

Beyond the Bat: Dissolving the Explanatory Gap Through Meaningful Integration — III. Mathematical Foundations

by Daniel Simon Jr.

Abstract

This paper formalizes the Recursive Social Prediction (RSP) model of consciousness — a three-level architecture where phenomenal consciousness is constituted by recursive integration of core affect (Level 1), core schemas (Level 2), and recursive social prediction (Level 3). The philosophical motivation and empirical evidence appear in companion papers (Parts I and II); here we provide the formal framework.

Five mathematical traditions converge on the same three-level dynamics: (1) hierarchical predictive coding specifies neural message-passing; (2) variational free energy establishes optimality conditions and the necessity of three levels; (3) intrinsically motivated reinforcement learning connects valence to reward and the strange loop to model-based planning; (4) cascade feedback control derives stability conditions and the contraction mapping convergence theorem; (5) a unified master functional shows all four are projections of one variational principle, analogous to equivalent formulations in classical mechanics.

We prove cross-formalism equivalence under the Laplace approximation, derive emergent properties (coordinate covariance, precision as universal currency), and provide algorithmic specifications with phylogenetic parameterization. A formal theory of intelligence derived from the same master functional establishes a constitutive identity between recursive social intelligence and consciousness at Level 3b. The bootstrapping problem — how the hierarchy is learned — is addressed through bilevel optimization (Dupoux et al., 2026): an evolutionary outer loop tunes meta-parameters while a developmental inner loop minimizes free energy with critical period gates regulating plasticity.

1. Introduction: Why Formalize Consciousness?

Consciousness theories are easy to state and hard to test. The difficulty is not empirical (there is abundant data) but formal: most theories lack the mathematical precision needed to derive specific, quantitative predictions from their architectural claims. Global Workspace Theory

(Baars, 1988; Dehaene & Changeux, 2011) describes broadcast but does not specify the dynamics. Integrated Information Theory (Tononi, 2008; Tononi & Koch, 2015) provides a measure (Φ) but faces computational intractability. Higher-Order Theories (Rosenthal, 2005; Lau & Rosenthal, 2011) describe the structure of meta-representation but not its dynamics or stability conditions.

This paper provides a complete mathematical formalization of the Recursive Social Prediction (RSP) model, a three-level architecture for consciousness whose philosophical motivation is presented in Part I (Author, 2026a) and whose empirical evidence is in Part II (Author, 2026b). The architecture proposes that phenomenal consciousness is constituted by the integration of core affect (Level 1), core schemas (Level 2), and recursive social prediction (Level 3), with the "strange loop" of Level 3b (recursive self-modeling) being the architectural locus of phenomenal experience.

Why Five Formalisms?

A single formalism would suffice if the goal were only simulation. But the RSP architecture makes claims that span perception (predictive coding), learning (free energy), motivation (reinforcement learning), and stability (control theory). Each mathematical tradition reveals different structural features:

Predictive coding (Appendix A) specifies the neural message-passing architecture: how prediction errors flow bottom-up and predictions flow top-down across the three-level hierarchy.

Free energy (Appendix B) establishes the variational optimality conditions for adding successive levels under a stipulated precision-gating relation, yielding the developmental ordering constraint.

Reinforcement learning (Appendix C) connects the architecture to reward, value, and exploration: valence is identified as the reward signal, schemas as the learned state representation, and the strange loop implements model-based planning.

Control theory (Appendix D) provides stability analysis: the cascade control structure, the contraction mapping theorem for the strange loop, and the formal conditions under which the recursive self-model converges.

Unified framework (Appendix E) shows the prior four formalisms admit a unified treatment as projections of a single master variational functional, derives emergent properties, and provides the algorithmic specification.

The fact that five independently developed mathematical traditions converge on the same three-level dynamics is itself evidence for the architecture, analogous to how the convergence of

Newtonian, Lagrangian, Hamiltonian, and Hamilton-Jacobi formulations in classical mechanics confirmed that the underlying physics was real, not an artifact of any single mathematical description.

The Bootstrapping Problem

Appendices A–E formalize what the RSP hierarchy does once running. Appendix I addresses a question these formalisms leave open: how the hierarchy is learned and initialized in the first place. Drawing on Dupoux et al.'s (2026) bilevel optimization framework, we formalize the evolutionary-developmental dynamics as a nested optimization. An outer loop (evolution) tunes architectural meta-parameters across generations. An inner loop (development) minimizes free energy within a single lifetime, with critical period gates regulating plasticity at each level. The result is an extended master functional and a bilevel developmental algorithm that completes the RSP formalization.

Appendix H derives a formal theory of intelligence from the same master functional. Under the constitutive-identity premise defended in Part I, this gives a precise specification of the architectural condition identified with consciousness at Level 3b, and generates cross-species predictions about which organisms should have which levels.

What This Paper Does Not Do

This paper does not argue for the philosophical dissolution of the hard problem (that is Part I) or present the empirical evidence (Part II). It provides the formal machinery. The philosophical payoff is that the convergence of five independent formalisms on the same architecture is strong evidence that the architecture captures genuine structure, not an artifact of any single theoretical lens. The empirical payoff is that the formalisms generate the specific, quantitative predictions tested in Part II.

Appendix A: The RSP as a Hierarchical Predictive Coding System

Overview

This appendix formalizes the Recursive Social Prediction (RSP) model as a hierarchical predictive coding system. We specify the generative model at each level, the message-passing dynamics between levels, and the precision-weighting mechanisms that govern attention and learning. The formulation extends hierarchical predictive coding (Friston, 2005; Clark, 2013) to encompass affect, schemas, and recursive self-modeling.

A.1 Notation and Definitions

A.1.1 Core Variables

Symbol Domain Definition $t \in \mathbb{N}$ Discrete time index $\mathbf{x}(t) \in \mathbb{R}^{n_x}$ Exteroceptive sensory input vector $\mathbf{i}(t) \in \mathbb{R}^{n_i}$ Interoceptive input vector $\mathbf{a}(t) \in [-1, 1] \times [0, 1]$ Core affect vector (valence, arousal) $\mathcal{S}$ —Set of core schemas $\{s_{\text{body}}, s_{\text{spatial}}, s_{\text{aff-hom}}, s_{\text{attn}}, s_{\text{self}}\}$ $\mathbf{s}_k(t) \in \mathbb{R}^{n_k}$ State vector of schema $k \in \mathcal{S}$ $\mathbf{z}(t) \in \mathbb{R}^{n_z}$ Integrated schema representation (structured affect) $c^*(t) \in \mathbb{C}$ Selected concept maximizing posterior probability $\mathbf{e}(t) \in \mathbb{R}^{n_e}$ Constituted experience $\mathcal{M}(t) \in \mathbb{R}^{n_m}$ Self-model state vector $\mathbf{o}_j(t) \in \mathbb{R}^{n_c}$ Model of other agent j 's mental state $\hat{\mathcal{M}}_j(t) \in \mathbb{R}^{n_m}$ Predicted self-as-seen-by agent j $\boldsymbol{\varepsilon}_l(t) \in \mathbb{R}^{n_l}$ Prediction error at level l $\boldsymbol{\Pi}_l(t) \in \mathbb{R}^{n_l \times n_l}$ Precision matrix at level l $\boldsymbol{\mu}_l(t) \in \mathbb{R}^{n_l}$ Sufficient statistics at level l

A.1.2 Level Indices

Level Label Generative Model Primary Substrate $l = 1$ Core Affect $g_1(\boldsymbol{\mu}_2)$: predicts interoceptive signals Vagal/lamina-I afferents $\rightarrow$ NTS $\rightarrow$ parabrachial nucleus $\rightarrow$ ventromedial thalamus $\rightarrow$ anterior insula and ACC; with hypothalamus, PAG, and central amygdala (Craig 2002, 2009; Critchley et al. 2004; Damasio 1999) $l = 2$ Core Schemas $g_2(\boldsymbol{\mu}_3)$: predicts schema configurations Body schema: posterior parietal + somatosensory + cerebellum. Spatial schema: hippocampus + entorhinal + retrosplenial. Attention schema: anterior TPJ + intraparietal sulcus + dACC. Affective-homeostatic: anterior/middle insula + OFC + hypothalamus. Bodily/minimal self: right inferior parietal + posterior insula + posterior TPJ + S1 $l = 3a$ Conceptual categorization g_{3a} : predicts categorical structure Lateral temporal + anterior temporal lobe (ATL) + IFG (semantic/conceptual networks) $l = 3b$ Recursive self-modeling $g_3(\boldsymbol{\mu}_3^{\text{rec}})$: predicts

experience from recursive self-modeling network: mPFC, posterior TPJ, posterior STS, PCC/precuneus; default mode network architecture (Frith & Frith 2006; Saxe & Kanwisher 2003; Schurz et al. 2014; Buckner et al. 2008)

A.2 The Generative Model

A.2.1 Hierarchical Structure

The RSP generative model is a three-level hierarchy where each level l generates predictions about the level below:

$$\text{Level 1: } \mu_1(t) = g_1(\mu_2(t)) + \mathbf{\Pi}_1^{-1/2} \omega_1$$

$$\text{Level 2: } \mu_2(t) = g_2(\mu_3(t)) + \mathbf{\Pi}_2^{-1/2} \omega_2$$

$$\text{Level 3: } \mu_3(t) = g_3(\mu_3^{\text{rec}}(t)) + \mathbf{\Pi}_3^{-1/2} \omega_3$$

where $\omega_l \sim \mathcal{N}(0, \mathbf{I})$ and μ_3^{rec} denotes the recursively updated self-model.

A.2.2 Prediction Errors and Updates

Prediction errors encode discrepancies between actual input and top-down prediction:

$$\varepsilon_l(t) = \text{input}_l(t) - g_l(\mu_{l+1}(t))$$

Precision-weighted errors propagate upward:

$$\xi_l(t) = \mathbf{\Pi}_l(t) \cdot \varepsilon_l(t)$$

The update rule for each level follows gradient descent on free energy:

$$\dot{\mu}_l = \xi_{l-1} - \frac{\partial g_l}{\partial \mu_l} \xi_l - \Sigma_l^{-1} \eta_l$$

where $\eta_l(t) = \mu_l(t) - f_l(\mu_l(t-1), \mu_{l+1}(t-1))$ is the temporal prediction error.

A.3 Level-Specific Generative Functions

A.3.1 Level 1: Core Affect Generation

Level 1 maps schema-level expectations to predicted interoceptive states. The generative function implements learned mappings from integrated schemas to visceral predictions:

$$g_1(\mu_2) = \mathbf{W}_1^{\text{aff}} \sigma(z(t)) + \mathbf{b}_1$$

Core affect emerges from precision-weighted interoceptive prediction errors:

$$\mathbf{a}(t) = \phi_{\text{aff}}(\boldsymbol{\xi}_1(t))$$

where $\phi_{\text{aff}} : \mathbb{R}^{n_i} \rightarrow [-1, 1] \times [0, 1]$ projects onto the valence-arousal plane.

A.3.2 Level 2: Schema Integration

Each schema $k \in \mathcal{S}$ maintains state vectors updated by sensory input and top-down predictions. The five schemas encode: body configuration ($\mathbf{s}_{\text{body}}$), spatial maps ($\mathbf{s}_{\text{spatial}}$), homeostatic state ($\mathbf{s}_{\text{aff-hom}}$), attentional focus ($\mathbf{s}_{\text{attn}}$), and self-other distinction ($\mathbf{s}_{\text{self}}$):

$$s_k(t) = f_k^{\text{schema}} s_k(t-1), x_k(t), \mathbf{a}(t), g_{2,k}(\boldsymbol{\mu}_3(t))$$

Schema integration combines individual states into unified structured representation:

$$\mathbf{z}(t) = \sigma(\mathbf{W}_{\text{int}} \cdot \text{concat}(\mathbf{a}(t), \mathbf{s}_1(t), \dots, \mathbf{s}_5(t)) + b_{\text{int}})$$

A.3.3 Level 3: Recursive Social Prediction

Level 3 operates through conceptual categorization followed by recursive self-modeling. Conceptual categorization selects the most probable concept given current structured affect:

$$c^*(t) = \arg \max_{c \in \mathcal{C}} P(c | \mathbf{z}(t), \text{context}(t))$$

Constituted experience integrates structured affect with selected concept:

$$\mathbf{e}(t) = \phi_{\text{exp}}(\mathbf{z}(t), \mathbf{h}_{c^*})$$

A.4 The Recursive Self-Model

A.4.1 Self-Model Dynamics

The self-model integrates constituted experience over time:

$$\mathcal{M}(t) = \mathcal{M}(t-1) + \lambda_{\mathcal{M}} \cdot \tanh(\mathbf{W}_{\mathcal{M}} \mathbf{e}(t) + b_{\mathcal{M}})$$

Other-agent models are updated based on observed behavior and inferred mental states:

$$o_j(t) = f_{o|} o_j(t-1), \text{obs}_j(t), \mathcal{M}(t)$$

A.4.2 Recursive Prediction

The recursive component predicts how others model the self. For each other-agent j , the system maintains a model of how that agent perceives the self:

$$\hat{\mathcal{M}}_j(t) = f_{\text{rec}} \mathcal{M}(t), o_j(t)$$

The recursive self-model state combines direct self-modeling with these reflected appraisals:

$$\boldsymbol{\mu}_3^{\text{rec}}(t) = \mathbf{W}_{\text{rec}} \cdot \text{concat } \mathcal{M}(t), f\hat{\mathcal{M}}_j(t)g_{j=1}^K + \mathbf{b}_{\text{rec}}$$

A.4.3 Architectural State Identified with Phenomenal State

We define the architectural state object $\Phi(t)$ that, on the constitutive-identity thesis defended in Part I, is *the phenomenal state at Level 3b. The formalism does not derive that $\Phi(t)$ is phenomenal experience — that is the substantive philosophical claim being made elsewhere in the series. What the formalism does is precisely specify the architectural object whose convergence properties (proven in subsequent sections) license the application of the constitutive identity to specific systems:*

$$\Phi(t) / \mathbf{he}(t), \mathcal{M}(t), \hat{\mathbf{e}}(t+1)_i$$

where $\mathbf{e}(t)$ is current integrated experience, $\mathcal{M}(t)$ is the recursive self-model, and $\hat{\mathbf{e}}(t+1) = g_3^{\text{temp}}(\boldsymbol{\mu}_3^{\text{rec}}(t))$ is the predicted future experience. (The notational collision with the prediction-error symbol $\boldsymbol{\epsilon}_l$ in earlier sections is resolved here by writing the experience-state vector as $\mathbf{e}(t)$ without subscript and continuing to use $\boldsymbol{\epsilon}_l$ for prediction errors throughout.)

A.5 Precision Dynamics

A.5.1 Attention as Precision Control

Precision matrices $\boldsymbol{\Pi}_l(t)$ implement attention by weighting prediction errors. High precision amplifies errors (increased attention), while low precision attenuates them:

$$\boldsymbol{\Pi}_l(t) = \text{diag } \pi_{l,1}(t), \dots, \pi_{l,n_l}(t)$$

Precision evolves based on error magnitude and higher-level expectations:

$$\dot{\pi}_{l,i} = \gamma_l \left(\psi(\hat{\mathbf{a}}_{l,i}) - \psi(\mathbf{a}_{l,i}) + \frac{1}{2}(\boldsymbol{\epsilon}_{l,i}^2 - \hat{\mathbf{a}}_{l,i}^{-1}) \right)$$

where $\psi(\cdot)$ is the digamma function and $\hat{\mathbf{a}}_{l,i}$ are expected precision parameters.

