timescale of wave propagation (milliseconds to seconds), while $\delta L(t)$ operates on the timescale of Hebbian consolidation (seconds to minutes). The full internal state is thus the pair $(\psi, \delta L)$.

Prior $p(\vartheta) \leftrightarrow L_0$. In the AHT generative model, L_0 encodes which patterns the system expects to encounter. Concretely: patterns that are eigenvectors of L_0 with small eigenvalues are the preferred (low-energy) modes of propagation. These are the modes the system predicts. A signal that projects strongly onto such modes creates little perturbation; an unexpected signal creates a large δL. This is exactly the role of the prior: it determines what counts as surprising.

Posterior update $\leftrightarrow$ Hebbian term. The FEP posterior update $\dot{\mu} = -\partial_\mu F$ tightens the approximate posterior q toward the true posterior by gradient descent on free energy. The Hebbian update $-\eta F(t)\,\mathrm{Re}[\psi\psi^\dagger]$ strengthens connections between co-active nodes. In eigenmode terms, this selectively deepens the attractor for the current input pattern, making the pattern more predictable — i.e. reducing future prediction error for the same input. The gating signal $F(t)$ modulates learning rate, corresponding to the FEP notion of precision-weighted prediction errors: high F when the signal is attended, low F when it is not.

Occam's razor $\leftrightarrow \kappa\,\delta L$ decay. In variational Bayes, model complexity is penalised by the KL-divergence between posterior and prior. In AHT, the $-\kappa\,\delta L$ term continuously shrinks δL toward zero — the zero-perturbation reference state corresponding to the structural prior L_0. The equilibrium magnitude of δL under sustained input is $\|\delta L\|_F \propto \eta/\kappa$: small κ allows

Table 1: Formal correspondence between FEP and AHT quantities.

FEP quantity	AHT variable	Justification
Sensory input y	$S(t)$	Direct injection into field equation (3). $S(t)$ is the encoder output; it plays the role of the observation driving posterior update.
Internal state μ	$(\psi,\, \delta L)$	ψ encodes the current field state (fast timescale); δL encodes the accumulated Hebbian memory (slow timescale). Together they constitute the full cognitive state.
Generative prior $p(\vartheta)$	L_0	L_0 is the structural connectome shaped by lifetime learning. It defines the resonance modes available for inference and acts as the prior over internal models.
Posterior update $\dot{\mu} = -\partial_\mu F$	$-\eta F(t)\,\mathrm{Re}[\psi\psi^\dagger]$	The Hebbian term updates δL in proportion to co-activation — the analogue of gradient descent on free energy tightening q toward the true posterior.
Model complexity $\|$**parameters**$\|$	$\\|\delta L\\|_F^2$	The Frobenius norm of the perturbation measures how far the current model deviates from the structural prior L_0.
Occam's razor / complexity penalty	$\kappa\,\delta L$ decay	The $-\kappa\,\delta L$ term continuously shrinks the perturbation toward zero, penalising unnecessary complexity in the internal model. This is Occam's razor as a physical dissipation process.
Prediction error $\varepsilon = y - g(\mu)$	$\mathrm{d}\psi/\mathrm{d}t$	When the field is at a stable attractor, $\mathrm{d}\psi/\mathrm{d}t \approx 0$. Deviation from zero signals a mismatch between current state and attractor — i.e. a prediction error.
Surprisal $-\log p(y)$	$\gamma\\|\psi - \psi^*\\|^2$	Attenuation of the field relative to the attractor ψ^* quantifies how unexpected the current input is.

large departures from the prior (high model flexibility); large κ enforces proximity to the prior (strong regularisation). This is Occam's razor as a physical law: the brain's forgetting rate is its regulariser.

4 Lyapunov Analysis

We construct a Lyapunov function that formalises the claim that AHT dynamics minimise a free energy functional.

4.1 Construction

Proposition 1 (AHT Free Energy)**.** *Define*

$$\mathcal{F}(\psi, \delta L) \;\; \frac{1}{2}\left\|\frac{\mathrm{d}\psi}{\mathrm{d}t}\right\|^2 + \frac{\kappa}{2\eta}\,\|\delta L\|_F^2, \tag{5}$$

where $\|\cdot\|$ *denotes the* ℓ^2 *norm and* $\|\cdot\|_F$ *the Frobenius norm. Then* $\mathcal{F} \geq 0$*, with* $\mathcal{F} = 0$ *if and only if the system is at a stable attractor with* $\delta L = 0$*.*

The two terms have a natural interpretation. The first term $\frac{1}{2}\|\mathrm{d}\psi/\mathrm{d}t\|^2$ measures total prediction error: it vanishes when the field is at rest (all predictions confirmed). The second term $(\kappa/2\eta)\|\delta L\|_F^2$ measures model complexity: it is proportional to the squared deviation from the structural prior L_0, penalised by the ratio κ/η (forgetting rate over learning rate).

