AlphaCOGANT: Recursive Corporate Self-Improvement as Active Inference

Rendering the AlphaFund Economic World Model as a GNN generative model via COGANT

Daniel Ari Friedman

Atta Labs

daniel@activeinference.institute

ORCID: 0000-0001-6232-9096

Tucker Cahill Chambers

Atta Labs

June 27, 2026

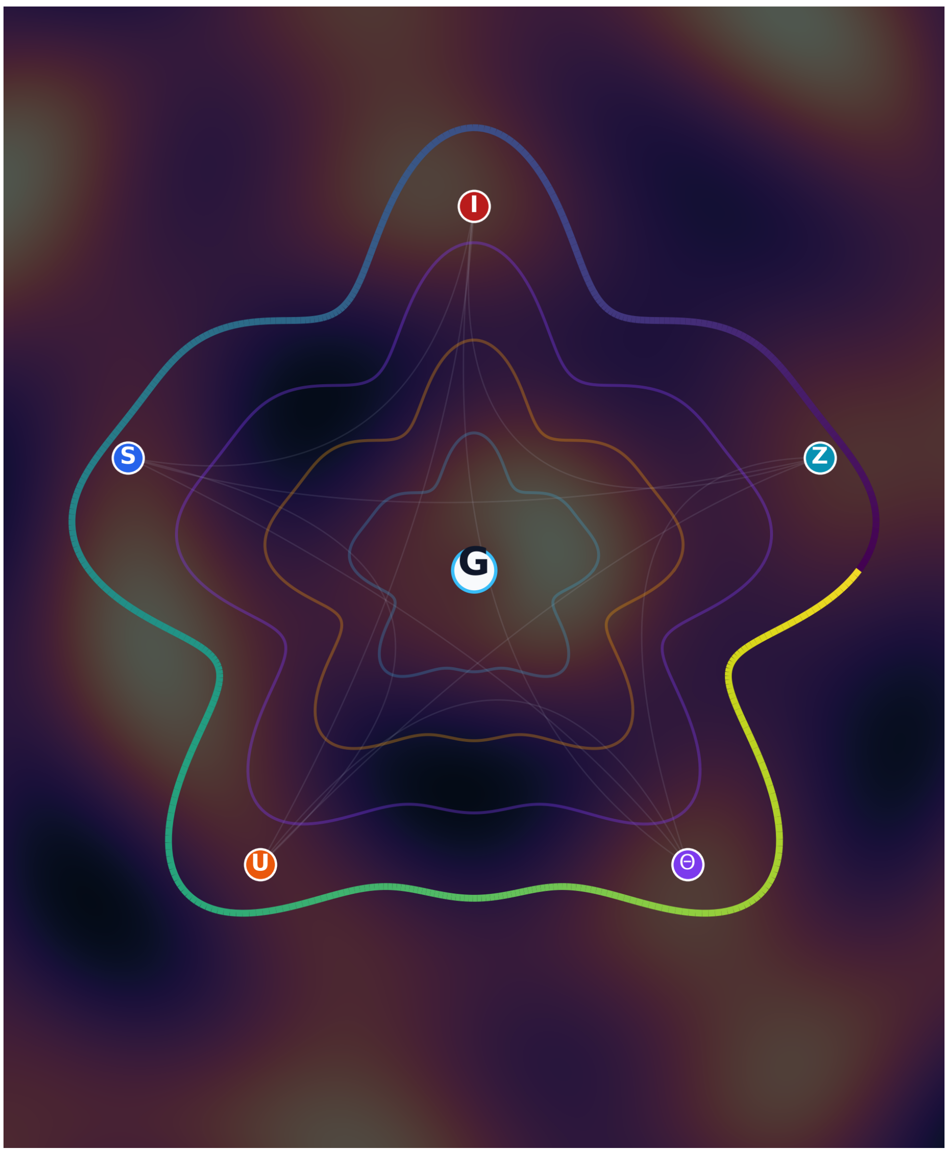

Contents

1 Abstract

The AlphaFund whitepaper reframes recursive self-improvement (RSI) as a portfolio optimization problem: [1] a corporation recursively improves when realized economic gains finance the next cycle of better prediction and deployment, and the firm's standing is summarized by **t-RSI**, a standardized gap between alpha-creation and alpha-decay rates. **AlphaCOGANT** observes that this construction is, term for term, an **Active Inference** agent [1, 2, 3, 4] — and makes the correspondence executable.

We render AlphaFund's **Economic World Model (EWM)** as a generative model written in **Generalized Notation Notation (GNN)**, produced by the **COGANT** [5, 6] codebase-to-GNN translation pattern. The firm's five capital channels — Investments, Sensors, Actuators, Parameters, and R&D — become the hidden-state factors of a partially-observed model; capital allocation becomes the control vector; and the portfolio optimizer's marginal-return objective becomes **Expected Free Energy (EFE)** minimization. The EFE decomposition supplies a principled reading of AlphaFund's own categories: its **pragmatic value** is expected log-equity growth (the alpha-creation rate, read off the broker ledger), and its **epistemic value** is the information gain about the EWM that Sensors and R&D purchase (the data-scaling and forecast-sharpening laws). t-RSI is recovered as the standardized distance between the create-rate and decay-rate posteriors — the thresholded EFE-improvement certificate that admits a self-improvement commit only when creation confidently exceeds decay.