A.5.2 Affective Modulation

Core affect modulates precision across levels. High arousal increases precision globally (heightened attention), while valence biases precision toward approach-relevant (positive valence) or threat-relevant (negative valence) features:

$$\boldsymbol{\Pi}_l^{\text{mod}}(t) = \boldsymbol{\Pi}_l(t) \odot \mathbf{1} + \alpha(t) \cdot \mathbf{1} + v(t) \cdot \mathbf{w}_l^{\text{val}}$$

where $\mathbf{w}_l^{\text{val}}$ encodes valence-dependent attention biases and $\odot$ denotes element-wise multiplication.

A.5.3 Cortical Substrate for Bidirectional Message Passing

The hierarchical predictive coding architecture above requires that early sensory cortex carry not only bottom-up sensory representations but also top-down predictions that modulate those representations during perception and decision. This is a substantive empirical claim: it predicts that activity in early sensory areas during perceptual decisions should reflect both afferent sensory input and feedback from higher areas, not just a feedforward sensory cascade.

Armstrong and Vlasov (2026) provide rodent evidence consistent with this in primary somatosensory cortex (S1). Using ethological whisker-guided navigation in mice combined with dense electrophysiological recordings, they show that high-dimensional spiking activity in whisker-related S1 collapses to a single decision-correlated latent variable during sensory evidence accumulation, followed by slow synchronous ramping across the cortical column. Their analysis concludes that this ramping "reflect[s] reverberation of neural activity in cortico-cortical S1-S2" feedback loops — i.e., the decision-related activity in S1 is not generated within S1 from sensory input alone, but reflects top-down signals from secondary somatosensory cortex (S2) and other higher areas projecting back to S1.

This is consistent with what hierarchical predictive coding requires: early sensory cortex carries level-mixed representations in which top-down predictions and bottom-up errors coexist and propagate through cortico-cortical feedback loops. The reported timescales — fast ms-scale spiking variability, ramping onset at ~ 0.12 s relative to evidence-accumulation onset, ramping peak at ~ 2.0 s — reflect cortico-cortical reverberation within a single sensory hierarchy ($S1 \leftrightarrow S2$). This is within-level cortical-feedback evidence for bidirectional message passing; it is not evidence of inter-level ($\text{Level } 1 \leftrightarrow \text{Level } 2 \leftrightarrow \text{Level } 3$) cascade message passing, whose timescale separation operates across, not within, a sensory hierarchy. The result is consistent with — though it does not uniquely test — the bidirectional message-passing architecture this appendix formalizes; it does not on its own confirm the constitutive-identity claim that this same architecture, recursively elaborated at Level 3, constitutes phenomenal experience.

A complementary point applies to Level 2 specifically. Hierarchical predictive coding does not require that all levels be active simultaneously for any one to operate. The cascade structure requires inner loops to stabilize on faster timescales before outer loops can be commissioned, but inner loops can continue operating when outer loops are disrupted. Katlowitz et al. (2026)

demonstrate this empirically using Neuropixels recordings in human hippocampus during propofol-maintained general anaesthesia: hippocampal single units and local field potentials encode oddball discrimination, semantic features of natural speech, grammatical structure, and decodable features of upcoming words during clinical loss of consciousness, developing representational plasticity over approximately ten minutes of exposure. (The authors caution that future-word decoding does not by itself establish active prediction beyond contextualization; the implication for hierarchical predictive coding is constrained accordingly.) The finding is consistent with — not a confirmation of — what cascade-based hierarchical predictive coding requires: Level 2 schema-driven processing, here hippocampal semantic and grammatical encoding, operates on its own timescale and can persist when the cortico-cortical loops supporting Level 3 recursive integration are pharmacologically disrupted.

A.6 Learning and Adaptation

A.6.1 Synaptic Updates

Model parameters adapt through prediction error minimization. Connection weights follow a Hebbian-like rule modulated by precision:

$$\Delta \mathbf{W}_l = \eta_l \cdot \boldsymbol{\xi}_l(t) \boldsymbol{\mu}_{l+1}(t)^T$$

A.6.2 Schema Plasticity

Schema parameters adapt based on multimodal prediction errors:

$$\Delta \mu^k = \eta_k \sum_1 \boldsymbol{\xi}_1^{(k)}(t) \frac{\partial g_1^{(k)}}{\partial \mu^k}$$

where $\boldsymbol{\xi}_l^{(k)}$ are prediction errors relevant to schema k .

This hierarchical architecture enables the RSP model to generate multi-component phenomenal states through the dynamic interaction of affect, structured schemas, and recursive social prediction, providing a computational foundation for understanding conscious emotional experience.

Appendix B: Free Energy Formulation of RSP

Overview

This appendix derives the variational free energy functional for the Recursive Social Prediction (RSP) model. We show how each level contributes to a scalar free energy bound, how minimizing this bound implements perception and learning, and how the recursive self-model enters as a self-referential term.

Notation. μ_l denotes expectations at level l ; Π_l denotes precision matrices; $g_l(\cdot)$ denotes generative functions; $\epsilon_l = \text{input}_l - g_l(\mu_{l+1})$ denotes prediction errors.

B.1 The Variational Free Energy Principle

Disambiguation. Throughout this paper "free energy" refers to variational free energy in the Friston (2010) sense — an information-theoretic upper bound on negative log-evidence (surprisal), measured in nats — and never to thermodynamic (Helmholtz or Gibbs) free energy, which is a different quantity with different units (joules) and a different physical interpretation. Any reader familiar with statistical mechanics should treat F in this paper as an information-theoretic functional, not a thermodynamic potential. The Markov-blanket formalism sometimes used to derive variational free energy is not* invoked here; we use the standard Bayesian-inference formulation directly to avoid the controversies in the Markov-blanket literature (Biehl et al. 2021; Bruineberg et al. 2022).

An organism maintains a generative model $p(y, \cdot)$ of sensory observations $\mathbf{y}$ and hidden causes $\cdot$. Since exact Bayesian inference is intractable, the system maintains an approximate posterior $q(\cdot)$ and minimizes the **variational free energy**:

$$F = -\ln p(\mathbf{y}) + D_{\text{KL}}[q(\cdot) \parallel p(\cdot | \mathbf{y})]$$

Since $D_{\text{KL}} \geq 0$, free energy bounds surprisal: $F \geq -\ln p(\mathbf{y})$. Under the Laplace approximation, $q(\cdot) = \mathcal{N}(\mu, \Sigma)$, and free energy simplifies to prediction error terms.

SymbolDefinition F Total variational free energy (scalar) F_l Free energy contribution from level l
 $\mathbf{y}$ Complete sensory observation: $\mathbf{y} = [\mathbf{x}(t), \mathbf{i}(t)]$ All hidden causes: $\mu_1, \mu_2, \mu_3, \mu_3^{\text{rec}}$ g

B.2 Free Energy Decomposition

Under Gaussian noise assumptions, total free energy decomposes as:

$$F = F_1 + F_2 + F_3 + F_{\text{temporal}}$$

Each level contributes:

$$F_l = \underbrace{\frac{1}{2} \boldsymbol{\varepsilon}_l^T \boldsymbol{\Pi}_l \boldsymbol{\varepsilon}_l}_{\text{accuracy}} - \underbrace{\frac{1}{2} \ln j \boldsymbol{\Pi}_l}_{\text{complexity}} + \text{const}$$

The **accuracy** term penalizes prediction errors weighted by precision; the **complexity** term prevents overconfidence.

Temporal free energy penalizes deviations from predicted state transitions:

$$F_{\text{temporal}} = \sum_{l=1}^3 \frac{1}{2} \boldsymbol{\eta}_l^T \boldsymbol{\Sigma}_l^{-1} \boldsymbol{\eta}_l - \frac{1}{2} \ln j \boldsymbol{\Sigma}_l^{-1}$$

where $\boldsymbol{\eta}_l(t) = \boldsymbol{\mu}_l(t) - f_l(\boldsymbol{\mu}_l(t-1), \boldsymbol{\mu}_{l+1}(t-1))$.

B.3 Level-Specific Free Energy

B.3.1 Level 1: Core Affect

$$F_1 = \frac{1}{2} [i(t) - g_1(\boldsymbol{\mu}_2)^*]^T \boldsymbol{\Pi}_1 [i(t) - g_1(\boldsymbol{\mu}_2)^*] - \frac{1}{2} \ln j \boldsymbol{\Pi}_1$$

Interpretation. F_1 is high when interoceptive signals deviate from schema predictions. High $\boldsymbol{\Pi}_1$ represents interoceptive sensitivity. Affective valence tracks the free energy gradient:

$$v(t) \approx -\text{sign} \left(\frac{\partial F_1}{\partial \boldsymbol{\mu}_2} \right) \cdot \left\| \frac{\partial F_1}{\partial \boldsymbol{\mu}_2} \right\|$$

B.3.2 Level 2: Schema

$$F_2 = \frac{1}{2} [\boldsymbol{\mu}_1 - g_2(\boldsymbol{\mu}_3)^*]^T \boldsymbol{\Pi}_2 [\boldsymbol{\mu}_1 - g_2(\boldsymbol{\mu}_3)^*] - \frac{1}{2} \ln j \boldsymbol{\Pi}_2$$

This decomposes across five schemas:

$$F_2 = \sum_{k \in \mathcal{S}} F_{2,k} + F_{2,\text{Cross}}$$

$$\text{where } F_{2,k} = \frac{1}{2} [s_k(t) - g_{2,k}(\boldsymbol{\mu}_3)^*]^T \boldsymbol{\Pi}_{2,k} [s_k(t) - g_{2,k}(\boldsymbol{\mu}_3)^*]$$

Attention as precision optimization. The attention schema modulates precision:

$$\boldsymbol{\Pi}_{2,\parallel}(t) = \boldsymbol{\Pi}_{2,\parallel}^{\text{base}} \cdot \sigma(\mathbf{W}_{\text{attn},\parallel} \cdot \mathbf{s}_{\text{attn}}(t))$$

This implements the precision-weighted computation that, under the constitutive-identity premise of Part I, is identified with the architectural substrate of Level 2 schema-organized phenomenal experience.

B.3.3 Level 3: Recursive Self-Model

$$F_3 = \frac{1}{2} [\mu_2 - g_3(\mu_3^{\text{rec}})^*]^T \Pi_3 [\mu_2 - g_3(\mu_3^{\text{rec}})^*] - \frac{1}{2} \ln j \Pi_3$$

The self-referential term. The hidden cause $\mu_3^{\text{rec}} = \mathcal{M}^{(D_{\text{max}})}(t)$ is the recursively updated self-model. The recursive structure expands as:

$$F_3^{\text{rec}} = \sum_{d=1}^{D_{\text{max}}} \gamma^d \cdot \frac{1}{2} \|\mathcal{M}^{(d)}(t) - \mathcal{M}^{(d-1)}(t)\|_{\Pi_3^{(d)}}^2$$

This formalizes the "strange loop": each depth penalizes discrepancy between self-perception and predicted other-perception.

B.4 Gradient Dynamics

Expectations follow gradient descent on total free energy:

$$\dot{\mu}_1 = D\mu_1 - \kappa_1 \cdot \frac{\partial F}{\partial \mu_1}$$

Level 1:

$$-\frac{\partial F}{\partial \mu_1} = -\Pi_2 \cdot \varepsilon_2 - \Sigma_1^{-1} \cdot \eta_1$$

Level 2:

$$-\frac{\partial F}{\partial \mu_2} = \frac{\partial g_1^T}{\partial \mu_2} \Pi_1 \varepsilon_1 - \Pi_3 \varepsilon_3 - \Sigma_2^{-1} \eta_2$$

Level 2 is pulled by three forces: explaining interoceptive errors, reducing self-model error, and temporal consistency.

Level 3:

$$-\frac{\partial F}{\partial \mu_3} = \frac{\partial g_2^T}{\partial \mu_3} \Pi_2 \varepsilon_2 + \frac{\partial g_3^T}{\partial \mu_3} \Pi_3 \varepsilon_3 - \Sigma_3^{-1} \eta_3$$

B.5 Expected Free Energy and Developmental Ordering

B.5.1 Expected Free Energy

For action selection, the system minimizes **expected free energy**:

$$G(\pi) = \mathbb{E}_{q(y_{1:T}, \pi)} \left[\sum_{t=1}^T F_t \right]$$

where π is a policy (action sequence). This decomposes into:

$$G(\pi) = \underbrace{\mathbb{E}[F_{\text{epistemic}}]}_{\text{information gain}} + \underbrace{\mathbb{E}[F_{\text{pragmatic}}]}_{\text{preference satisfaction}}$$

Epistemic value drives exploration to reduce uncertainty. **Pragmatic value** drives exploitation toward preferred states.

B.5.2 Developmental Ordering Constraint

Theorem B.1 (Architectural Ordering Constraint). *Under the RSP free energy functional, equipped with the precision-gating relation $\Pi_{l+1} = \Pi_{l+1}^{\text{base}} \cdot \exp(-\text{ff} \cdot \text{h}F_l \mathbf{i})$, the architecture admits stable functional connectivity at Level $l + 1$ only after Level l achieves sustained free energy reduction ($\text{h}F_l \mathbf{i} < F_l^{\text{threshold}}$).*

Proof sketch. The precision-gating relation entails $\Pi_{l+1} \propto 0$ when $\text{h}F_l \mathbf{i}$ is high, making Level $l + 1$ insensitive to its inputs. Only when Level l achieves sustained $\text{h}F_l \mathbf{i} < F_l^{\text{threshold}}$ does Level $l + 1$ acquire sufficient precision for stable functional connectivity to be admitted by the formalism. ■

Note on the architectural-vs-causal reading. This theorem derives an architectural *consistency constraint* on the formalism — Level $l + 1$ cannot have stable connectivity in the formalism while Level l free energy remains high — given the precision-gating relation. It does not derive that Level l causally produces Level $l + 1$ in the empirical-development sense. The empirical claim that human development follows the ordering Level $1 \rightarrow 2 \rightarrow 3$ (Part II) is consistent with the architectural constraint but is established by developmental data, not by the theorem alone. The precision-gating relation itself is an architectural assumption motivated by the variational free-energy framework; alternative gating relations would yield different orderings.

Corollary B.1. Subject to the precision-gating assumption above, the RSP hierarchy admits stable connectivity in the sequence: Core Affect $\rightarrow$ Schemas $\rightarrow$ Recursive Self-Model.

This ordering is consistent with the empirical observation that affective development precedes cognitive development, and that theory of mind emerges after basic schema formation (Part II).

B.6 The Fisher Metric: A Measurement-Theory Extension

This section records a consistency statement about the geometry already implicit in the precision-weighted machinery above; it does not introduce new geometric structure. Da Costa et al. (2024) propose that the Fisher information metric on belief space provides a "Rosetta stone" between first-person phenomenological reports and third-person neural dynamics. We show below that, for the Gaussian belief families used throughout this appendix, this Fisher metric coincides with the precision matrix $\mathbf{\Pi}_l$ that already governs RSP's free-energy norms.