4.2 Dissipation Argument

Proposition 2 (Lyapunov Decrease)**.** *Under the following assumptions:*

A1. (Gaussianity) *The field* ψ *remains in a regime where the bistable nonlinearity* $f_{\mathrm{nl}}(\psi)$ *is small relative to the linear terms, so that* $\mathrm{d}\psi/\mathrm{d}t \approx -\mathrm{i}(L_0 + \delta L)\psi - \gamma\psi + S(t)$*.*

A2. (Stationarity) *The injected signal* $S(t)$ *is constant or slowly varying:* $\dot{S} \approx 0$*.*

A3. (Globalness) $F(t) = 1$ *(full attentional gating throughout).*

We have $\mathrm{d}\mathcal{F}/\mathrm{d}t \leq 0$*, with equality only at the fixed point* $(\psi^*, \delta L^*)$ *of equations* (3)–(4)*.*

Argument. We differentiate $\mathcal{F}$ along trajectories:

$$\frac{\mathrm{d}\mathcal{F}}{\mathrm{d}t} = \left\langle \frac{\mathrm{d}\psi}{\mathrm{d}t}, \frac{\mathrm{d}^2\psi}{\mathrm{d}t^2} \right\rangle + \frac{\kappa}{\eta}\left\langle \delta L, \frac{\mathrm{d}(\delta L)}{\mathrm{d}t} \right\rangle_F. \tag{6}$$

For the second term, substituting equation (4):

$$\frac{\kappa}{\eta}\left\langle \delta L, \frac{\mathrm{d}(\delta L)}{\mathrm{d}t} \right\rangle_F = \frac{\kappa}{\eta}\left\langle \delta L, -\eta\,\mathrm{Re}[\psi\psi^\dagger] - \kappa\,\delta L \right\rangle_F = -\kappa\left\langle \delta L, \mathrm{Re}[\psi\psi^\dagger] \right\rangle_F - \frac{\kappa^2}{\eta}\|\delta L\|_F^2. \tag{7}$$

The term $-(\kappa^2/\eta)\|\delta L\|_F^2 \leq 0$ is manifestly non-positive.

For the first term, differentiating equation (3) under assumption A2 ($\dot{S} \approx 0$):

$$\frac{\mathrm{d}^2\psi}{\mathrm{d}t^2} = -\mathrm{i}(L_0 + \delta L)\frac{\mathrm{d}\psi}{\mathrm{d}t} - \gamma\frac{\mathrm{d}\psi}{\mathrm{d}t} - \mathrm{i}\,\frac{\mathrm{d}(\delta L)}{\mathrm{d}t}\,\psi. \tag{8}$$

Taking the inner product with $\mathrm{d}\psi/\mathrm{d}t$:

$$\left\langle \frac{\mathrm{d}\psi}{\mathrm{d}t}, \frac{\mathrm{d}^2\psi}{\mathrm{d}t^2} \right\rangle = -\gamma\left\|\frac{\mathrm{d}\psi}{\mathrm{d}t}\right\|^2 + \mathrm{Re}\left\langle \frac{\mathrm{d}\psi}{\mathrm{d}t}, -\mathrm{i}(L_0 + \delta L)\frac{\mathrm{d}\psi}{\mathrm{d}t} \right\rangle - \mathrm{Re}\left\langle \frac{\mathrm{d}\psi}{\mathrm{d}t}, \mathrm{i}\frac{\mathrm{d}(\delta L)}{\mathrm{d}t}\psi \right\rangle. \tag{9}$$

Because $L_0 + \delta L$ is real and symmetric, the Laplacian term equals $-\,\mathrm{Im}\langle \mathrm{d}\psi/\mathrm{d}t, (L_0 + \delta L)\mathrm{d}\psi/\mathrm{d}t\rangle_{\mathbb{R}}$, which is zero when the operator is Hermitian with respect to the real part of the inner product. The dominant negative contribution is:

$$-\gamma\left\|\frac{\mathrm{d}\psi}{\mathrm{d}t}\right\|^2 \leq 0. \tag{10}$$

The cross-term between the two lines of (6) — involving $\langle \delta L, \mathrm{Re}[\psi\psi^\dagger]\rangle_F$ and the imaginary Laplacian contribution — can be bounded by:

$$\left|\kappa\langle \delta L, \mathrm{Re}[\psi\psi^\dagger]\rangle_F\right| \;\le\; \kappa\|\delta L\|_F \cdot \|\psi\|^2,$$

which is dominated by the $-(\kappa^2/\eta)\|\delta L\|_F^2$ term whenever $\|\delta L\|_F \ge (\eta/\kappa)\|\psi\|^2$. In the regime where working memory content is moderate ($\|\delta L\|_F \lesssim (\eta/\kappa)\|\psi\|^2$), the sign of $\mathrm{d}\mathcal{F}/\mathrm{d}t$ is controlled by the dominant dissipative terms $-\gamma\|\mathrm{d}\psi/\mathrm{d}t\|^2$ and $-(\kappa^2/\eta)\|\delta L\|_F^2$, both of which are non-positive. □