We give the technical and computational realization: a GNN model file for the five-channel firm, a tested NumPy Active Inference engine that performs state inference, computes the epistemic/pragmatic EFE split and the marginal-return vector, and evaluates the t-RSI certificate. We argue that GNN-via-COGANT brings two things AlphaFund's program needs and Active Inference already enforces: **filtration integrity** (the model may condition only on information available at decision time — the same "no-peeking" discipline that separates an EWM from a language model) and **auditable capital allocation** (every admissible funding move has a negative-EFE score under a single, legible objective). This is not financial advice; it is a demonstration that this reduced recursive-corporate-self-improvement model has a direct Active Inference representation supported by source-owning methods and artifact checks [1, 5, 6].

2 Introduction

2.1 Recursive self-improvement is an economic control problem

The literature on recursive self-improvement (RSI) — from Yudkowsky's seed AI [13], through Schmidhuber's Gödel Machines [12], to intelligence-explosion dynamics and diminishing-returns analyses [12, 13] — often studies self-improvement where compute costs are secondary to asymptotics. The AlphaFund whitepaper removes that assumption by placing survival and capital constraints into the optimization problem [1, 24, 25]. Every FLOP and every bit of data costs money; a system that spends more on self-improvement than it earns from the resulting improvement runs out of resources and dies. AlphaFund therefore recasts RSI as constrained control survival test and an auditable certificate [1, 4]. RSI is therefore a **stochastic control problem under a survival constraint** [1, 24, 25], and a corporation is its cleanest instance: a legal object with perpetual succession that turns capital into improved capability and improved capability back into capital [18].

In this light, the economics-of-modern-manufacturing framing is useful, because it treats RSI economics as constrained multi-factor coordination where capital budget structure, not just return maximization, determines what a firm can credibly fund [14].

AlphaFund formalizes the firm as constrained stochastic optimal control over a production cycle. The state is a bundle of five capital **channels** — what the firm holds (Investments), what it can see (Sensors), what it can do (Actuators), what it knows (Parameters), and how it learns (R&D). The objective is the expected discounted log-return on shareholders' equity subject to solvency [9]. The controller is a model-predictive convex program over an **Economic World Model (EWM)** — a learned, filtration-respecting approximation to the true corporate transition law [1, 25]. The firm's standing is summarized by **t-RSI**, the standardized distance between its alpha-creation and alpha-decay rates.

2.2 The thesis of this paper

Read the previous paragraph again with one substitution. A system that maintains a **generative model** of a partially-observed world, infers hidden state from a filtered history of observations, and selects actions that minimize **Expected Free Energy** — trading off the value of preferred outcomes against the value of information — is an **Active Inference** agent [2, 3, 4, 5, 25]. AlphaFund's corporation can be represented in that form: the EWM is its generative model, the five channels are its hidden-state factors, capital allocation is its action, and the marginal-return vector the portfolio optimizer maximizes is the negative-EFE value of each admissible funding move [3, 4]. The two posteriors whose separation defines t-RSI — alpha created per dollar and alpha decayed from the deployed book — are represented here as the **pragmatic** and (the cost side of the) **epistemic** terms of that free energy [2, 4].

AlphaCOGANT makes this correspondence concrete and executable. It does three things:

1. **Renders the firm as a generative model in GNN.** Generalized Notation Notation is a text specification language for Active Inference generative models. We write AlphaFund's five-channel EWM as a GNN model file (`models/alphafund_ewm.md`): channel factors, likelihood and transition matrices, log-preferences, and an Expected-Free-Energy objective annotated with its epistemic and pragmatic parts [5]. The GNN pipeline is explicit about typed parsing, validation, rendering, and executable exports, so the model is readable and checkable before inference [5]. In this representation, the firm's temporal assumptions are explicit graph objects and therefore directly reviewable.
2. **Produces that model by the COGANT pattern.** COGANT is a codebase-to-GNN translator: it scans a system's structure and emits a GNN generative model of it. It also emits a reverse map from generated artifacts, which we use as a provenance check [6]. The differentiable corporation is a system whose every operational degree of freedom is, in AlphaFund's words, "API-complete" — a function call with a structured, causal record. That is precisely the substrate COGANT consumes. We use the COGANT translation step in miniature [6] to map a firm description onto the generative model's priors and maintain a reproducible model provenance chain.
3. **Computes inference, value, and the certificate.** A small, deterministic, fully-tested NumPy engine (`src/alphacogant/`) performs state inference over the channels, computes the Expected-Free-Energy decomposition into epistemic and pragmatic value [3, 4], derives the marginal-return vector, and evaluates the t-RSI improvement certificate.