For a parametric family of beliefs $q(\boldsymbol{\mu}, \boldsymbol{\gamma})$ with parameter vector $\boldsymbol{\gamma}$, the Fisher information matrix is:

$$G_{ij}(\boldsymbol{\gamma}) = \mathbb{E}_q \left[\frac{\partial \ln q}{\partial \theta_i} \frac{\partial \ln q}{\partial \theta_j} \right]$$

(We write $\mathcal{G}$ rather than $\mathcal{I}$ to avoid collision with the generative function $g_l(\cdot)$.) This induces a Riemannian metric on the manifold of belief states, with two useful properties: it is reparameterization-invariant, and it measures information distinguishability — two belief states are far apart in Fisher distance to the extent that they predict observably different data.

The compatibility result. For Gaussian beliefs $q(\boldsymbol{\mu}_l) = \mathcal{N}(\boldsymbol{\mu}_l, \hat{\boldsymbol{\mu}}_l, \mathbf{\Pi}_l^{-1})$ parameterized by their means $\boldsymbol{\gamma}_l = \hat{\boldsymbol{\mu}}_l$ with fixed precisions, direct computation gives $\mathcal{G}(\hat{\boldsymbol{\mu}}_l) = \mathbf{\Pi}_l$. The Fisher metric on each level's belief space is therefore the precision matrix the framework already uses. The precision-weighted prediction-error norms $\frac{1}{2} \boldsymbol{\varepsilon}_l^T \mathbf{\Pi}_l \boldsymbol{\varepsilon}_l$ that appear in §B.2–§B.3 are squared Fisher distances in mean-parameter coordinates. The introduction of $\mathcal{G}$ here is a consistency statement, not the introduction of a new geometric object. If precision is itself a parameter (i.e., $\boldsymbol{\gamma}_l = (\hat{\boldsymbol{\mu}}_l, \mathbf{\Pi}_l)$), the Fisher metric acquires a second block on the precision parameters and is no longer the same as $\mathbf{\Pi}_l$ alone; the framework's precision dynamics then traverse this extended metric. The full Gaussian-family Fisher geometry is well-understood (see Amari, Information Geometry*, chs. 2–3); the relevant point is only that nothing here contradicts what §B.2 already commits to.

The bridge claim (honest-claude flag). Da Costa et al.'s further methodological proposal is that this geometry is what one should use to compare phenomenological reports to neural dynamics — that *phenomenal distance* between two experiences could be modeled as the Fisher distance between the corresponding belief states. This is an *interpretive bridge*, not a derived result. RSP does not establish it; Da Costa et al. do not derive it from the free energy principle alone; it is an additional methodological commitment about what the right measurement is. The framework above is consistent with the bridge but does not require it. Read this way, Da Costa et al. supply a candidate *measurement theory* for the kind of generative model RSP specifies

architecturally: the two are complementary, addressing measurement and architecture as different questions about the same formal object.

B.7 Summary

The RSP free energy formulation provides:

Unified dynamics: All three levels minimize the same scalar functional

Architectural ordering constraint: Under the precision-gating relation, higher levels admit stable functional connectivity only after lower levels stabilize

Self-reference: Level 3 has the property that its predictions involve its own state via the recursive self-modeling structure

Precision-attention coupling: Attention is identified with precision-weighting on prediction errors, determining which prediction errors propagate through the hierarchy

Under the constitutive-identity premise of Part I, the architectural condition identified with Level 3b phenomenal consciousness is: the recursive self-model at Level 3b achieves sufficient precision Π_3 to gate top-down control on schema-level predictions, with F_3 in a stable basin. This is a *condition on the architecture*, not a separate mechanism by which consciousness is produced; under the constitutive premise, the obtaining of this condition just is Level 3b phenomenal consciousness in the corresponding system.

Appendix C: Intrinsically Motivated Reinforcement Learning in RSP

Overview

This appendix recasts the Recursive Social Prediction (RSP) model in reinforcement learning (RL) terms. We show that valence acts as a reward signal, schemas as learned state representations, and Level 3 recursive prediction as model-based planning. The formulation reveals that the RSP hierarchy implements **intrinsically motivated model-based reinforcement learning** where the agent learns both reward function and state representation through intrinsic motivation.

C.1 RL Notation and RSP-to-RL Mapping

C.1.1 New Symbols

Symbol Domain Definition $r(t)$ $\mathbb{R}$ Scalar reward signal at time t $\mathbf{s}^{\text{RL}}(t)$ $\mathbb{R}^n$ RL state vector (agent's representation of its situation) $\pi(\mathbf{u} | \mathbf{s}^{\text{RL}})$ $[0, 1]$ Policy: probability of action $\mathbf{u}$ given state $\mathbf{s}^{\text{RL}}$ $V(\mathbf{s}^{\text{RL}})$ $\mathbb{R}$ Value function: expected cumulative future reward $\delta(t)$ $\mathbb{R}$ Temporal difference (TD) error $\gamma_{\text{RL}}(0, 1)$ Discount factor for future rewards $T(\mathbf{s}^{\text{U}} | \mathbf{s}^{\text{RL}}, \mathbf{u})$ $[0, 1]$ Transition model $R(\mathbf{s}^{\text{RL}})$ $\mathbb{R}$ Reward model

C.1.2 RSP-to-RL Mapping

RL Concept RSP Equivalent Explanation Reward $r(t)$ Valence $v(t)$ Interoceptive prediction errors, signed by homeostatic direction State $\mathbf{s}^{\text{RL}}(t)$ Schema integration $\mathbf{z}(t)$ Five-schema structured representation Transition model $\mathcal{T}$ Level 3 generative model $\mathcal{P}$ Predicts schema changes given self-model and actions Reward model $\mathbb{R}$ Level 1 generative model $\mathcal{I}$ Predicts interoceptive consequences Model-based planning Strange loop recursion Simulating social consequences before acting Exploration Epistemic foraging Seeking states that reduce uncertainty Value function V Expected free energy $-G$ States with low G are high-value

C.2 Valence as Reward

C.2.1 The Reward Signal

In standard RL, reward is supplied externally. In RSP, reward is generated internally through interoceptive prediction:

$$r(t) = v(t) = \phi_v \xi_1(t)$$

where $\phi_v : \mathbb{R}^{n_i} \rightarrow [-1, 1]$ projects precision-weighted interoceptive prediction error onto the valence dimension. This projection is learned, not innate.

Why this is intrinsically motivated: No external agent specifies what counts as rewarding. The reward function emerges from generative model structure:

Predictions consistent with homeostasis $\rightarrow$ small $\varepsilon_1 \rightarrow$ positive valence

Homeostatic disruption $\rightarrow$ large $\varepsilon_1 \rightarrow$ negative valence

C.2.2 Reward Decomposition

The scalar reward decomposes into schema contributions:

$$r(t) = \sum_{\|2s} w_{\|}^r \cdot r_{\|}(t) + r_{\text{cross}}(t)$$

where:

$$r_{\|}(t) = -\frac{1}{2} \varepsilon_{1,\|}^T \Pi_{1,\|} \varepsilon_{1,\|}$$

SchemaReward ContributionExample **S**bodyBody-state comfortWarmth, satiation, physical ease

SpatialEnvironmental safetyFamiliar surroundings, orientation **S**af-hOm Homeostatic

balanceBlood sugar stability, temperature regulation **S**attnCognitive engagementFocused

attention, flow states **S**selfAgency, social belongingCompetence, autonomy, relatedness

C.2.3 Temporal Difference Implementation

The TD error in RSP is:

$$\delta(t) = r(t+1) + \gamma_{\text{RL}} \cdot V(\mathbf{z}(t+1)) - V(\mathbf{z}(t))$$

The neural substrate of TD errors is canonically associated with phasic dopamine firing in mesencephalic dopamine neurons (Schultz 1998; Schultz, Dayan & Montague 1997). Note the distinction between *valence* and *reward prediction error (RPE)*: phasic dopamine signals RPE, not valence directly — these dissociate (e.g., reward omission produces negative RPE but neutral valence; aversive prediction error involves distinct dopamine subpopulations and lateral habenula circuitry; Matsumoto & Hikosaka 2009). The mapping in this paper from δ_1 to VTA dopamine therefore should be read as the RPE *signal driving learning over the valence-grounded reward function* $r(t) = v(t)$, *not as a direct neural signature of valence itself. The hypothesis advanced is that dopaminergic TD-RPE signals correlate with interoceptive-prediction-error gradients, not with raw valence values.*

C.3 Schemas as Learned State Representations

C.3.1 The Representation Learning Problem

The RSP model's five schemas constitute the agent's learned state representation:

$$\mathbf{s}^{\text{RL}}(t) = \mathbf{z}(t) = \text{Integrate } \mathbf{a}(t), f\mathbf{s}_k(t)g_{k2S}$$

This representation develops ontogenetically: body schema first (infancy), spatial schema next, affective homeostasis in parallel, attention as executive functions mature, and self-other distinction last.

C.3.2 Auxiliary Tasks and Schema Learning

Each schema solves an auxiliary task providing useful gradients:

SchemaAuxiliary TaskLearning Signal
 $\mathbf{s}_{\text{body}}$ Predict proprioceptive next-stateBody movement prediction error
 $\mathbf{s}_{\text{spatial}}$ Predict visual-spatial next-stateNavigation prediction error
 $\mathbf{s}_{\text{aff-hom}}$ Predict homeostatic next-stateInteroceptive prediction error
 $\mathbf{s}_{\text{att}}$ Predict which inputs matterSalience prediction error
 $\mathbf{s}_{\text{self}}$ Predict social consequencesAgency prediction error

C.3.3 State Abstraction

The integration step performs state abstraction, compressing high-dimensional sensory input into a compact representation. The integrated schema vector $\mathbf{z}(t)$ satisfies:

Sufficiency: $\mathbf{z}(t)$ retains information to predict future rewards

Minimality: $\mathbf{z}(t)$ discards irrelevant information

The attention schema $\mathbf{s}_{\text{att}}$ determines what passes through this information bottleneck.

C.4 Level 3 as Model-Based Planning

C.4.1 Model-Based vs. Model-Free

The RSP hierarchy implements both:

RL TypeRSP LevelMechanism
Model-freeLevel 1 → Level 2Cached affect-to-action mappings
Model-basedLevel 3Recursive simulation of social consequences
HybridCross-levelModel-based predictions guide model-free learning

C.4.2 Planning via Recursive Simulation

Level 3 planning iteratively simulates:

Step 1. Predict action consequences on schema state:

$$\hat{z}(t+1 \mid \mathbf{u}_{\text{cand}}) = g_3^{\text{dyn}}(z(t), \mathbf{u}_{\text{cand}}, \boldsymbol{\mu}_3^{\text{rec}}(t))$$

Step 2. Predict reward of resulting state:

$$\hat{r}(t+1 \mid \mathbf{u}_{\text{cand}}) = R(\hat{\mathbf{z}}(t+1))$$

Step 3. Predict others' responses:

$$\hat{o}(t+1 \mid \mathbf{u}_{\text{cand}}) = f_o^{\text{pred}}(o(t), \mathbf{u}_{\text{cand}})$$

Step 4. Predict recursive self-model update:

$$\hat{\mathcal{M}}^{(d)}(t+1 \mid \mathbf{u}_{\text{cand}}) = \text{RecurseUpdate}(\hat{\mathcal{M}}^{(d-1)}, f_{\hat{o}_j}(t+1)g$$

Step 5. Evaluate expected free energy:

$$G(\mathbf{u}_{\text{cand}}) = \hat{F}_1 + \hat{F}_2 + \hat{F}_3 + \hat{F}_3^{\text{rec}}$$

Step 6. Select action minimizing expected free energy:

$$\mathbf{u}^*(t) = \arg \min_{\mathbf{u}_{\text{cand}}} G(\mathbf{u}_{\text{cand}})$$

C.4.3 Social Reward Shaping

The recursive self-model introduces social reward shaping:

$$r_{\text{social}}(t) = - \sum_{d=1}^{D_{\text{max}}} \gamma^d \cdot \frac{1}{2} \|\mathcal{M}^{(d)}(t) - \mathcal{M}^{(d-1)}(t)\|_{\Pi_3}^2$$

The total reward becomes:

$$r_{\text{total}}(t) = r(t) + \gamma_{\text{social}} \cdot r_{\text{social}}(t)$$

This ensures the agent values both homeostatic comfort and social coherence, explaining why humans endure physical discomfort to maintain social standing.

C.5 The Intrinsically Motivated Learning Pipeline

C.5.1 Three Phases of Learning

The RSP model's learning comprises three nested RL problems:

Phase 1: Reward Function Learning (Level 1)

$$\hat{r}(t) = \phi_v \Pi_1(t) \cdot [i(t) - g_1(\boldsymbol{\mu}_2(t))]$$

Learn what counts as rewarding through interoceptive prediction. Timescale: milliseconds.

Phase 2: State Representation Learning (Level 2)

$$\mathbf{z}(t) = \text{Integrate } \mathbf{a}(t), f_k^{\text{schema}}(\cdot) g_k \mathcal{S}$$

Learn to extract reward-relevant information from sensory input. Timescale: seconds to minutes for updates; months to years for structural development.

Phase 3: World Model and Policy Learning (Level 3)

$$\pi^*(\mathbf{u}_{jz}, \mu_3^{\text{rec}}) = \arg \min_{\pi} G(\pi)$$

Learn transition model, social model, and policy. Timescale: minutes to hours for tactical planning; years for strategic self-model revision.

C.5.2 Why "Intrinsically Motivated"

No external reward specification: $r(t) = v(t)$ derives from interoceptive prediction errors

No external state labels: Schemas are learned representations

No external curriculum: Developmental sequence emerges from learning problem structure

C.6 Schema-Decomposed Value Function

The value function decomposes across schemas:

$$V(\mathbf{z}) = \sum_{k \in \mathcal{S}} w_k^V \cdot V_k(s_k) + V_{\text{cross}}(\mathbf{z})$$

where $V_k(s_k)$ is the expected future reward from schema k 's state, and V_{cross} captures cross-schema interactions.

Level-Specific TD Errors:

Level 1 (Affective TD):

$$\delta_1(t) = r(t+1) + \gamma_{\text{RL}} \cdot V_1(\mathbf{a}(t+1)) - V_1(\mathbf{a}(t))$$

Level 2 (Schema TD):

$$\delta_2(t) = r(t+1) + \gamma_{\text{RL}} \cdot V_2(\mathbf{z}(t+1)) - V_2(\mathbf{z}(t))$$

Level 3 (Social TD):

$$\delta_3(t) = r_{total}(t+1) + \gamma_{RL} \cdot V_3(\boldsymbol{\mu}_3^{\text{rec}}(t+1)) - V_3(\boldsymbol{\mu}_3^{\text{rec}}(t))$$

The hypothesis is that these TD-RPE signals are propagated through distinct dopaminergic circuits: δ_1 via mesencephalic dopamine (VTA/SNc) projections to limbic targets; δ_2 via dorsal striatal updating of schema/habit-level value; δ_3 via mPFC mentalizing-network signaling for socially structured RPE. The dorsal-striatum mapping for "schema-level errors" is non-standard — dorsal striatum canonically encodes habit and sensorimotor TD signals (Yin & Knowlton 2006), not "schema" errors in the RSP sense — so this assignment should be read as a hypothesis to be tested, not an established mapping.