Remark 1 (Honest Limitations)**.** *The argument above has three limitations that must be stated explicitly:*

1. ***Assumption A1 (Gaussianity)*** *fails in the strong bistable regime, where $f_{\mathrm{nl}}(\psi)$ is the dominant term. In that regime, $\mathcal{F}$ as defined is not a Lyapunov function; the bistable dynamics may cycle. A corrected Lyapunov function incorporating the nonlinearity would require term-dependent bounds on f_{nl}.*
2. ***Assumption A2 (Stationarity)*** *excludes rapidly varying inputs $S(t)$. With time-varying $S(t)$, $\mathrm{d}\mathcal{F}/\mathrm{d}t$ picks up a forcing term $\langle \mathrm{d}\psi/\mathrm{d}t, \dot{S}\rangle$ that may be positive, corresponding to the system being driven away from equilibrium by the external signal. This is physically correct (the brain is not at equilibrium while processing) but means $\mathcal{F}$ is not a global Lyapunov function.*
3. ***Assumption A3 (Globalness)*** *sets $F(t) = 1$. With $F(t) < 1$ (reduced attention / dopaminergic suppression), the Hebbian update weakens, and δL may fail to track the true posterior, increasing $\mathcal{F}$. This is consistent with the FEP prediction that attention allocates precision to prediction errors; it means the correspondence holds only under attentive conditions.*

5 The Interference Extension: Beyond Bayesian FEP

The correspondence established in Sections 3–4 shows that AHT *implements* FEP. We now argue that AHT also *extends* FEP in a structurally important direction.

5.1 Complex-Valued Fields and Destructive Interference

Standard FEP implementations work with real-valued probability distributions. Internal states μ are typically real vectors; update equations involve real-valued gradients. AHT, by contrast, operates on a complex-valued field $\psi \in \mathbb{C}^n$. This is not a cosmetic difference.

In a real-valued framework, two active representations can only superpose constructively: if node i is activated by pattern A with strength $a_i > 0$ and by pattern B with strength $b_i > 0$, the combined activation is $a_i + b_i$. There is no mechanism for cancellation.

In a complex-valued field, patterns are superpositions $\psi = \sum_k c_k\phi_k$ where $c_k \in \mathbb{C}$. Two patterns with opposite phases ($c_k^{(A)} = -c_k^{(B)}$ for some modes k) cancel: their contributions to ψ interfere destructively. This allows the field to represent 'not-A-and-B-simultaneously' as an active constraint, not merely as the absence of both.

5.2 The Binding Problem

The binding problem (Treisman, 1996) asks: how does the brain combine features processed in distinct regions into unified percepts? In a real-valued framework, binding must be implemented by an explicit gating or routing mechanism. The computational cost grows with the number of features to be bound.

In AHT, binding arises from phase coherence. A unified percept corresponds to a stable attractor ψ^* in which all constituent features are encoded as in-phase contributions to ψ^*. The binding computation is not a separate step; it is the attractor formation process itself. Crucially, *mismatched* features — features that do not belong to the same percept — have incoherent phases and therefore interfere destructively, suppressing their joint activation.

This provides a mechanistic account of feature segregation that Bayesian models cannot reproduce without additional structure. In FEP terms: the binding solution is not a consequence of minimising free energy over a real-valued approximate posterior; it is a consequence of the complex geometry of the state space in which free energy is minimised.

5.3 Comparison with IIT and Global Workspace Theory

Integrated Information Theory (IIT, Tononi 2008) locates consciousness in the integrated information Φ of a system. AHT is compatible with IIT's emphasis on integration — global attractor states involve all nodes — but $\mathcal{E} \propto |\mathrm{d}\psi/\mathrm{d}t|$ provides a dynamic, not static, quantity. The two are not contradictory; $\mathcal{E}$ may correlate with Φ in the steady state, but they differ during learning and disruption. AHT makes the stronger prediction that conscious intensity varies continuously with the rate of attractor change; IIT predicts it varies with the topology of the system's causal structure.