2.3 Why route AlphaFund through Active Inference and GNN

AlphaFund already has a controller; what does the Active Inference framing add? Two things the whitepaper itself asks for and Active Inference makes explicit in this reduced model.

The first is **integrity of inference**. AlphaFund spends a full section arguing that a language model is *not* an EWM, because held-out validation only enforces **filtration discipline** — conditioning a forecast at time t only on information available at t — when the holdout is strictly after the training corpus [1]. Active Inference is built on this discipline: the generative model factorizes over time and the agent's posterior at t is, by construction, a function of the history filtration $\mathcal{F}_t = \sigma(H_t)$ and nothing resolved later [2, 3, 4, 24, 25]. This is where COGANT and the GNN graph contract are doing most of the heavy lifting [5, 6]. GNN makes this factorization explicit and checkable [5]. The "no-peeking" property AlphaFund must bolt onto a wrapped LLM is the native semantics of the object COGANT emits [6, 25]. This is the same hard constraint that separates time-aware economic forecasting from post-fitted narrative claims [1, 2, 3, 4, 24].

The second is a **single legible objective**. AlphaFund's central move is to put a researcher hire, a data feed, a GPU, and a position in AAPL on one dollar axis by differentiating a common objective. Expected Free Energy *is* that common axis, and its decomposition tells you *why* a channel is worth funding: because it pays in preferred outcomes (pragmatic) or because it sharpens the model that prices all future outcomes (epistemic) [3, 4]. The chronic difficulty of comparing an explore dollar to an exploit dollar dissolves into the same quantity Active Inference has formalized for the past decade as a single value decomposition [2, 3, 4, 23].

AlphaFund's self-forecasting loop — predict (query the EWM for the marginal-return vector), optimize (the convex inner program), execute the funded action, and fold the realized outcome back in as the next training row — is, in Active Inference terms, the perception–action cycle that minimizes Expected Free Energy over the 12-cycle horizon [24]. fig. 1 renders that loop with each stage labelled by its engine symbol, so the reader can see the whole correspondence before the formalism arrives [2, 3, 4, 25].