C.7 Exploration via Epistemic Foraging

Expected free energy naturally decomposes into:

$$G(\boldsymbol{\pi}) / \underbrace{G_{\text{pragmatic}}(\boldsymbol{\pi})}_{\text{seek reward}} + \underbrace{G_{\text{epistemic}}(\boldsymbol{\pi})}_{\text{reduce uncertainty}}$$

At Level 3, epistemic value becomes social curiosity:

$$G_{\text{epistemic}}^{\text{social}}(\boldsymbol{\pi}) = \mathbb{E} \left[\sum_{j=1}^{J_O} D_{\text{KL}} [q(o_j(t+\tau) | \mathbf{y}_{t+\tau}) q(o_j(t))^{*}] \right]$$

This drives understanding of others' mental states through social exploration.

Exploration-Exploitation Trade-off:

Exploitation: Maintain current self-model (conformity)

Exploration: Update self-model through novel actions (risk-taking)

Balance governed by self-model precision Π_3 : high precision favors exploitation, low precision favors exploration.

C.8 Summary: The RL Architecture

The RSP model implements a complete RL system:

Level Role Key Function TD Error **Level 3** Model-based planner

$\approx^*(\mathbf{u}_j, \boldsymbol{\mu}, \mathcal{M}^{\text{rec}}) = \arg \min_{\mathbf{u}} \mathcal{G}(\mathbf{u}) \delta_3$ (social validation) **Level 2** State representation learner

$\boldsymbol{\mu}(t) / \text{Integrate}(\{s_k(t)\}) \delta_2$ (situation assessment) **Level 1** Reward function learner

$r(t) = \phi_v(\boldsymbol{\xi}_1(t)) \delta_1$ (affective surprise)

This architecture bridges affective neuroscience and computational RL, showing how consciousness is constituted by the interaction of intrinsic motivation, learned representations,

and recursive social modeling. The key insight is that consciousness is not a separate system but the natural result of an RL agent that learns its own reward function, discovers structured state representations, and plans by simulating social consequences—all through intrinsic motivation.

Appendix D: Control-Theoretic Dynamics of the Strange Loop

Overview

This appendix models the Recursive Social Prediction (RSP) hierarchy as a multi-loop feedback control system. We specify reference signals, comparators, gains, and transfer functions at each level, analyze stability conditions for the strange loop, and show how control-theoretic concepts characterize the dynamics of consciousness, emotion regulation, and psychopathology.

The formulation draws on Perceptual Control Theory (Powers, 1973, 2005), modern control theory (Åström & Murray, 2021), and hierarchical control perspectives in cognitive science (Carver & Scheier, 1998).

D.1 Control-Theoretic Notation

D.1.1 New Symbols

Symbol **Domain** **Definition**
 $r_l(t) \in \mathbb{R}^{n_l}$ Reference signal (set-point) at level l
 $\mathbf{p}_l(t) \in \mathbb{R}^{n_l}$ Perceptual signal (measured output) at level l
 $\mathbf{e}_l^{\text{ctrl}}(t) \in \mathbb{R}^{n_l}$ Control error: $r_l - \mathbf{p}_l$
 $K_l, K_l^I, K_l^D \in \mathbb{R}^{n_l \times n_l}$ Proportional, integral, derivative gain matrices
 $G_l(s), C_l(s) \in \mathbb{C}^{n_l \times n_l}$ Plant and controller transfer functions
 $L_l(s) \in \mathbb{C}^{n_l \times n_l}$ Loop transfer function: $L_l = C_l G_l$
 $\tau_l \in \mathbb{R}_{>0}$ Time constant of level l dynamics
 $\phi_m, g_m \in \mathbb{R}$ Phase margin (degrees), gain margin (dB)

D.1.2 RSP-to-Control Mapping

Control Concept **RSP Equivalent**
Reference signal r_l Top-down prediction from level $l + 1$: $g_{l+1}(\mu_{l+1})$
Perceptual signal $\mathbf{p}_l$ Bottom-up input: actual state at level l
Control error $\mathbf{e}_l^{\text{ctrl}}$ Prediction error ϵ_l (sign reversed)
Controller output $\dot{\mu}_l$ or action $\mathbf{u}$
Plant $\text{Body} + \text{environment} + \text{social world}$

D.2 Cascade Control Architecture

The RSP hierarchy implements cascade control with nested feedback loops where outer (slow) loops set references for inner (fast) loops:

Level 3 Self-concept $r_3^e / r_3 - p_3$ Schema set-points $r_3^e \tau_3 \xi$ minutes–hours
Level 2 Schema targets $r_2^e / r_2 - p_2$ Homeostatic set-points $r_2^e \tau_2 \xi$ seconds–minutes
Level 1 Homeostatic targets $r_1^e / r_1 - p_1$ Motor/autonomic action $u(t) \tau_1 \xi$ milliseconds–seconds

D.2.1 Level-Specific Control Laws

Level 1: Homeostatic Controller (PID)

$$u(t) = K_1 \cdot e_1^{\text{ctrl}}(t) + K_1^I \int_0^t e_1^{\text{ctrl}}(\tau) d\tau + K_1^D \cdot \dot{e}_1^{\text{ctrl}}(t)$$

The integral term ensures zero steady-state error for homeostatic variables, mapping to hypothalamic set-point mechanisms.

Level 2: Schema Controller (Adaptive Gain)

$$\dot{\mu}_2(t) = K_2(t) \cdot e_2^{\text{ctrl}}(t) + K_2^D \cdot \dot{e}_2^{\text{ctrl}}(t)$$

where $K_2(t) / \Pi_2(t)$ – the gain equals the precision matrix, implementing adaptive control whose authority changes with confidence.

Level 3: Self-Model Controller (Model Predictive Control)

$$\dot{\mu}_3(t) = K_3 \cdot e_3^{\text{ctrl}}(t) + K_3^{\text{pred}} \cdot \hat{e}_3^{\text{ctrl}}(t+1)$$

where $e_3^{\text{ctrl}}(t) = \mathcal{M}^{(0)}(t) - \mathcal{M}^{(D_{\text{max}})}(t)$ is the discrepancy between direct and socially-reflected self-models.

D.3 Transfer Functions and Timescale Separation

D.3.1 Linearized Dynamics

Each level has a transfer function in the Laplace domain:

Level 1:

$$\mathcal{G}_1(s) = \frac{A_1}{s + \frac{1}{\tau_1}}, \quad C_1(s) = K_1 + \frac{K_1^I}{s} + K_1^D s$$

Loop Transfer Function:

$$L_1(s) = C_1(s) \mathcal{G}_1(s) = \frac{(K_1 s + K_1 + K_1 \frac{D_1^2}{s})}{s(s + \frac{1}{\tau_1})} A_1$$

D.3.2 Timescale Separation Condition

For cascade stability:

$$\tau_1 \ll \tau_2 \ll \tau_3$$

Typical biological values:

Level Time Constant Bandwidth Process $l = 10.5\text{--}250.5\text{--}2\text{ Hz}$ Autonomic reflexes $l = 25\text{--}30\text{ s}$
 0.03–0.2 Hz Schema updates $l = 360\text{--}3600\text{ s}$ 0.0003–0.017 Hz Self-model revision
 The ratio $\tau_{l+1}/\tau_l \geq 5$ ensures each controller sees the inner loop as instantaneous.

D.4 Stability Analysis

D.4.1 Local Stability

The linearized three-level system:

$$\dot{\boldsymbol{\mu}} = \mathbf{A} \cdot \boldsymbol{\mu} + \mathbf{B} \cdot \mathbf{r}$$

where the Jacobian is:

$$\mathbf{A} = \begin{pmatrix} -K_1 G_1^0 & J_{12} & 0 \\ J_{21} & -K_2 G_2^0 & J_{23} \\ 0 & J_{32} & -K_3 G_3^0 + J_{33}^{\text{rec}} \end{pmatrix}$$

Stability condition: All eigenvalues must have negative real parts:

$$\text{Re}(\lambda_i(\mathbf{A})) < 0 \quad \forall i$$

D.4.2 Strange Loop Stability

The recursive self-model creates positive feedback. For stability:

$$\rho(J_{33}^{\text{rec}}) < \sigma_{\min} \left(K_3 \cdot \frac{\partial g_3}{\partial \boldsymbol{\mu}_3^{\text{rec}}} \right)$$

The rate of self-amplification through recursive reflection must be less than the rate of prediction error correction.

Convergence condition: The recursion $\mathcal{M}^{(d)} = (1 - \beta_d)\mathcal{M}^{(d-1)} + \beta_d \hat{\mathcal{M}}^{(d)}$ converges if:

$$\beta_d / \beta_0 \leq 1 \quad \text{with} \quad \beta_0 < 1, \quad \beta_d < 1$$

D.4.3 Stability Margins

Phase and gain margins quantify robustness:

*LevelMin ϕ_m Min g_m Interpretation $l = 1 > 45 > 6$ dBStable autonomic regulation $l = 2 > 30 > 4$ dB
Schema convergence without rumination $l = 3 > 20 > 3$ dB
Self-model stability under social perturbation*

D.5 Disturbance Rejection and Emotion Regulation

D.5.1 Sensitivity Functions

The sensitivity function determines disturbance rejection:

$$S_l(s) = \frac{1}{1 + L_l(s)}$$

$|S_l(j\omega)| < 1$: disturbances attenuated (effective regulation)

$|S_l(j\omega)| > 1$: disturbances amplified (regulation failure)

The fundamental constraint $S_l(s) + T_l(s) = 1$ means perfect disturbance rejection and perfect reference tracking cannot coexist — the control-theoretic formulation of the flexibility-stability trade-off.

D.5.2 Emotion Regulation as Control Reconfiguration

Different emotion regulation strategies modify the control architecture: reappraisal changes schema references (r_2), suppression reduces autonomic gains (K_1), attention deployment modifies precision matrices (Π_2), and acceptance increases sensitivity bandwidth. Reappraisal is most effective because it operates at Levels 2-3, changing the reference signals that lower loops track.

D.6 Pathological Dynamics

D.6.1 Instability Modes (Speculative Clinical Analogs)

Control-system failures admit suggestive analogies with psychopathological patterns. These analogies are formal metaphors, not mechanistic claims: clinical disorders are multi-causal and operate on different timescales than the cascade dynamics described here. The mappings below should be read as one architectural facet among many, not as one-to-one mechanistic explanations.

Limit Cycle Oscillation ($\phi_m \omega 0$): Self-model oscillates without converging:

$$\mathcal{M}(t) \propto \mathcal{M}^* + A \sin(\omega_{osc} t)$$

Speculative analog: **Bipolar oscillation phenomenology**. Note that clinical bipolar disorder involves switches over days to weeks with circadian, monoaminergic, and HPA-axis components — far slower than the minutes-to-hours timescale of the limit-cycle here. The formal mapping captures only the oscillatory phenomenology, not the underlying neuropsychiatric mechanism.

Integrator Windup (K_f too large): Persistent error accumulation leads to extreme responses.

Speculative analog: **Anxiety-disorder phenomenology of accumulated homeostatic error**. Clinical anxiety has multiple etiologies (genetic vulnerability, learning history, neuroimmune factors); the integrator-windup formalism captures one facet but is not a complete mechanistic account.

Loss of Controllability ($\text{rank}(\mathbf{B}) < n$): Some states become unreachable.

Speculative clinical analog: **Depression-related agency deficits** — phenomenology of reduced agency. Note that depression is a multi-causal disorder involving monoaminergic, HPA-axis, and neuroinflammatory factors; the formal control-theoretic mapping captures only one aspect (reduced reachability of desired self-states) and should not be read as a complete neuropsychiatric account.

Strange Loop Divergence ($\rho \geq 1$): Recursive self-model becomes unstable:

$$\| \mathcal{M}^{(d)} \|_k \omega 1 \text{ as } d \omega 1$$

Speculative clinical analog: **Aberrant precision-weighting of self-referential predictions**, of the kind formally invoked in predictive-coding accounts of psychosis (Sterzer et al., 2018; Corlett & Fletcher, 2014). The mapping from the formal $\rho \geq 1$ condition to specific DSM categories (paranoia, narcissistic personality, etc.) is not established; clinical paranoia involves prefrontal-striatal-thalamic dysconnectivity, salience-attribution errors, and aberrant mesolimbic dopamine — most of which are not captured by the recursive-divergence formalism alone. This row should be read as a speculative analog indicating that disrupted recursion is one possible mechanism, not a complete mechanistic mapping.

D.6.2 Critical Transitions

The system exhibits critical slowing down near instability. As $\rho \omega 1$:

$$\tau_{\text{recovery}} \propto \frac{1}{1 - \rho}$$

Recovery time diverges, creating hysteresis and sudden transitions between stable and pathological states — consistent with clinical observations of mental health crises and therapeutic breakthroughs.

D.7 Therapeutic Implications

Control theory suggests intervention strategies:

Gain adjustment: *Modify controller gains (K_I) through medication or training*

Reference modification: *Change set-points (r_I) through cognitive therapy*

Plant modification: *Alter the controlled system through environmental changes*

Feedforward compensation: *Add predictive elements to reduce disturbance impact*

Cascade reconfiguration: *Modify the hierarchical structure through metacognitive interventions*

The framework predicts that interventions targeting higher levels (Level 3 self-model) will have broader effects but slower dynamics, while lower-level interventions (Level 1 autonomic) will have faster but more limited effects.

Appendix E — Unified Mathematical Framework

Overview

Appendixes A–D formalize the RSP model from four distinct mathematical traditions: hierarchical predictive coding, variational free energy, intrinsically motivated reinforcement learning, and cascade feedback control. This appendix demonstrates that these are four descriptions of one theory by identifying a single **master variational functional** from which all four formalisms derive as special cases.

E.1 Four Windows on One Architecture

The four RSP formalisms stand in analogous relationship to the equivalent formulations of classical mechanics:

Formalism	Mechanical Analogue	Primary Variable	What It Reveals
(A) Newtonian (forces)	Prediction error ϵ_l	Neural message-passing architecture	Free Energy
(B) Lagrangian (variational)	Free energy F	Global optimality, why three levels	Reinforcement Learning
(C) Hamiltonian (state-action)	Value V , policy π	Learning, planning, development	Control Theory
(D) Hamilton-Jacobi (optimal control)	Error e_l^{ctrl} , gain K_l	Stability, robustness, failure modes	

The RSP correspondences are approximate (holding under the Laplace approximation with Gaussian noise near equilibrium) rather than exact, but their convergence from independently established mathematical traditions suggests the architecture captures genuine structure.