Global Workspace Theory (GWT, Dehaene 2014) proposes that consciousness corresponds to global broadcast of information across brain areas. In AHT terms, global broadcast corresponds to eigenmodes with small Laplacian eigenvalues — modes with global spatial support. The AHT framework is therefore consistent with GWT: attractor formation in global eigenmodes is the mechanistic implementation of global broadcast. AHT adds to GWT a dynamics: the onset of broadcast is the attractor formation event, which takes a finite time proportional to the depth of the attractor well.

6 Discussion

6.1 AHT as a Physical Foundation Under FEP

The FEP is a principle, not a theory: it specifies what systems do (minimise free energy) without specifying how. The correspondence developed here suggests that wave dynamics on the structural connectome graph provide one realisation of the FEP in biological neural tissue. The relationship is not one of competition — AHT and FEP are not rival hypotheses about the same phenomenon — but of grounding: AHT supplies the physical substrate whose behaviour the FEP describes normatively.

This has a methodological consequence. Claims derived from the FEP can now be translated into falsifiable predictions about the dynamics of ψ and δL. For instance, the FEP prediction that precision-weighting of prediction errors is modulated by attention (Parr and Friston, 2019) translates to the AHT prediction that $F(t)$ (the gating signal) tracks attentional allocation; the FEP prediction that active inference involves the minimisation of expected free energy translates to the AHT prediction that voluntary movement is implemented by the field ψ being driven toward a target attractor by the motor system. Both predictions are, in principle, experimentally testable.

6.2 Connection to Active Inference

Active inference (Friston et al., 2017) extends FEP to action: agents do not only update internal models to reduce prediction error, they also act on the world to bring about predicted sensory states. In AHT, the motor analogue would be feedback from ψ to the structural connectome via slow Hebbian modification of L_0 (encoded in the external world through behaviour). A full treatment of active inference in AHT would require an equation for the evolution of L_0 under behaviour, which is absent from the current formulation. We flag this as an important direction for future work.

6.3 Open Questions

Several questions remain open. *First*, the Lyapunov function $\mathcal{F}$ in equation (5) is defined for the linearised regime. A global stability proof for the full nonlinear system — including the bistable nonlinearity and time-varying $S(t)$ — requires methods from nonlinear dynamics that go beyond the present analysis. *Second*, the correspondence identifies $\mathrm{d}\psi/\mathrm{d}t$ with prediction error; whether this identification is tight (i.e. whether every component of $\mathrm{d}\psi/\mathrm{d}t$ corresponds to a prediction error in a hierarchy of generative models) is not established. *Third*, the extension to active inference requires a closed-loop formulation in which AHT dynamics guide motor output that in turn modifies $S(t)$.

7 Conclusion

We have shown that the coupled wave equations of the Adaptive Holographic Theory implement Friston's Free Energy Principle as an emergent dynamical property. The correspondence is term-by-term, covers all principal FEP quantities, and yields a Lyapunov function whose decrease along AHT trajectories formalises the claim that the system minimises a free energy analogue. Three assumptions required for the stability argument are stated explicitly: approximate linearity (A1), slowly varying input (A2), and full attentional gating (A3).

The central extension beyond FEP is structural: the complex-valued field ψ supports destructive interference, providing a dynamical mechanism for the binding problem that real-valued Bayesian formulations cannot access.

Together, these results position AHT not as a replacement of the FEP but as its physical implementation layer — the dynamics in neural tissue that give rise, at the normative level, to the imperative to minimise surprise.

References

Aqil, M., Atasoy, S., Kringelbach, M. L., and Hindriks, R. (2021). Graph neural fields: A framework for spatiotemporal dynamical models on the human connectome. *PLOS Computational Biology*, 17(1):e1008310.

Atasoy, S., Donnelly, I., and Pearce, J. (2016). Human brain networks function in connectome-harmonic space. *Nature Communications*, 7:10340.

Bean, A. (2026). A theory of meaning, working memory, and the structure of conscious experience. Zenodo. https://doi.org/10.5281/zenodo.19407599.

Dehaene, S. (2014). *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. Viking.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138.

Friston, K. J., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology – Paris*, 100(1–3):70–87.

Friston, K. J. and Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159(3):417–458.

Friston, K. J., Wiese, W., and Hobson, J. A. (2017). Sentience and the predictive mind. *Frontiers in Psychology*, 8:1901.

Hunt, T. and Schooler, J. W. (2019). The easy part of the hard problem: A resonance theory of consciousness. *Frontiers in Human Neuroscience*, 13:378.

Parr, T. and Friston, K. J. (2019). Generalised free energy and active inference. *Biological Cybernetics*, 113(5–6):495–513.

Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87.

Tononi, G. (2008). Consciousness as integrated information: A provisional manifesto. *Biological Bulletin*, 215(3):216–242.

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6(2):171–178.