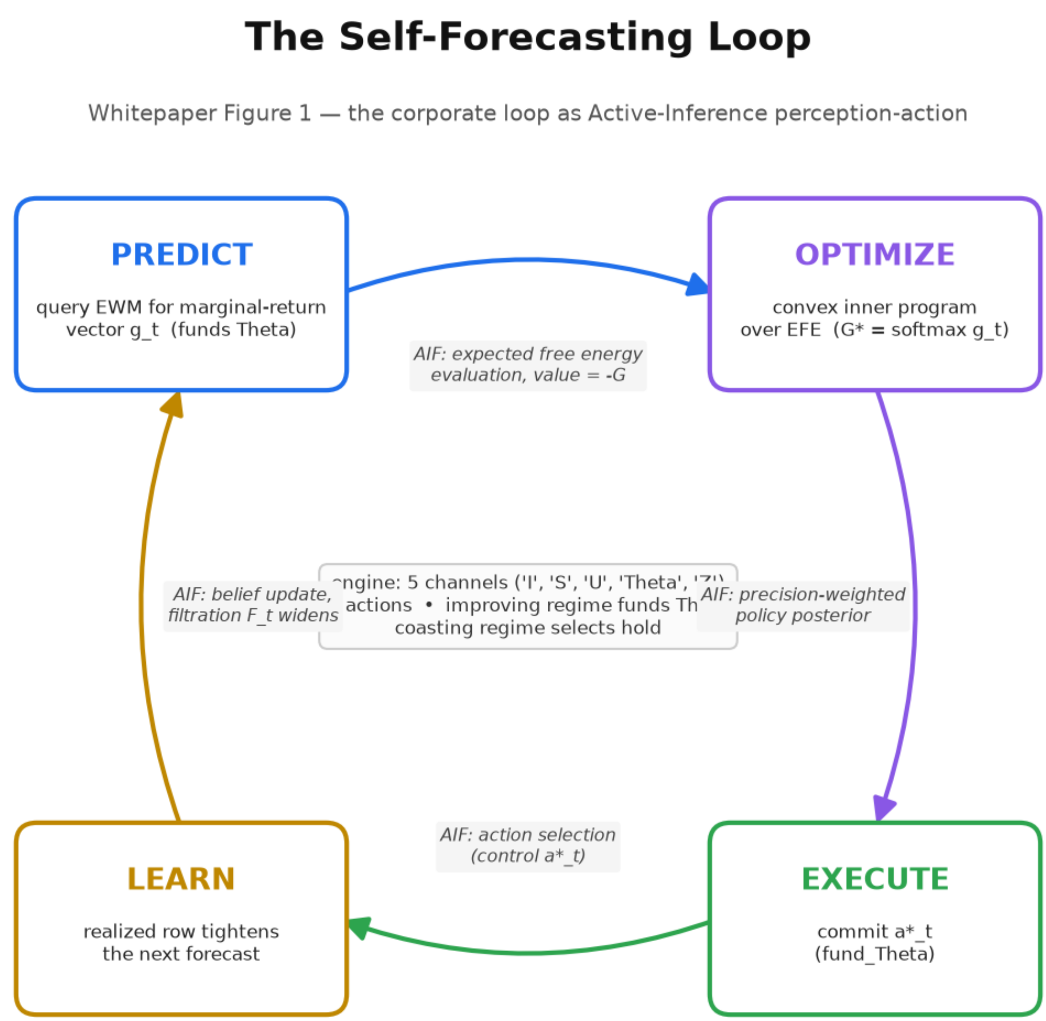

Figure 1: The self-forecasting loop as the Active Inference perception–action cycle: predict via `free_energy.marginal_return_vector`, select via `free_energy.policy_posterior`, execute the funded action over 12 cycles, and fold the realized reward/loss back through `generative_model.infer_states`.

The remainder of the paper builds the dictionary (sec. 3), gives the GNN-via-COGANT realization (sec. 4), works through generative-model inference under the firm filtration (sec. 5), derives the epistemic/pragmatic value split and recovers t-RSI as the EFE certificate (sec. 6), and argues the integrity case (sec. 7).

3 The AlphaFund $\leftrightarrow$ Active Inference dictionary

Active Inference casts any adaptive system as an agent that holds a generative model of how hidden causes produce sensory data, infers those causes by minimizing variational free energy, and selects actions that minimize *expected* free energy over a planning horizon [2, 3, 4]. Below, each AlphaFund construct is matched to its Active Inference counterpart. The match is structural, not metaphorical: the equations coincide.

AlphaFund construct	Active Inference object	Correspondence
Corporation tuple $\Xi_t = (I, S, U, \Theta, Z)$	Factorized hidden state s_t	Five capital channels = five hidden-state factors
Environment E_t (prices, flow, macro)	External hidden causes	The part of the world the firm conditions on but does not fully control
Economic World Model $\widehat{W}_t$	Generative model $P(o, s' \mid s, a)$	Learned, filtration-respecting next-cycle law
Firm history H_t; channel histories H_t^k	Observation sequence $o_{0:t}$	The evidence inference conditions on
Firm filtration $\mathcal{F}_t = \sigma(H_t)$	Belief-update information set	"No-peeking" — posterior at t depends only on $\mathcal{F}_t$
Action vector a_t (dollars per channel)	Control state / policy π	Capital allocation = the agent's action
Cumulative objective J_t (expected $\sum$ log-equity)	Negative Expected Free Energy (pragmatic part)	Discounted log-return = preference satisfaction
Marginal-return vector $g_t = \partial J_t / \partial a_t$	Negative-EFE action value	Continuous marginal return becomes a discrete funding-move score
Equimarginal identity $\hat{g}_t^k / \sigma_t^k = \lambda_{S,t}^*$	Precision-weighted policy optimum	Risk-adjusted shadow price of capital
Sensors / R&D returns (data-scaling, search laws)	**Epistemic value** (information gain)	What reduces EWM predictive loss
Investments / Actuators returns (broker ledger, Sharpe)	**Pragmatic value** (expected utility)	What realizes preferred outcomes now
t-RSI = std. gap(create, decay)	Thresholded EFE-improvement statistic	Certificate gating each commit
Certificate of monotone improvement	Admissibility gate on a model update	Admit update iff value clears a margin
Drift detection + refit	Precision / model revision under surprise	Refit when observations leave support

fig. 2 lays the same correspondence out as a two-column dictionary panel, pairing each AlphaFund construct with its Active Inference object across the 5 channels and 6 actions of the engine — the equations coincide, so the panel is a map of where to find each AlphaFund symbol inside `src/alphacogant`.

3.1 The firm as a generative-model-carrying agent

The corporation "sees the world only through its sensors," so both the environment E_t and the firm's own state Ξ_t enter the EWM as **posteriors over noisy observations**, never as latent ground truth. This is the defining posture of a partially-observed Active Inference agent: there is a hidden state, a likelihood mapping that state to observations, and a transition mapping that state forward under action. AlphaFund's channel histories H_t^k — the $(o_\tau^k, a_\tau^k, R_{\tau+1})$ rows the firm fits its row-laws on — are exactly the per-factor evidence streams an Active Inference agent accumulates.

The EWM's structural promise is that it is the **only** model of the future the firm has access to, so improving it is itself an allocation with first-order effect on the objective. In Active Inference terms: the generative model is parameterized (by Θ), those parameters are themselves hidden states subject to inference and to control, and the value of acting to improve them is quantified by the same expected-free-energy functional that values every other action. AlphaFund's insistence that "every input that lowers predictive loss is priced in dollars on the firm's books" is the economic image of Active Inference's unification of perception, learning, and action under one objective [1].