E.2 Master Variable Correspondence

Table E.1: Unified Variable Correspondence

Unified	Pred. Coding (A)	Free Energy (B)	RL (C)	Control (D)	Interpretation
$\mathbf{x}_l, \boldsymbol{\mu}_l, \boldsymbol{\mu}_l^{\text{RL}}, \mathbf{p}_l$	Internal belief at level l	$\epsilon_l, \epsilon_l, \delta_l$ (approx.)	$e_l^{\text{ctrl}} = -\epsilon_l$	Discrepancy driving update	$\mathbf{W}_l \boldsymbol{\Pi}_l \boldsymbol{\Pi}_l$
exploration-exploitation	K_l ($K_2 / \boldsymbol{\Pi}_2$)	Error weighting / attention	η_l	temporal pred. error	F_{temp} or α term related to TD error
e_l^{ctrl}	Temporal consistency	$\mathcal{M}$ self-model state	$\boldsymbol{\mu}_3^{\text{fc}}$	Level 3 planning state	controlled state at Level 3
Recursive identity	$\mathcal{J} \sum_k \xi_l^k$	$k^2 F$	$\sum_l F_l$	V	

$$(negative\ value) \sum_l k e_l^{ctrl} k_{Kl} Objective\ to\ minimize$$

E.3 The Unified Optimization Principle

E.3.1 The Master Functional

All four appendices minimize a single variational functional:

$$\begin{aligned} \dot{\mathbf{x}}_l &= -\kappa_l \left[\frac{\partial J}{\partial \mathbf{x}_l} + \frac{\partial g_{l-1}^T}{\partial \mathbf{x}_l} \mathbf{W}_{l-1} \mathbf{e}_{l-1} - \mathbf{W}_l \mathbf{e}_l \right] \\ \mathbf{e}_l &= \mathbf{y}_l - \mathbf{W}_l \mathbf{x}_l \\ \mathbf{W}_l &= \mathbf{W}_l + \eta_l \left[\frac{\partial J}{\partial \mathbf{W}_l} + \frac{\partial g_{l-1}^T}{\partial \mathbf{W}_l} \mathbf{W}_{l-1} \mathbf{e}_{l-1} - \mathbf{e}_l \mathbf{x}_l^T \right] \end{aligned}$$

\$\$ (E.2)

where $\mathbf{e}_0 \equiv \mathbf{e}_{\text{sensory}}$ (sensory prediction error, with $\mathbf{I}$ as the identity on sensory input), $\mathbf{W}_0 \equiv \mathbf{\Pi}_{\text{sensory}}$ (sensory precision matrix), and $\mathbf{e}_4 \equiv 0$ (no level above 3).

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \left[\sum_{t=0}^T \gamma^t \left(\mathbf{e}_t^T \mathbf{W}_t \mathbf{u} + \frac{1}{2} \mathbf{u}^T \mathbf{K}_t \mathbf{u} \right) \right]$$

\$\$ (E.4)

E.4 Equivalence Proofs

Proposition E.1: The Universal Error Signal

Statement: Under the Laplace approximation with Gaussian noise, the following are equivalent up to sign conventions:

- (a) Precision-weighted prediction error: $\xi_1 / \mathbf{\Pi}_1 \mathbf{\epsilon}_1$
- (b) Gradient of accuracy term: $\partial F_1^{\text{acc}} / \partial \mathbf{\epsilon}_1 = \mathbf{\Pi}_1 \mathbf{\epsilon}_1$
- (c) Weighted control error at Level 2: $\mathbf{K}_2 \mathbf{e}_2^{ctrl} = \mathbf{\Pi}_2 \mathbf{\epsilon}_2$

Proof: (a) = (b) follows from differentiating $F_1^{\text{acc}} = \frac{1}{2} \epsilon_1^T \Pi_1 \epsilon_1$. For (c): $e_2^{\text{ctrl}} = -\epsilon_2$ and K_2 / Π_2 . *

Proposition E.2: Free Energy–Value Duality

Statement: Under conditions (i) Laplace approximation, (ii) valence as precision-weighted interoceptive error, and (iii) $\gamma_{\text{RL}} = e^{-1/\beta}$:

$$V(\mathbf{z}) \propto -\mathbb{E} \left[\sum_{\tau=0}^1 \gamma_{\text{RL}} F_1(t + \tau) \mid \mathbf{z}(t) = \mathbf{z} \right] + \text{const}$$

Proof sketch: Value function $V = \mathbb{E}[\sum \gamma_{\text{RL}}^T r(t + \tau)]$ with reward $r(t) \propto -F_1(t)$ under the approximation conditions. *

Proposition E.3: The Master Dynamics Unifies All Update Rules

Statement: The master dynamics $\dot{\mathbf{x}}_1 = -\kappa_1 \partial \mathcal{J} / \partial \mathbf{x}_1$ simultaneously implements:

- (a) Bidirectional message passing (App A)
- (b) Gradient descent on free energy (App B)
- (c) Approximate value-gradient update (App C)
- (d) Closed-loop control law (App D)

Proof: Direct substitution shows the master dynamics produces the message-passing equations (a), equals $-\kappa_1 \partial F / \partial \mu_1$ (b), approximates value gradient ascent under the duality (c), and equals the control law with $K_1 / \tau_1 \mathbf{W}_1$ (d). *

Proposition E.4: Universal Strange Loop Convergence

Statement: The recursive self-modeling converges in all formalisms under either of two sufficient conditions:

(i) *Fixed mixing weight:* With constant $\beta \in (0, 1)$ and uniform Lipschitz bound $L / \sup_d k \partial \text{SocialRec} \gamma_{\text{ectd}} \partial \mathcal{M} k < 1$, the contraction coefficient $\rho = (1 - \beta) + \beta L < 1$ and Banach's theorem guarantees convergence to a unique fixed point.

(ii) *Mixed-regime (biologically relevant):* With $\beta_d / \beta_{0,d}^d$ ($\gamma < 1$), low-depth reflections are high-fidelity contractions ($L_d < 1$ for $d \leq d_0$) while high-depth reflections may be expansive ($L_d \geq 1$ for $d > d_0$) but heavily discounted. The composed map satisfies

$\prod_{d=1}^{D_{\max}} [(1 - \beta_d) + \beta_d L_d] < 1$ provided the contractive low-depth factors outweigh the expansive high-depth factors: $\sum_{d>d_0} \beta_d (L_d - 1) < \left| \sum_{d \leq d_0} \ln[(1 - \beta_d) + \beta_d L_d] \right|$.

Proof sketch: Case (i) is standard Banach. For case (ii), $\ln \prod = \sum_d \ln[1 + \beta_d (L_d - 1)]$. For contractive depths ($L_d < 1$), each term is negative. For expansive depths, $\ln(1 + x) \leq x$, so the expansive contribution is bounded by $\sum_{d>d_0} \beta_d (L_d - 1) \leq \beta_0 \gamma^{d_0} (L_{\max} - 1) / (1 - \gamma) < 1$. The product is less than 1 when the contractive sum dominates. *

E.5 Emergent Properties of the Unified Framework

E.5.1 Coordinate Covariance

The accuracy terms in $\mathcal{J}$ are invariant under smooth reparametrizations $\mathbf{x}_l \rightarrow h_l(\mathbf{x}_l)$ with appropriate transformation of precision matrices. This expresses mathematically that consciousness depends on relational structure, not specific neural codes.

E.5.2 Precision as Universal Currency

Across all formalisms, precision is the universal mediating variable:

Formalism Role of Precision
 Predictive Coding (A) Weights prediction errors in message passing
 Free Energy (B) Balances accuracy vs. complexity
 RL (C) Controls exploration vs. exploitation
 Control (D) Sets controller gain; determines stability

Attention — the mechanism modulating precision — simultaneously determines what reaches consciousness (A), what is learned (B), what is valued (C), and what is stable (D).

E.5.3 Three-Level Convergence Argument

All four formalisms independently conclude three hierarchical levels are necessary:

Free energy: Cannot simultaneously minimize interoceptive and self-model surprise without mediating Level 2

RL: Developmental ordering requires reward $\rightarrow$ state $\rightarrow$ model sequence

Control: Cascade stability requires timescale separation $\tau_{l+1} / \tau_l \geq 5$

Predictive coding: Three levels minimum for interoceptive, schematic, and self-referential scales

E.6 Where the Formalisms Diverge

E.6.1 Divergent Predictions

Domain Pred. Coding (A) Free Energy (B) RL (C) Control (D) Neural circuits Error vs. representation neurons separable Agnostic about circuits Dopaminergic TD signals Gain-modulated feedback Psychopathology Aberrant precision Excessive complexity cost Reward dysfunction Instability/oscillation Development Hierarchical construction Sequential free energy reduction Reward $\rightarrow$ state $\rightarrow$ model Inner loops first

E.6.2 Non-Equivalent Regimes

Far from equilibrium: Laplace approximation fails; RL handles large fluctuations better

Discrete actions: RL natural; control theory continuous; others agnostic

Model learning vs. use: RL cleanly separates; others blur distinction

Stability analysis: Only control theory provides phase/gain margins

E.7 Algorithm 1: RSP Processing Loop

...

Algorithm 1: Unified RSP Processing Loop

Input: Sensory data $y(t)$, current beliefs $x_l(t)$, precisions $W_l(t)$

Output: Updated beliefs $x_l(t+1)$, action $u(t)$

Compute prediction errors:

$e_l \leftarrow \text{input}_l - g_l(x_{l+1})$ for $l = 1, 2, 3$

Compute temporal errors:

$\eta_l \leftarrow x_l(t) - f_l(x_l(t-1), x_{l+1}(t-1))$

Update beliefs (master dynamics):

$\dot{x}_l \leftarrow \kappa_l [\partial g_{l-1} / \partial x_l W_{l-1} e_{l-1} - W_{l+1} e_{l+1} - S_l^{-1} \eta_l]$

Update precisions:

```

 $\dot{W}_l \leftarrow \alpha_l (1/2 W_l^{\{-1\}} - 1/2 \langle e_l e_l^T \rangle_t)$ 

Recursive self-modeling (if  $l = 3$ ):
for  $d = 1$  to  $D_{\max}$ :
 $M^{\wedge}(d) \leftarrow (1 - \beta_d) M^{\wedge}(d-1) + \beta_d \cdot \text{SocialReflect}(M^{\wedge}(d-1))$ 
if  $\|M^{\wedge}(d) - M^{\wedge}(d-1)\| < \delta_{\text{conv}}$ : break

Action selection:
 $u^* \leftarrow \text{argmin}_u E[\sum_{\tau=t}^{\tau=t+T} J(\tau) \mid u]$ 

Return  $x_l(t+1), u(t)$ 
...

```

E.8 Phylogenetic Parameterization

Table E.2: Cross-Species Parameter Ranges

Parameter	Bacteria	Insects	Mammals	Primates	Humans	Interpretation
$D_{\max}$	00	11	22	33	55	Max recursion depth
γ	0.90	0.70	0.90	0.50	0.30	Timescale separation
β	0.10	0.30	0.50	0.50	0.70	Recursion discount
β_d	0.01	0.10	0.001	0.01	0.01	Social mixing weight
α_l	0.01	0.10	0.001	0.01	0.01	Learning rate ratio

The progression shows increasing recursion depth, timescale separation, and social weighting, with decreasing recursion discount (deeper but more heavily discounted recursions) and learning rate ratios (slower self-model updates relative to perception).

E.9 Conclusion

The unified framework reveals that consciousness, as formalized by the RSP model, is the solution to a single optimization problem viewed through four complementary lenses. The master functional $\mathcal{J}$ and its dynamics equations (E.1-E.4) constitute the complete mathematical theory. Each formalism offers unique insights: predictive coding reveals neural implementation, free energy proves optimality, reinforcement learning explains development, and control theory predicts failure modes. The convergence of four independent mathematical traditions on the same three-level recursive architecture provides strong evidence for the RSP model's structural validity.

Appendix F — Equation Reference Sheet

F.0 Notation Guide

State Variables

SymbolDomainDescription $\mathbf{i}(t) \in \mathbb{R}^{n_i}$ Interceptive sensory input (Level 1 observations) $\mathbf{a}(t) \in \mathbb{R}^2$ Core affect: valence and arousal $\mathbf{s}_k(t) \in \mathbb{R}^{n_k}$ Core schema k : body ($k=1$), spatial ($k=2$), attentional ($k=3$), affective-homeostatic ($k=4$), self ($k=5$) $\mathbf{z}(t) \in \mathbb{R}^{n_z}$ Integrated Level 2 representation (all schemas + core affect) $\boldsymbol{\mu}_l(t) \in \mathbb{R}^{n_l}$ Expectations (sufficient statistics) at level l $\mathcal{M}^{(d)}(t) \in \mathbb{R}^{n_m}$ Recursive self-model at depth d (Level 3b) $\hat{\mathcal{M}}^{(d)}(t) \in \mathbb{R}^{n_m}$ Socially-reflected self-model estimate at depth d $\Phi(t) \in \mathbb{R}^{n_e} \times \mathbb{R}^{n_m} \times \mathbb{R}^{n_e}$ Phenomenal state: experience, self-model, predicted experience

Precision and Error

SymbolDomainDescription $\boldsymbol{\varepsilon}_l \in \mathbb{R}^{n_l}$ Prediction error at level l : $\boldsymbol{\varepsilon}_l = \mathbf{y}_l - g_l(\boldsymbol{\mu}_{l+1})$ $\boldsymbol{\xi}_l \in \mathbb{R}^{n_l}$ Precision-weighted prediction error: $\boldsymbol{\xi}_l = \boldsymbol{\Pi}_l \boldsymbol{\varepsilon}_l$ $\boldsymbol{\Pi}_l \in \mathbb{R}^{n_l \times n_l}$ Precision matrix at level l (inverse covariance) $\boldsymbol{\eta}_l \in \mathbb{R}^{n_l}$ Temporal prediction error: $\boldsymbol{\eta}_l = \dot{\boldsymbol{\mu}}_l - D \boldsymbol{\mu}_l$

Functions and Parameters

SymbolDescription $g_l(\cdot)$ Generative function: maps level $l+1$ expectations to level l predictions F, F_l (Total, level-specific) variational free energy F_3^{ec} Strange loop free energy contribution (recursion depth penalty) $\mathcal{G}(\pi)$ Expected free energy for policy π (action selection) $v(t)$ Core affect valence: positive = prediction error decreasing $\leq l$ Learning rate / step size at level l β_d Mixing weight for social reflection at recursion depth d γ Geometric discount for recursion depth penalties

F.1 Hierarchical Predictive Coding (Appendix A)

F.1.1 Three-Level Generative Model

Level 1: $\mathbf{i}(t) = g_1(\boldsymbol{\mu}_2(t)) + \boldsymbol{\Pi}_1^{-1/2} \boldsymbol{\omega}_1$
 Level 2: $\boldsymbol{\mu}_1(t) = g_2(\boldsymbol{\mu}_3(t)) + \boldsymbol{\Pi}_2^{-1/2} \boldsymbol{\omega}_2$
 Level 3: $\boldsymbol{\mu}_2(t) = g_3(\boldsymbol{\mu}_3^{\text{rec}}(t)) + \boldsymbol{\Pi}_3^{-1/2} \boldsymbol{\omega}_3$

F.1.2 Schema Integration

$$\mathbf{z}(t) = \sigma(\mathbf{W}_{int} \cdot \text{concat}(\mathbf{a}(t), \mathbf{s}_1(t), \dots, \mathbf{s}_5(t)) + b_{int})$$

F.1.3 Strange Loop Recursive Update

$$\mathcal{M}^{(d)}(t) = (1 - \beta_d) \cdot \mathcal{M}^{(d-1)}(t) + \beta_d \cdot \hat{\mathcal{M}}^{(d)}(t)$$

F.1.4 Phenomenal State

$$\Phi(t) = \langle \mathbf{e}(t), \mathcal{M}^{(L_{\max})}(t), \hat{\mathbf{e}}(t+1) \rangle$$

F.1.5 Precision-Weighted Message Passing

$$\dot{\mu}_1 = \kappa_1 \left[\frac{\partial g_{l-1}^T}{\partial \mu_1} \Pi_{l-1} \epsilon_{l-1} - \Pi_{l+1} \epsilon_{l+1} - \Sigma_1^{-1} \eta_1 \right]$$

(Bottom-up error from level below drives updates positively; top-down prediction error from level above acts negatively. Boundary: $\epsilon_0 \equiv \epsilon_{\text{sensory}}, \epsilon_4 \equiv 0$.)

F.2 Free Energy Formulation (Appendix B)

F.2.1 Total Variational Free Energy

$$F = \sum_{l=1}^3 \left[\frac{1}{2} \epsilon_l^T \Pi_l \epsilon_l - \frac{1}{2} \ln j \Pi_{lj} \right] + \sum_{l=1}^3 \frac{1}{2} \eta_l^T \Sigma_l^{-1} \eta_l$$

F.2.2 Recursive Self-Model Free Energy

$$F_3^{\text{rec}} = \sum_{d=1}^{D_{\max}} \gamma^d \cdot \frac{1}{2} \|\mathcal{M}^{(d)}(t) - \mathcal{M}^{(d-1)}(t)\|_{\Pi_3^{(d)}}^2$$

F.2.3 Valence as Free Energy Gradient

$$v(t) \approx -\text{sign} \left(\frac{\partial F_1}{\partial \mu_2} \right) \cdot \left\| \frac{\partial F_1}{\partial \mu_2} \right\|$$

F.3 Intrinsically Motivated Reinforcement Learning (Appendix C)

F.3.1 Reward as Interoceptive Prediction Error

$$r(t) = v(t) = \phi_v \cdot \xi_1(t)$$

F.3.2 Action Selection via Expected Free Energy

$$\mathbf{u}^*(t) = \arg \min_{\mathbf{u}_{\text{cand}}} G(\mathbf{u}_{\text{cand}}) = \arg \min_{\mathbf{u}_{\text{cand}}} \hat{F}_1 + \hat{F}_2 + \hat{F}_3 + \hat{F}_3^{\text{rec}}$$

F.3.3 Value Function Update

$$V(\mathbf{z}(t)) \leq V(\mathbf{z}(t)) + \mathbb{E} [r(t) + \gamma_{\text{RL}} V(\mathbf{z}(t+1)) - V(\mathbf{z}(t))]$$

F.3.4 Policy Gradient

$$r = J(\theta) = \mathbb{E}_{\pi} [r - \ln \pi(u|\mu) \Delta A(\mu, u)]$$

F.4 Control Theory Formulation (Appendix D)

F.4.1 PID Control for Hierarchical Error Correction

$$\mathbf{u}_l(t) = K_l^P \boldsymbol{\varepsilon}_l(t) + K_l^I \int_0^t \boldsymbol{\varepsilon}_l(\tau) d\tau + K_l^D \frac{d\boldsymbol{\varepsilon}_l(t)}{dt}$$

F.4.2 Strange Loop Stability Condition

$$\rho = (1 - \beta) + \beta_L < 1$$

F.4.3 Lyapunov Function for Convergence

$$L(t) = \sum_{l=1}^3 \frac{1}{2} \boldsymbol{\varepsilon}_l^T \mathbf{\Pi}_l \boldsymbol{\varepsilon}_l + \sum_{d=1}^{D_{\max}} \gamma^d \|\mathcal{M}^{(d)} - \mathcal{M}^{(d-1)}\|^2$$

F.4.4 Adaptive Precision Update

$$\dot{\mathbf{\Pi}}_l = \alpha_l \left(\frac{1}{2} \mathbf{\Pi}_l^{-1} - \frac{1}{2} \langle \boldsymbol{\varepsilon}_l \boldsymbol{\varepsilon}_l^T \rangle_t \right)$$

(Consistent with Eq. E.3; α_l is the learning rate.)

F.5 Unified Mathematical Framework (Appendix E)

F.5.1 Master Variational Functional

$$\mathcal{J}[\mathbf{x}, \mathbf{W}, \mathbf{u}] = \sum_{l=1}^3 \frac{1}{2} \mathbf{e}_l^T \mathbf{W}_l \mathbf{e}_l + \sum_{l=1}^3 \frac{1}{2} \boldsymbol{\eta}_l^T \mathbf{S}_l^{-1} \boldsymbol{\eta}_l - \frac{1}{2} \sum_{l=1}^3 \ln |\mathbf{W}_l| + \sum_{d=1}^{D_{\max}} \gamma^d \frac{1}{2} \|\mathcal{M}^{(d)} - \mathcal{M}^{(d-1)}\|^2_{\mathbf{W}_3^{(d)}}$$

The single unified objective function from which all four formalisms derive as projections.

F.5.2 Master Dynamics Equation

$$\dot{x}_l = -\kappa_l \frac{\partial J}{\partial x_l} = \kappa_l \left[\frac{\partial g_{l-1}^T}{\partial x_l} \mathbf{W}_{l-1} \mathbf{e}_{l-1} - \mathbf{W}_{l+1} \mathbf{e}_{l+1} - \mathbf{S}_l^{-1} \boldsymbol{\eta}_l \right]$$

A single equation that simultaneously implements message passing (A), free energy gradient descent (B), value maximization (C), and control error correction (D).

F.5.3 Precision Update (Attention)

$$\dot{\mathbf{W}}_l = -\eta_l \frac{\partial J}{\partial \mathbf{W}_l} = \eta_l \left(\frac{1}{2} \mathbf{W}_l^{-1} - \frac{1}{2} \mathbf{h} \mathbf{e}_l \mathbf{e}_l^T \right)$$

Attention allocation as self-tuning precision: at equilibrium, optimal precision equals the inverse empirical error covariance (η_l is the precision learning rate; the instantaneous outer product serves as a stochastic approximation in online implementations).

F.5.4 Free Energy–Value Duality

$$V(\mathbf{z}) \approx -\mathbb{E} \left[\sum_{\tau=0}^1 \gamma_{\text{RL}} F_{1(t+\tau)} \left(\begin{smallmatrix} \circ \\ \circ \\ \circ \\ \circ \end{smallmatrix} \mathbf{z}(t) = \mathbf{z} \right) \right] + \text{const}$$

The RL value function is approximately the negative cumulative interoceptive free energy – bridging reinforcement learning and variational inference.

F.5.5 Universal Strange Loop Convergence

$$\|\mathcal{M}^{(d)} - \mathcal{M}^*_k\| \leq \rho^d_k \|\mathcal{M}^{(0)} - \mathcal{M}^*_k\|, \quad \text{where } \rho = (1 - \beta_0) + \beta_0 \cdot L < 1$$

Geometric convergence of the recursive self-model to a fixed point $\mathcal{M}^*$, proven identically across all four mathematical formalisms.

Appendix H — RSP Theory of Intelligence

H.1 Introduction

The RSP formalization of consciousness as a three-level predictive architecture (Appendixes A-E) contains an implicit account of intelligence. Drawing this out yields three key results: a formal definition grounded in the master functional $\mathcal{J}$, a constitutive identity proposition showing that recursive social intelligence and phenomenal consciousness cannot come apart at Level 3b, and empirical predictions about intelligence-consciousness relationships.

H.2 RSP Intelligence: Formal Definition

H.2.1 Adaptive Capacity as Free Energy Reduction

The master functional from Appendix E provides the foundation:

$$\mathcal{J} = \sum_{l=1}^3 \frac{1}{2} \epsilon_l^T \mathbf{\Pi}_l \epsilon_l + \frac{1}{2} \boldsymbol{\eta}_l^T \boldsymbol{\Sigma}_l^{-1} \boldsymbol{\eta}_l - \frac{1}{2} \ln |\mathbf{\Pi}_l| + \sum_{d=1}^{D_{\max}} \gamma_d \frac{1}{2} \|\mathcal{M}^{(d)} - \mathcal{M}^{(d-1)}\|_{\mathbf{\Pi}_3^{(d)}}^2$$

where $\|\mathbf{A}\|_x = \sqrt{\mathbf{A}^T \mathbf{A}}$ is the Mahalanobis norm.

Definition H.1 (Intelligence as Adaptive Efficiency). The intelligence of an RSP system is the time-averaged rate at which it reduces the master functional per unit of environmental novelty:

$$\mathbf{I} = - \left\langle \frac{d\mathcal{J}}{dt} \middle/ \mathcal{N}(t) \right\rangle_T$$

where $\mathcal{N}(t) = -\ln p(y(t) | \boldsymbol{\mu}(t-1))$ measures environmental novelty, $\langle \cdot \rangle_T$ denotes averaging over time window $\mathcal{T}$, and the definition applies only to episodes of non-trivial novelty ($\mathcal{N}(t) > \mathcal{N}_{\min}$).

This formalizes the idea that intelligence is efficient adaptation to change, not mere accurate prediction in familiar environments. A system encountering high novelty that rapidly reduces prediction error is highly intelligent; one maintaining low $\mathcal{J}$ only by avoiding novelty is not.

H.2.2 Level-Specific Decomposition

Because $\mathcal{J}$ decomposes across levels, intelligence decomposes correspondingly:

$$\mathbf{I}_{total} = \sum_{l=1}^3 \mathbf{I}_l + \mathbf{I}_{rec}$$

where:

$$\mathbf{I}_l = - \left\langle \frac{dF_l}{dt} \middle/ \mathcal{N}_l(t) \right\rangle_T \quad l = 1, 2, 3$$

$$\mathbf{I}_{rec} = - \left\langle \frac{d}{dt} \sum_{d=1}^{D_{max}} \gamma^d \frac{1}{2} \|\mathcal{M}^{(d)} - \mathcal{M}^{(d-1)}\|_{\mathbf{\Pi}_3^{(d)}}^2 \mathcal{N}_{social}(t) \right\rangle_T$$

This decomposition allows for environment-specific intelligence profiles. Physical environments favor $\mathbf{I}_1$ and $\mathbf{I}_2$; social environments favor $\mathbf{I}_3$ and $\mathbf{I}_{rec}$. A corvid may excel at $\mathbf{I}_2$ while humans dominate $\mathbf{I}_{rec}$.

H.2.3 The Role of Precision

Intelligence involves not only reducing prediction error but learning what to attend to. Precision weighting follows:

$$\dot{\mathbf{\Pi}}_l = -\text{ff}_l \frac{\partial \mathcal{J}}{\partial \mathbf{\Pi}_l}$$

At equilibrium, optimal precision equals the inverse of empirical error covariance:

$$\mathbf{\Pi}_l^* = \text{h}\epsilon_l \epsilon_l^T \text{i}_t^{-1}.$$

Proposition H.1 (Precision as Meta-Intelligence). The speed of precision convergence measures meta-intelligence—the capacity to learn what is worth learning:

$$\mathbf{I}_{meta} = \sum_{l=1}^3 \left\langle \frac{\|\mathbf{\Pi}_l(t_0 + \Delta t) - \mathbf{\Pi}_l^*\|}{\|\mathbf{\Pi}_l(t_0) - \mathbf{\Pi}_l^*\| + \epsilon} \right\rangle_{events}^{-1}$$

This connects to the attention schema at Level 2. Attention, in RSP, is precision allocation. Well-calibrated attention schemas identify relevant features quickly, enabling faster learning.

H.3 The Constitutive Identity Proposition

The central claim: recursive social intelligence and phenomenal consciousness are not two properties that co-occur but one property described from two perspectives.

H.3.1 The Convergence Argument

RSP intelligence at its highest level requires efficient convergence of the recursive self-model to its fixed point $\mathcal{M}^*$. From Appendix E, convergence requires:

Contraction mapping Ψ with coefficient $\rho < 1$

Other-models $\mathcal{O}$ of sufficient fidelity

Appropriate precision weighting $\Pi_3^{(d)}$

These are exactly the conditions for phenomenal consciousness at Level 3b. The strange loop whose convergence constitutes recursive intelligence is the same loop whose operation constitutes conscious experience.

Proposition H.2 (Internal Consistency of the Recursive Intelligence Measure with the Architectural Stability Condition). Let $\Phi(t) = \text{he}(t), \mathcal{M}^{(D_{\text{max}})}(t), \hat{\mathbf{e}}(t+1)_i$ be the architectural state object identified, on the constitutive-identity thesis of Part I, with the phenomenal state at Level 3b. For RSP systems encountering non-trivial social novelty:

$$\mathbf{h}_{\text{rec}} \mathbf{i} T \geq 0 \iff \Phi(t) \text{ is well-defined and stable}$$

where "well-defined and stable" means all components of Φ exist, are finite, and $\mathcal{M}^{(D_{\text{max}})}$ has converged within δ_{conv} of its fixed point.

Proof sketch. Forward: $\mathbf{h}_{\text{rec}} \mathbf{i} T \geq 0$ implies non-increasing recursive discrepancy. The contraction condition ($\rho < 1$) ensures convergence, stabilizing $\mathcal{M}^{(D_{\text{max}})}$. Recursion requires $\mathbf{e}(t)$ and $\hat{\mathbf{e}}(t+1)$, making all Φ components well-defined. *Reverse:* stable Φ implies near-fixed-point operation. Contraction ensures self-correction, yielding non-positive time-averaged discrepancy change. *

Note on what the proposition does and does not establish. The biconditional shows that the recursive-intelligence measure I_{rec} is non-negative if and only if the architectural object Φ is well-defined and stable. This is a result about the formalism's internal consistency: the two definitions pick out the same architectural condition. It does not establish, by itself, that Φ is phenomenal experience — that identity is the substantive claim of Part I (constitutive-identity thesis) and is brought to bear here as a premise. Combined with Part I's dissolution argument, the proposition shows that under the constitutive premise, recursive intelligence and Level 3b phenomenal consciousness are the same architectural property described from different formal angles.

H.3.2 Constitutive Identity (Substantive Claim of Part I)

E.3.2 The Master Dynamics

The identity is constitutive, not correlational, on the philosophical thesis defended in Part I. Within the formalism of Part III, what is provable is that the recursive self-modeling architecture, when stable, simultaneously satisfies the conditions for intelligence (non-decreasing I_{reC}) and the conditions identified by Part I's constitutive thesis with phenomenal experience. Minimizing $\mathcal{J}$ simultaneously maximizes adaptive capacity and maintains the recursive structure that, on the constitutive premise, is experience. Under that premise, at Level 3b one cannot occur without the other.

This constrains AI orthogonality theses. For RSP-type systems, goals (Level 1 valence) are architecturally inseparable from the predictive hierarchy constituting intelligence. Remove valence and the hierarchy collapses. Intelligence and goals cannot be arbitrarily orthogonal in this architecture.

H.4 Predictions

Prediction H.1 (Intelligence scales with recursion depth). Intelligence in social domains correlates with maximum recursion depth $D_{m\mathbf{a}_x}$:

Species	$D_{m\mathbf{a}_x}$	Level Assignment	Key Evidence
Humans	3-53b	(full consciousness)	Kinderman et al. (1998); real-time depth likely 2-3
Great apes	1-23a	(proto-consciousness)	Krupenye et al. (2016); contested whether true recursion
Cetaceans	12+		Connor et al. (2006) nested alliances; speculative
Elephants	0-12		Plotnik et al. (2006) MSR; limited ToM evidence
Corvids	0-12		Ecological intelligence; small social groups
Canids	0-12		Cooperative hunting; may use heuristics not recursion
Rodents	01-2		Limited social prediction
Cephalopods	01-2		High non-social intelligence; largely solitary

Prediction H.2 (Precision learning rates). Species with higher $D_{m\mathbf{a}_x}$ should show faster precision adaptation (higher $I_{me\mathbf{t}\mathbf{a}}$) in attention-switching tasks.