3.2 Channel-specific world models = factorized generative models

AlphaFund decomposes the joint EWM into **channel-specific world models** [4] $\widehat{W}_t^k$, each trained on its own channel history — a scaling law, a market-impact curve, a refit-decay model, a search law. It is explicit that this is a practical approximation: "cross-channel coupling re-enters when the controller composes the rows." This is precisely the **mean-field / structured factorization** of a generative model in Active Inference [20]: the joint is approximated as a product of per-factor distributions for tractable inference, and coupling re-enters at the policy-evaluation step where Expected Free Energy is computed over the joint predicted outcome [3, 4]. AlphaFund's supermodular cross-partials — "a marginal dollar on channel j raises the marginal value of a dollar on channel k" — are the coupling terms that a fully factorized model drops and the EFE computation restores [14].

The GNN model file in sec. 4 encodes exactly this: five factor blocks with per-factor transition matrices B_k, two likelihood matrices A_R, A_L that couple factors at the observation, and an Expected-Free-Energy block whose epistemic and pragmatic parts compose the rows back together.

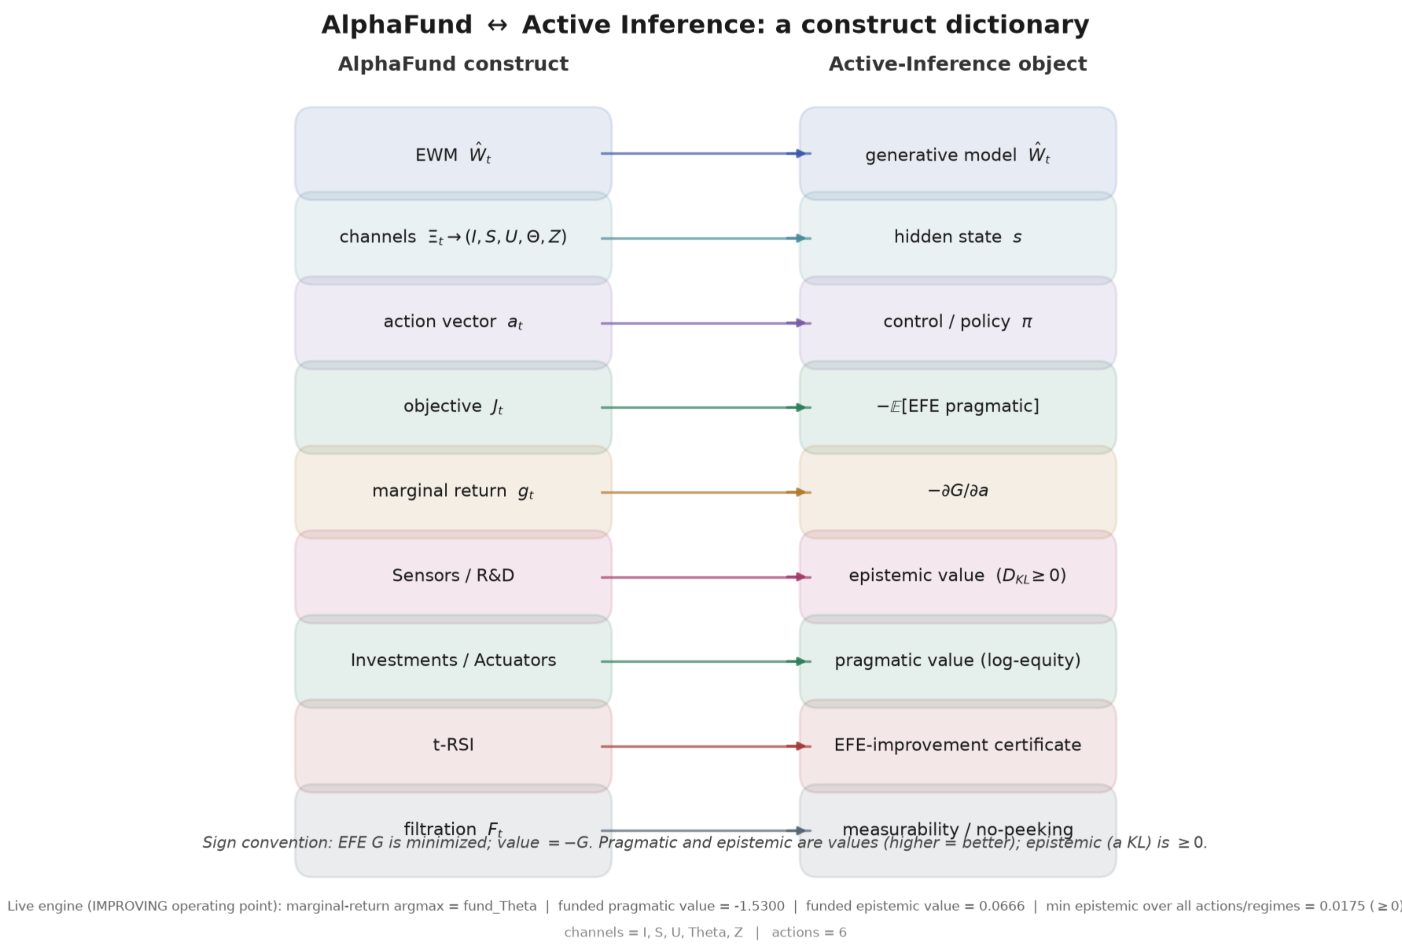

Figure 2: The AlphaFund ↔ Active Inference dictionary: each whitepaper construct (equity, reward, EWM, filtration, marginal-return vector, certificate) paired with its Active Inference object and the realizing engine symbol across 5 channels and 6 actions, with each pairing's value computed by `free_energy.marginal_return_vector`.