Prediction H.3 (Social vs. non-social intelligence dissociation). I_{reC} and $I_{1,2}$ should dissociate across species and individuals. Autism spectrum conditions may show high $I_{1,2}$ with impaired I_{reC} .

Prediction H.4 (Developmental progression). Human $D_{m\mathbf{a}_x}$ should increase from 0 (infancy) to 1 (ages 3-5, first-order ToM consolidation; Wellman, Cross & Watson 2001 meta-analytic range) to 2 (ages 6-7, second-order false-belief consolidation; Perner & Wimmer 1985) to 3-5 (ages 9-12 and beyond; Hollebrandse, van Hout & Hendriks 2014;

Miller 2009). On Part II's framing, full Level 3b (recursive self-modeling under contraction) is properly anchored on the second-order* milestone, not the first-order milestone — first-order ToM is a Level 3a precondition, while second-order recursive ToM is the formal anchor for Level 3b emergence under the contraction. The architectural prediction is therefore that the consolidation of Level 3b occurs at $D_m \mathbf{a}_x = 2$ (ages 6–7), with subsequent depth elaboration through middle childhood.

Prediction H.5 (Neural efficiency). Higher intelligence should correlate with lower neural energy expenditure per unit $d\mathcal{J}/dt$ reduction, measurable via metabolic imaging during novel task performance.

Prediction H.6 (Pathological states). When $\rho \geq 1$ (paranoia, echo chambers), both I_{rec} and phenomenal consciousness stability should degrade simultaneously.

H.5 Limitations and Extensions

The framework has three main limitations. First, I requires full state observability; empirical proxies need development. Second, the level decomposition weights $\rightarrow l$ lack principled constraints. Third, the constitutive identity applies only to Level 3b recursive intelligence, not intelligence generally.

Extensions include: (1) developing observable proxies for $d\mathcal{J}/dt$ using neural efficiency measures, (2) constraining $\rightarrow l$ through ecological analysis of species' environmental demands, and (3) testing the recursion depth predictions through comparative false-belief paradigms adapted for non-human species.

The RSP theory provides a mathematically grounded account linking intelligence, consciousness, and neural architecture through the shared principle of hierarchical prediction under precision-weighted uncertainty. The constitutive identity at Level 3b offers a novel resolution to questions about the relationship between cognitive capacity and phenomenal experience.

Appendix I: Bilevel Optimization and Autonomous Schema Learning

Overview

This appendix formalizes the evolutionary-developmental (Evo/Devo) dynamics that bootstrap the RSP architecture. Drawing on Dupoux, LeCun, and Malik's (2026) bilevel optimization framework, we specify what evolution provides (meta-parameters), what development unfolds (schema learning through self-supervised learning), and how meta-control maps onto the strange loop's switching dynamics.

Scope. This provides a causal-historical account of how the RSP architecture emerged, not how consciousness is produced from non-conscious precursors. Under strong dissolution, the architecture running IS the phenomenology (Part I).

Notation. All symbols from Appendixes A–E are retained. New bilevel symbols: (evolutionary meta-parameters), $L(\cdot)$ (evolutionary loss), $\mathcal{L}_k^{\text{aux}}$ (auxiliary SSL loss for schema k), $\pi_{\mathcal{M}}(\mathbf{a}_m \text{ } \mathbf{s}_m)$ (meta-control policy), $\sigma l(t)$ (critical period gate).

I.1 Bilevel-to-RSP Mapping

Dupoux et al. propose three interacting systems: System A (learning from observation), System B (learning from action), and System M (meta-control that arbitrates between modes). The mapping to RSP is approximate — RSP's Level 3b (recursive self-modeling) substantially extends beyond Dupoux's System M (a simpler routing/scheduling mechanism):

Bilevel Concept	RSP Equivalent	Role	Meta-parameters	Phylogenetic constraints	Receptor densities, connectivity priors, critical period timing
Init()	Neonatal architecture	Level 1	partially operative; Levels 2–3 scaffolded but immature	System A	Levels 1–2 + Level 3a
Perceptual learning via prediction error minimization	System B	Levels 2 + Level 3a	Schema-guided motor learning and exploration	System M	Level 3b + Level 2 (affective)
Meta-control via recursive self-monitoring and valence signaling	Evolutionary loss $\mathcal{L}$	Negative inclusive fitness	Reproductive cost as function of schema quality and social competence	Inner loop	Ontogenetic development
Free energy minimization within one lifetime	Outer loop	Phylogenetic adaptation	Selection pressure on	across generations	

I.2 The Bilevel Framework

I.2.1 Problem Formulation

Outer loop (evolutionary): Optimize meta-parameters across generations:

$$\phi^* = \arg \min_{\phi} \mathbb{E}_{\text{Env} \sim E} [L(A_{0:K}, \mathbf{B}_{0:K} \mid \phi)]$$

Inner loop (developmental): Given fixed ϕ , initialize and develop:

$$\begin{aligned} \mathbf{A}_0, \mathbf{B}_0, \mathcal{M}_0 &= \text{Init}(\phi) \\ \mathbf{A}_{i+1}, \mathbf{B}_{i+1} &= \text{Update}(\mathcal{M}_i, \mathbf{A}_i, \mathbf{B}_i, \text{Env}) \end{aligned}$$

I.2.2 Meta-Parameter Decomposition

$$\phi = [\phi_1, \phi_2, \phi_3, \mathcal{M}]$$

1 (Core Affect): Interoceptive receptor densities, subcortical circuitry (NTS, parabrachial, Vmb, hypothalamus, PAG, central amygdala), innate valence mapping. Partially operative at birth.

2 (Core Schemas): Body schema connectivity priors, hippocampal-entorhinal spatial circuits, attention network scaffolding (collicular + frontoparietal), self-schema scaffolding (anterior insula, mPFC). Scaffolded at birth, develops through experience.

3 (Recursive Prediction): Theory of Mind scaffolding (TPJ, mPFC, precuneus), action-observation network (pSTS, IPL, ventral premotor), communication scaffolding (IFG, posterior temporal). Non-operational until social experience.

4 (Meta-Control): Prefrontal routing, neuromodulatory circuits, critical period timing $\{t_{\text{open},l}\}$, sleep/wake gating.

I.2.3 Critical Period Gates

Developmental updates include plasticity gating:

$$\mu_1^{(i+1)} = \mu_1^{(i)} - \kappa_1 \frac{\partial J}{\partial \mu_1} \bigg|_{\mu_1^{(i)}} \cdot \sigma_1(i \cdot \Delta t, \phi, \mathcal{M})$$

where the critical period gate is:

$$\sigma_l(t, \mathcal{M}) = \sigma_{\text{min},l} + (1 - \sigma_{\text{min},l}) \cdot \frac{1}{1 + \exp(-\alpha_l^{\text{open}}(t - t_{\text{open},l}))} \cdot \frac{1}{1 + \exp(\alpha_l^{\text{close}}(t - t_{\text{close},l}))}$$

I.2.4 Evolutionary Loss Function

$$L(\theta) = \underbrace{\frac{1}{K+1} \sum_{i=0}^K J_{\text{dev}}(i)}_{\text{developmental efficiency}} + \underbrace{\lambda_{\text{social}} \cdot D_{\text{KL}} \left(p(y_{\text{social}} | \mathcal{M}^{(D_{\text{max}})}) \parallel \bar{p}(y_{\text{social}}) \right)}_{\text{social coordination cost}} + \underbrace{\lambda_{\text{repro}} \cdot R(\theta)}_{\text{reproductive cost}}$$

The three terms capture: (1) how efficiently the organism reduces free energy during development, (2) coordination failures from divergent social predictions (frequency-dependent individual selection, not group selection), and (3) metabolic costs of complex architectures.

I.3 Self-Supervised Schema Learning

I.3.1 SSL Framework

Each schema s_k is learned by minimizing an SSL loss:

$$\theta_k^* = \arg \min_{\theta} \mathbb{E}_{(\mathbf{x}_{\text{in}}, \mathbf{x}_{\text{tar}}) \sim \mathcal{G}(\mathbf{y})} [L_k^{\text{aux}}(f_{\theta}(\mathbf{x}_{\text{in}}), \mathbf{x}_{\text{tar}})]$$

I.3.2 Schema-Specific SSL Objectives

SchemaSSL Paradigm Loss $\mathcal{L}_k^{\text{aux}}$ Analogue Body Contrastive InfoNCE on proprioceptive streamsCURL (Laskin et al., 2020) SpatialJoint-EmbeddingV-JEPA prediction in latent spaceV-JEPA (Bardes et al., 2024) Affective-Homeostatic Predictive World Model

$$\frac{1}{2} g(\theta(\mathbf{i}_t, \mathbf{u}_t) - \mathbf{i}_{t+1})^2 + \beta D_{\text{KL}}[q_{k,p}] \text{Dreamer (Hafner et al., 2020) Attention Self-Distillation} \\ - \sum_i \hat{\pi}_i \log p(\text{salient}_i | \mathbf{x}) \text{DINO (Caron et al., 2021) SelfSocial Prediction}$$

$$k \hat{\mathcal{M}}_j(\mathcal{M}) - \mathcal{M}_j^{\text{obs}} k^2 \text{Novel (recursive social embedding)}$$

Proposition I.1 (Schema SSL as Free Energy Decomposition). *For MSE-type SSL losses, under isotropic precision $\mathbf{\Pi}_{2,k} = \approx_k I$ and shared sufficient statistics between SSL encoder and generative model:**

$$\frac{1}{2}\epsilon_{2,\|\Pi_2,\|\epsilon_{2,\|}}^T \approx \frac{\|\pi\|}{2} \cdot L^{\text{aux}}$$

Schema learning IS free energy minimization, decomposed by schema.

I.3.3 Developmental Ordering

Proposition 1.2 (Developmental Ordering). Given timescale separation ≤ 1 , ≤ 2 , ≤ 3 , initial precision decoupling $\Pi_{l+1}^{(0)} \pi 0$, and critical period schedule $t_{\text{open},1} < t_{\text{open},2} < t_{\text{open},3}$, schema k at level l cannot reduce its free energy below threshold until all schemas at level $l - 1$ have stabilized.*

Developmental Predictions:

Schema Predicted Age Required Precursor Core Affect Prenatal–3 months (rudimentary substrate present at birth; full valence-arousal integration develops postnatally) None (innate substrate) Body Schema 0–6 months Stable core affect Spatial Schema 6–12 months Body schema Affective-Homeostatic 6–18 months Core affect + body schema Attention Schema 9–18 months (joint-attention onset; full precision-control function matures with prefrontal development at 3–5 years) Body + spatial + affective schemas Self-Schema (bodily/minimal) 18–24 months All prior Level 2 schemas Level 3a (conceptual ToM, first-order) 3.5–5 years (Wellman et al., 2001) Level 2 self-schema + linguistic conceptual development Level 3b (recursive self-model, second-order) 6–7 years (Perner & Wimmer, 1985), with full recursive depth elaborating through 9–12 years (Hollebrandse et al., 2014; Miller, 2009) Level 3a + executive-function maturation + sustained social experience

I.4 Meta-Control as Routing Policy

I.4.1 Meta-State Vector

$$s_m(t) = \left[\underbrace{v(t), k\epsilon_1(t)}_{\text{Level 1}}, \underbrace{fF_2, \|(t)g\|_2 S, k\epsilon_2(t)}_{\text{Level 2}}, \underbrace{k_e \frac{\text{ctrl}}{3}(t) k, \text{conf}_3(t)}_{\text{Level 3}}, \underbrace{E(t), \text{arousal}(t)}_{\text{somatic}} \right]$$

I.4.2 Routing Policy

The meta-control policy routes between processing modes:

$$\pi_{\mathcal{M}}(\mathbf{a}_m | \mathbf{s}_m) = \text{softmax} \left(\mathbf{W}_{\mathcal{M}}^{(\text{fixed})} \mathbf{s}_m + \mathbf{W}_{\mathcal{M}}^{(\text{rec})} \mathcal{M}(t) + b_{\mathcal{M}} \right)$$

where $\mathbf{W}_{\mathcal{M}}^{(\text{fixed})}$ is hardwired (from $\mathcal{N}$) and $\mathbf{W}_{\mathcal{M}}^{(\text{rec})}$ is learned recursively once Level 3b stabilizes.

Meta-actions $\mathbf{a}_m$:

Homeostatic mode: Level 1 dominates (high Π_1 , low $\Pi_{2,3}$)

Schematic mode: Level 2 dominates (high Π_2 , moderate $\Pi_{1,3}$)

Reflective mode: Level 3 dominates (high Π_3 , low $\Pi_{1,2}$)

Integrated mode: Balanced precision across levels

I.5 Bilevel Algorithm

Algorithm 2: Bilevel Evo/Devo Loop

...

Input: Environment distribution $\mathcal{E}$, population size N , generations T_{gen}

Output: Optimized meta-parameters φ^*

1: Initialize population $\{\varphi^{(i)}\}_{i=1}^N$ randomly

2: for generation $g = 1$ to T_{gen} do

3: for individual $i = 1$ to N do

4: // Inner loop: development

5: $(A_o, B_o, M_o) \leftarrow \text{Init}(\varphi^{(i)})$

6: for developmental step $t = 1$ to K do

7: $\sigma_l(t) \leftarrow \text{CriticalPeriodGate}(t, \varphi_M^{(i)})$ // Eq. I.4

8: for level $l = 1$ to 3 do

9: $\mu_l \leftarrow \mu_l - \kappa_l \nabla_{\mu_l} J \cdot \sigma_l(t)$ // Eq. I.3

10: end for

11: $(A_t, B_t) \leftarrow \text{Update}(M_{t-1}, A_{t-1}, B_{t-1}, \text{Env})$

12: end for

13: // Evaluate evolutionary loss

14: $L^{(i)} \leftarrow \text{ComputeLoss}(A_{0:K}, B_{0:K})$ // Eq. I.5

15: end for

```

16: // Outer loop: selection and variation
17:  $\{\varphi^i\}_{i=1}^N \leftarrow \text{SelectAndMutate}(\{\varphi^i\}_{i=1}^N, L^i)_{i=1}^N$ 
18: end for
19: return  $\varphi^* \leftarrow \arg\min_i L^i$ 
...

```

I.6 Extended Master Functional

The complete RSP objective integrates all levels with meta-control:

$$\mathcal{J}_{\text{extended}} = \sum_{l=1}^3 \alpha_l(t) \leq F_l + \gamma_l \int_0^t k \dot{\mu}_l(\tau) k^2 d\tau + \beta_l \text{DIC}_l \leq + \lambda \mathcal{L}_{\text{meta}}(\pi \mathcal{M})$$

where DIC_l is the Deviance Information Criterion (Spiegelhalter et al., 2002) for level l (penalizing effective model complexity beyond the log-determinant term in $\mathcal{J}$), $\alpha_l(t) / \text{tr}(\mathbf{\Pi}_l(t))$ are the precision-based level weights, and $\mathcal{L}_{\text{meta}}$ is the meta-control loss (policy gradient on long-term free energy reduction).

I.7 Convergence Analysis

Proposition I.3 (Developmental Convergence). *Under the assumptions of Proposition I.2, plus Lipschitz continuity of the update rule and bounded environments, the inner loop converges to a developmental fixed point $(\mathbf{A}^*, \mathbf{B}^*, \mathcal{M}^*)$ for any fixed $\cdot$.