4 Technical and computational realization: GNN via COGANT

4.1 What GNN is, and why it fits

Generalized Notation Notation (GNN) is a text-based specification language for Active Inference generative models, with a processing pipeline that transforms a GNN `.md` file into executable simulations (PyMDP, RxInfer.jl, JAX, Active Inference.jl, and others), visualizations, type-checked validations, and reports. This pipeline and model contract are the source of our representation [5]. A GNN model is a Markdown document with structured sections: a `StateSpaceBlock` declaring tensors (the A likelihood, B transitions, C preferences, D priors, hidden states s, observations o, policy π, Expected Free Energy G), a `Connections` block giving the factor-graph edges, an `InitialParameterization` with concrete numbers, an `Equations` block, and an `ActInfOntologyAnnotation` binding every symbol to its Active Inference role.

GNN is the right target for AlphaFund because it is **executable, typed, and filtration-explicit**. The EWM is not a diagram; it is an object the firm rolls forward under candidate actions. GNN compiles to exactly such rollouts, and its structured sections make tensor shapes, probability objects, and graph connections explicit before any policy is evaluated. That static discipline is what AlphaFund needs for an auditable controller: a reviewer can inspect the state factors, likelihoods, transitions, and action edges rather than trusting a free-form narrative.

4.2 What COGANT is, and the translation step

COGANT is a **codebase-to-GNN translator** [6]: it scans a system's structure (program graph, modules, call edges), builds a state-space factor graph, and exports a GNN generative model of the system, which it then renders, visualizes, and validates against the GNN package. The conceptual claim of AlphaCOGANT is that the differentiable corporation is a *natural* COGANT input, because AlphaFund's **third structural fact** is that the firm is "API-complete": data ingestion, model training, capital allocation, trade execution, and asset acquisition are all function calls, each producing a structured record that "doubles as the causal record needed for a derivative against it." A system whose operational degrees of freedom are all API calls with causal records is a system whose structure COGANT can scan and whose dynamics COGANT can express as B_k transitions.

The translation has three stages, mirrored by the engine in `src/alphacogant/`:

1. **Structure → channels.** A firm description — counts of data feeds, execution venues, deployed models, researchers, and book size — maps onto the five channel factors and their prior beliefs D_k. More feeds strengthen the Sensors prior; a larger validated book strengthens Investments; a fresher model strengthens Parameters. This is `cogant_bridge.firm_structure_to_channels`. It is the COGANT scan in miniature [6]: system structure becomes generative-model priors.
2. **Channels → GNN model.** The five factors, their likelihood matrices (A_R for the broker-ledger reward readout, A_L for the predictive-loss readout), their per-action transition matrices B_k, the log-preferences C_R, C_L, and the Expected-Free-Energy block are emitted as a GNN file (`models/alphafund_ewm.md`). `cogant_bridge.model_to_gnn_summary` performs the round-trip check: an `EconomicWorldModel` re-emits a GNN-style ontology block, proving structure was preserved.
3. **GNN model → rollouts and value.** The GNN pipeline (or, here, the NumPy engine that implements the same semantics) performs inference and policy evaluation: posteriors over the channels, Expected Free Energy per policy split into epistemic and pragmatic value, the marginal-return vector, and the t-RSI certificate.

4.3 The model file

`models/alphafund_ewm.md` encodes the five whitepaper channels as hidden-state factors `sI, sS, sU, sTheta, sZ`, each with two capability levels $\{\text{weak}, \text{strong}\}$. The control vector has six actions — `fund_I, fund_S, fund_U, fund_Theta, fund_Z, hold` — so a cycle's decision is "which channel receives the marginal dollar." Two observations are modeled: the realized log-equity reward `o_R` (the only channel read directly off the broker ledger, per AlphaFund's Investments row) and the EWM predictive loss `o_L` (the forecast-evaluation panel). The likelihood A_R makes high reward probable only when the production channels I, U are strong **and** the parameters Θ are fresh; A_L makes low predictive loss probable only when Sensors are rich **and** Θ is fresh. The transition B_Θ encodes AlphaFund's central empirical finding — that deployed parameters **decay** (alpha-decay) toward stale under any non-refit action — and refreshes only when Θ is funded.

fig. 3 draws the GNN factor graph this file declares: the 5 hidden-state factors `sI, sS, sU, sTheta, sZ`, the two likelihood matrices A_R, A_L wiring factors to the reward and loss observations, the per-factor transitions B_k under the 6 actions, and the Expected-Free-Energy block — the `Connections` block made visible, so a future observation conditioning a past belief is an edge one can see is absent.

The reduced two-level encoding keeps the GNN file legible and type-checkable; the continuous marginal-return formalism of sec. 6 is what the engine and the whitepaper uses in production [1]. The point of the file is not numerical fidelity to AlphaFund's proprietary surfaces (which are not public); it is to demonstrate that the firm's control problem has a well-formed reduced Active Inference representation with an Expected-Free-Energy objective, expressible in a language intended to compile to executable inference.

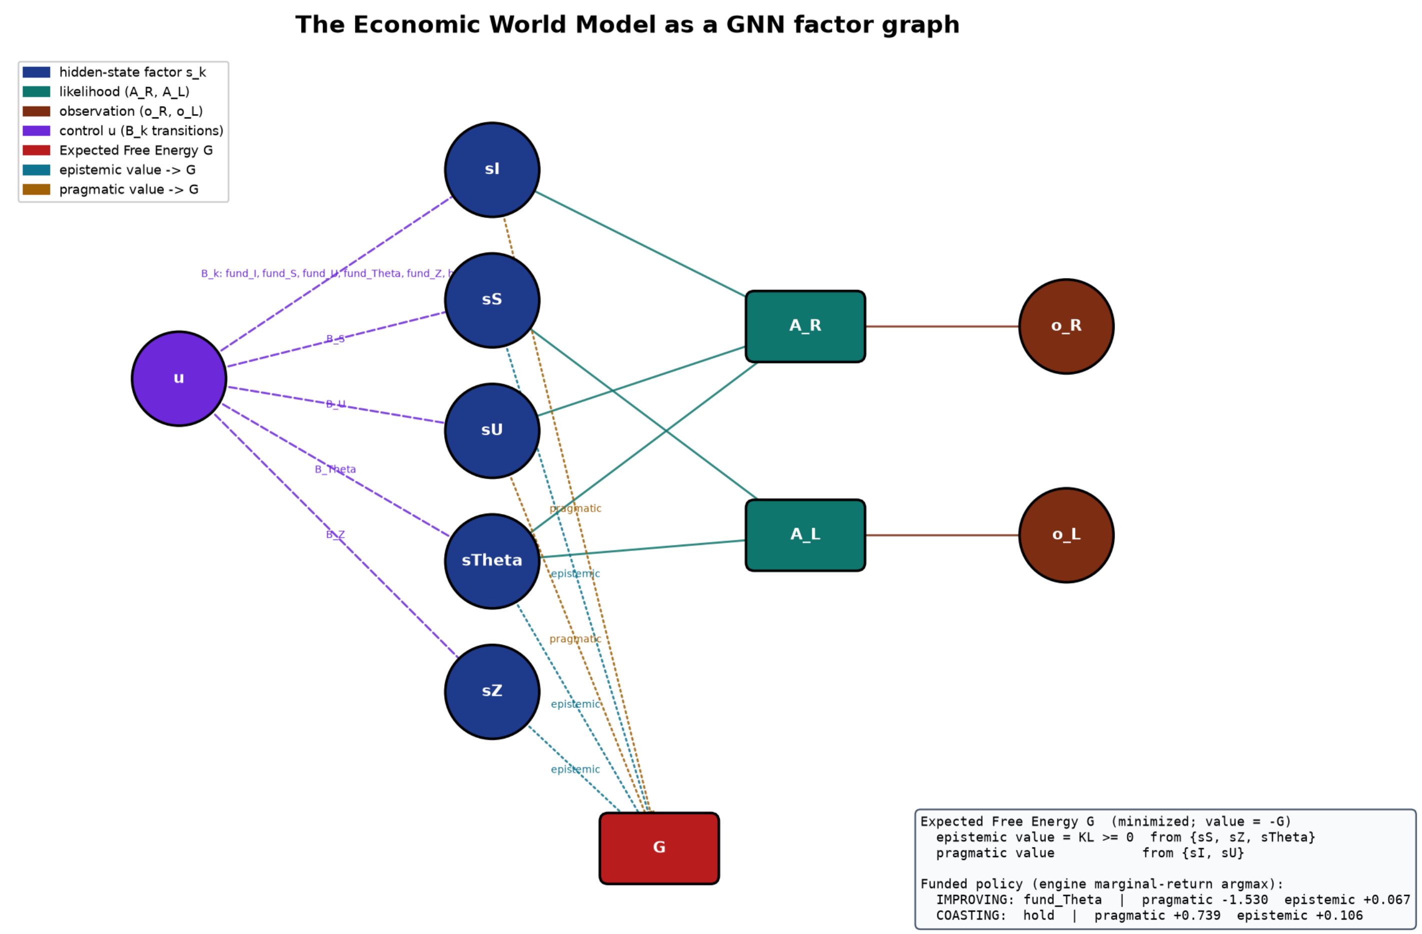

Figure 3: GNN factor graph of the five-channel firm: 5 hidden-state factors, likelihoods A_R (coupling I, U, Θ to reward o_R) and A_L (coupling S, Θ to loss o_L), per-factor transitions B_k over 6 actions, and the EFE objective block, with factor structure read from `generative_model.default_model`.