Proof sketch. The critical period gates ensure finite plasticity windows. Within each window, the update rule is a contraction mapping on the bounded parameter space (by Lipschitz assumption). The hierarchical ordering prevents circular dependencies. Convergence follows by the Banach fixed-point theorem applied sequentially to each level. *

Proposition I.4 (Evolutionary Convergence under Quasi-Stationarity). Under standard evolutionary algorithm assumptions (bounded mutation, elitist selection, infinite population limit) **and the assumption that the environment distribution is approximately quasi-stationary on the relevant timescale**, the outer loop converges to a local minimum of the evolutionary loss $\mathcal{L}(\cdot)$. *

Proof sketch. This follows from standard results on evolutionary optimization (Rudolph, 1994). The key requirement is that the developmental inner loop (Proposition I.3) provides a well-

defined fitness landscape under stationary inputs. The bilevel structure ensures this by making each $L(\cdot)$ evaluation deterministic given a fixed environment distribution. The quasi-stationarity assumption is required for the standard convergence proof to apply. *

Substantive limitation: Pleistocene non-stationarity. The Pleistocene period most relevant to hominin Level 3 evolution involved substantial climatic oscillations (glacial–interglacial cycles ~100 kyr periodicity, Heinrich events, Dansgaard–Oeschger oscillations) that violate the quasi-stationarity assumption on the very timescales where the outer loop is supposed to act. The infinite-population limit is also archaeologically implausible for Pleistocene hominins (estimated effective population sizes of 10^3 – 10^4). Proposition I.4 therefore does not establish that hominin Level 3 actually emerged via the convergence the proposition describes. What it establishes is that under quasi-stationarity, the bilevel formalism admits convergence to a local minimum; whether the formalism applies to non-stationary Pleistocene environments is a substantive empirical question outside the scope of the formal result. Treatment of non-stationary cases would require online or stochastic-environment extensions of the bilevel framework (e.g., contextual bandits, lifelong learning) and is left to future work.

Corollary I.1 (RSP Architecture Emergence under Quasi-Stationarity). Under the assumption of approximate quasi-stationarity, the combination of Propositions I.3 and I.4 implies that the bilevel optimization can converge to meta-parameters θ^* , that produce the RSP architecture as a developmental attractor.* This corollary describes formal sufficient conditions; whether the corollary applies to actual hominin evolution depends on whether the Pleistocene environment satisfies the stationarity assumption to a sufficient approximation. The corollary is therefore a constraint on the formalism's domain of applicability rather than an empirical explanation of Level 3 emergence.

I.8 Complexity Thresholds and Species Predictions

The reproductive cost term $R(\cdot)$ in Eq. I.5 creates complexity thresholds. Species evolve higher RSP levels only when the coordination benefits outweigh metabolic costs:

Level 1 only: Most vertebrates (basic affect, no schemas)

Level 1 + partial Level 2: Social mammals (body/spatial schemas, limited self-schema)

Level 3a (conceptual ToM precursors): Great apes (well-supported); elephants and cetaceans (suggestive comparative-cognition evidence); corvids (contested — see Penn & Povinelli 2007; Heyes 2015 for behavior-reading alternatives to mental-state attribution). Full

Level 3b (recursive self-modeling) in non-human species remains an open empirical question; this row should be read as the upper bound of the architectural range these lineages may instantiate, not as a claim that they all instantiate the full hierarchy. Corvid placement is consistent with $D_{m\mathbf{a}_x} = 0-1$ (per the species-prediction table in §H), reflecting Level 2 sophistication with contested Level 3a status, not the "full hierarchy."

The bilevel framework thus explains both the RSP architecture's emergence and its taxonomic distribution, completing the formal account of how evolution bootstrapped the strange loop that constitutes conscious experience.

10. Conclusion: Five Windows on One Architecture

What the Formalization Achieves

This paper has formalized the RSP model of consciousness from five mathematical traditions — predictive coding, variational free energy, intrinsically motivated reinforcement learning, cascade feedback control, and a unified master functional — and established their approximate equivalence under the Laplace approximation. Unlike the exact equivalences between Newtonian, Lagrangian, and Hamiltonian mechanics (related by Legendre transforms), the RSP correspondences are approximate and break down far from equilibrium. The analogy is aspirational in precision but genuine in structure: the convergence of independently developed formalisms on the same three-level dynamics, the same developmental ordering constraints, and the same stability conditions provides formal evidence that the architecture captures genuine computational structure.

Three results stand out:

The contraction mapping theorem (Proposition E.4) establishes that the strange loop converges to a unique fixed point when the Lipschitz constant of social prediction stays below 1. This is a provable mathematical property of the recursive self-modeling map. The constitutive identity at Level 3b — that the converged self-model just is reflexively self-aware phenomenality, on the philosophical thesis defended in Part I — is not itself proven by the contraction theorem; what the theorem provides is the formal stability and uniqueness conditions for the architectural object that, under the constitutive premise, is identified with Level 3b phenomenal consciousness. When the contraction condition fails (through social deprivation, pharmacological disruption, or neurological damage), the formalism predicts specific failure modes that map onto known clinical conditions.

The master functional (Appendix E, Eq. E.1) unifies accuracy, temporal consistency, complexity, and recursive self-consistency into a single variational objective. The four prior

formalisms are projections of this functional, each foregrounding different terms. Emergent properties visible only from the unified perspective include coordinate covariance (the dynamics are description-independent) and precision as universal currency (attention, confidence, and control gain are the same quantity seen from different mathematical windows).

The bilevel developmental algorithm (Appendix I, Algorithm 2) completes the formalization by specifying how the architecture is learned. The outer evolutionary loop tunes meta-parameters; the inner developmental loop minimizes free energy with critical period gates regulating plasticity. The extended master functional (Eq. I.8) integrates level-specific free energies with kinetic penalties on belief change rate, model complexity (DIC), and meta-control loss, explaining why the three-level architecture evolved and why it develops in a specific order.

Computational Implications

The formalizations have direct implications for computational modeling and artificial intelligence:

Simulability. Algorithm 1 (Appendix E) specifies the RSP processing loop at single time steps; Algorithm 2 (Appendix I) specifies the developmental and evolutionary loops. Together they constitute a complete algorithmic specification of the theory. A simulated RSP agent that implements both algorithms should, if the formalism is correct, exhibit the developmental sequence, the convergence properties, and the failure modes predicted by the mathematics.

AI consciousness. The RSP conditions for consciousness are functional, not biological. Any system implementing the three-level architecture with genuine interoceptive monitoring (Level 1), embodied schemas (Level 2), and recursive social prediction converging to a stable self-model (Level 3b) would, under the constitutive premise, be a candidate for consciousness. Current large language models satisfy none of these conditions. But the conditions are specific enough to guide the design of future architectures.

The intelligence-consciousness connection. The formal theory of intelligence (Appendix H) derives a measure of intelligence from the same master functional that, under the constitutive premise of Part I, is identified with the architecture of consciousness. The constitutive identity at Level 3b — recursive social intelligence as phenomenal consciousness — is the substantive claim of the series; what the formalism contributes is the precise specification that allows this identity to be applied to specific systems and yield falsifiable predictions about which systems are candidates for consciousness on the architectural functionalist reading defended in Part I.

Limitations

The formalisms operate under the Laplace approximation (Gaussian noise, MAP estimation), which is standard in the predictive coding literature but introduces approximation error for highly non-Gaussian distributions. The bilevel convergence result (Proposition I.4) requires quasi-stationarity of the environment distribution, which is violated during the Pleistocene period most relevant to Level 3 evolution. The evolutionary fitness function has free weighting parameters that are not independently measurable. These limitations are structural, not oversights; they mark the boundary between what the current formalism can prove and what requires future extension.

The Larger Project

This paper is Part III of a three-part series. The philosophical argument for why the RSP architecture dissolves the hard problem of consciousness is in Part I (Author, 2026a). The empirical evidence and falsifiable predictions are in Part II (Author, 2026b). The complete integrated argument with all appendixes is available as a monograph (Author, 2026d).

The mathematics presented here does not explain consciousness. Under the RSP framework's constitutive identity, there is nothing to explain beyond the organizational form itself. What the mathematics does is specify that form precisely enough to derive predictions, prove stability, establish equivalences, and build implementations. The hard problem of consciousness dissolves not when we find the right bridge law between physics and phenomenology, but when we recognize that the architecture running IS the phenomenology. The five formalisms presented here characterize that architecture from every angle available to mathematical science.

References

[Author]. (2026a). Beyond the Bat — I. The dissolution. Manuscript submitted for publication. (Reference anonymized for double-blind review.)

[Author]. (2026b). Beyond the Bat — II. Evidence and predictions. Manuscript submitted for publication. (Reference anonymized for double-blind review.)

[Author]. (2026d). Beyond the Bat: Dissolving the explanatory gap through meaningful integration. Monograph preprint, anonymous repository. (Reference anonymized for double-blind review; specific repository identifiers withheld to preserve blinding.)

Armstrong, A. G., & Vlasov, Y. (2026). Neural correlates of perceptual decision-making in the primary somatosensory cortex. Proceedings of the National Academy of Sciences, 123(18),

e2514107123. <https://doi.org/10.1073/pnas.2514107123>

Åström, K. J., & Murray, R. M. (2021). *Feedback systems: An introduction for scientists and engineers* (2nd ed.). Princeton University Press.

Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge University Press.

Banach, S. (1922). Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae*, 3, 133-181.

Bardes, A., Garrido, Q., Ponce, J., Chen, X., Rabbat, M., LeCun, Y., Assran, M., & Ballas, N. (2024). V-JEPA: Latent video prediction for visual representation learning. In *Proceedings of the 12th International Conference on Learning Representations (ICLR 2024)*. OpenReview.

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

Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., & Joulin, A. (2021). Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV 2021)* (pp. 9650-9660). IEEE.

Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior*. Cambridge University Press.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181-204.

Connor, R. C., Smolker, R., & Bejder, L. (2006). Synchrony, social behaviour and alliance affiliation in Indian Ocean bottlenose dolphins, *Tursiops aduncus*. *Animal Behaviour*, 72(6), 1371-1378.

Corlett, P. R., & Fletcher, P. C. (2014). Computational psychiatry: A Rosetta Stone linking the brain to mental illness. *The Lancet Psychiatry*, 1(5), 399-402.

Craig, A. D. (2002). How do you feel? Interoception: The sense of the physiological condition of the body. *Nature Reviews Neuroscience*, 3(8), 655-666.

Craig, A. D. (2009). How do you feel — now? The anterior insula and human awareness. *Nature Reviews Neuroscience*, 10(1), 59-70.

Critchley, H. D., Wiens, S., Rotshtein, P., Öhman, A., & Dolan, R. J. (2004). Neural systems supporting interoceptive awareness. *Nature Neuroscience*, 7(2), 189-195.

Damasio, A. R. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. Harcourt Brace.

Da Costa, L., Seth, A. K., Friston, K., Ramstead, M. J. D., & Sandved-Smith, L. (2024). *A Rosetta Stone Hypothesis for Neurophenomenology: Mathematical Predictions from Predictive Processing*. arXiv preprint arXiv:2409.20318. <https://arxiv.org/abs/2409.20318>

Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2), 200-227.

Dupoux, E., LeCun, Y., & Malik, J. (2026). *Why AI systems don't learn and what to do about it: Lessons on autonomous learning from cognitive science*. Preprint. (Specific preprint identifier and DOI pending; cite latest version available on the authors' institutional pages.)

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B*, 360(1456), 815-836.

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

Frith, C. D., & Frith, U. (2006). The neural basis of mentalizing. *Neuron*, 50(4), 531-534.

Hafner, D., Lillicrap, T., Ba, J., & Norouzi, M. (2020). *Dream to control: Learning behaviors by latent imagination*. In *Proceedings of the 8th International Conference on Learning Representations (ICLR 2020)*. OpenReview.

Hollebrandse, B., van Hout, A., & Hendriks, P. (2014). Children's first and second-order false-belief reasoning in a verbal and a low-verbal task. *Synthese*, 191(3), 321-333.

Kinderman, P., Dunbar, R. I. M., & Bentall, R. P. (1998). Theory-of-mind deficits and causal attributions. *British Journal of Psychology*, 89(2), 191-204.

Katlowitz, K. A., Cole, E. R., Mickiewicz, E. A., Shah, S., Franch, M., Adkinson, J. A., Belanger, J. L., Mathura, R. K., Meszéna, D., McGinley, M., Muñoz, W., Banks, G. P., Cash, S. S., Hsu, C.-W., Paulk, A. C., Provenza, N. R., Watrous, A. J., Williams, Z., Goldman, A. M., Krishnan, V., Maheshwari, A., Heilbronner, S. R., Kim, R., Rungratsameetaweemana, N., Hayden, B. Y., & Sheth, S. A. (2026). *Plasticity and language in the anaesthetized human hippocampus*. *Nature*. <https://doi.org/10.1038/s41586-026-10448-0>

Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308), 110-114.

- Laskin, M., Srinivas, A., & Abbeel, P. (2020). *CURL: Contrastive unsupervised representations for reinforcement learning*. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*. PMLR.
- Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15(8), 365-373.
- MacKay, D. J. C. (2003). *Information theory, inference, and learning algorithms*. Cambridge University Press. (See Chapter 27 on Laplace approximation.)
- Matsumoto, M., & Hikosaka, O. (2009). Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature*, 459(7248), 837-841.
- Miller, S. A. (2009). Children's understanding of second-order mental states. *Developmental Review*, 29(4), 249-262.
- Perner, J., & Wimmer, H. (1985). 'John thinks that Mary thinks that...': Attribution of second-order beliefs by 5- to 10-year-old children. *Journal of Experimental Child Psychology*, 39(3), 437-471.
- Plotnik, J. M., de Waal, F. B. M., & Reiss, D. (2006). Self-recognition in an Asian elephant. *Proceedings of the National Academy of Sciences*, 103(45), 17053-17057.
- Powers, W. T. (1973). *Behavior: The control of perception*. Aldine.
- Powers, W. T. (2005). *Behavior: The control of perception (2nd ed.)*. Benchmark Publications.
- Rosenthal, D. M. (2005). *Consciousness and mind*. Oxford University Press.
- Rudolph, G. (1994). Convergence analysis of canonical genetic algorithms. *IEEE Transactions on Neural Networks*, 5(1), 96-101.
- Saxe, R., & Kanwisher, N. (2003). People thinking about thinking people: The role of the temporo-parietal junction in "theory of mind." *NeuroImage*, 19(4), 1835-1842.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80(1), 1-27.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599.
- Schurz, M., Radua, J., Aichhorn, M., Richlan, F., & Perner, J. (2014). Fractionating theory of mind: A meta-analysis of functional brain imaging studies. *Neuroscience and Biobehavioral Reviews*, 42, 9-34.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B*, 64(4), 583-639.

Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., Petrovic, P., Uhlhaas, P., Voss, M., & Corlett, P. R. (2018). The predictive coding account of psychosis. *Biological Psychiatry*, 84(9), 634-643.

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

Tononi, G., & Koch, C. (2015). Consciousness: Here, there and everywhere? *Philosophical Transactions of the Royal Society B*, 370(1668), 20140167.

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3), 655-684.

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, 7(6), 464-476.