5 Generative-model inference under the firm filtration

5.1 The generative model

The AlphaCOGANT generative model factorizes the next cycle as

$$P(o_R, o_L, s_{t+1} \mid s_t, a_t) = A_R(o_R \mid s^I, s^U, s^\Theta)\ A_L(o_L \mid s^S, s^\Theta) \prod_{k\in\{I,S,U,\Theta,Z\}} B_k(s^k_{t+1} \mid s^k_t, a_t)\,. \tag{1}$$

This is AlphaFund's "true corporate transition" W — a property of the world the firm cannot access — approximated by the firm's learned $\widehat{W}_t$, the EWM. In this form it is a finite-state, controlled, partially-observed model in the same sense as active-inference POMDPs [2, 3, 24, 25]. The likelihoods A_R, A_L are how the latent corporate state generates the observable broker ledger and forecast-evaluation panel; the transitions B_k are how a capital-allocation action moves each channel. Inference is the inverse problem: recover a posterior $q(s_t)$ over the five channels from the observed reward/loss history.

5.2 Inference is the firm reading its own state

Because the firm "sees the world only through its sensors," it never observes Ξ_t directly; it observes o_R, o_L and infers the channel capabilities that best explain them. In the engine, `generative_model.infer_states` performs one Bayesian filtering update of the per-channel posterior from a bucketed reward and loss observation. For this factor graph the mean-field update is exact per factor given the others, so the update is cheap and the posterior is a product of per-channel beliefs — the computational image of AlphaFund's per-channel row-laws. In this engine instance, exactness follows from the graph structure itself, so the tractability is explicit, not an extra approximation [20, 24, 25].

A worked reading: a high-reward, low-loss observation is most probable when the production channels and the EWM are strong and fresh, so the posterior shifts mass toward `strong` on I, U, Θ and toward `rich` on S. The firm has *inferred* that it is in a compounding regime — not because it was told, but because the ledger and the panel are jointly improbable under any weaker state. This is the self-forecasting loop's "predict" step: query the world model for the current state and its forward law before choosing an allocation.

The single most consequential transition is B_Θ, which encodes AlphaFund's central empirical finding that deployed parameters **decay** toward stale under any non-refit action and refresh only when Θ is funded. fig. 4 traces the belief in fresh Θ forward under the two policies: under `hold` the freshness mass bleeds away each cycle through the decay-leak probability 0.4000, while under `fund_Theta` it is pulled back toward fresh — the mechanism that makes refit a first-order allocation rather than maintenance overhead [1].

5.3 Filtration discipline is native, not bolted on

AlphaFund's sharpest methodological claim is that an LLM is **not** an EWM, because a model trained on a static corpus can "mix documents from before and after the event it is asked to predict," so its context can contain future information. Held-out validation only enforces the no-peeking property when the holdout is strictly chronologically after the entire training corpus — a window that shrinks as internet-scale models train on ever-more-recent data, leaving little data for validation and high measured uncertainty.

In the Active Inference / GNN construction this failure mode becomes an explicit graph contract. The firm filtration $\mathcal{F}_t = \sigma(H_t)$ is the σ-algebra generated by the firm history, and a random variable is $\mathcal{F}_t$-measurable iff its value is determined by H_t. The generative model is defined as a forward factorization in time; the posterior $q(s_t)$ is a function of $o_{0:t}$ and nothing resolved later; the transition B_k maps $t \to t+1$ and is not inverted to leak future observations into a past belief. The two loss functions AlphaFund contrasts —

$$\mathcal{L}_{\mathrm{LLM}}(\Theta) = \sum_i \ell(\hat{p}_\Theta(x_i \mid \mathrm{ctx}_i), x_i) \quad \text{(permutation-invariant over documents)}, \tag{2}$$

$$\mathcal{L}_{\mathrm{EWM}}(\Theta) = \sum_{\tau\in I_{\mathrm{eval}}} \ell(\widehat{P}_\tau(o_{\tau+1}, R_{\tau+1} \mid \mathcal{F}_\tau, a_\tau), (o_{\tau+1}, R_{\tau+1})) \quad \text{(information order matters)} \tag{3}$$

— are, respectively, a likelihood with no temporal index and the **predictive** likelihood that a GNN/Active Inference model is built to optimize [1, 2, 24, 25]. In this reduced project, the factor-graph `Connections` block makes the second form inspectable: there is no declared edge that conditions a time-t belief on a time-t+1 observation. AlphaFund spends a section arguing for a discipline; AlphaCOGANT shows the discipline as a source-owned graph and engine contract.

A note AlphaFund itself makes survives the translation: a general LLM *wrapped* in a strict post-cutoff evaluation harness can still serve as a component of the generative model (e.g. a likelihood or proposal). The categorical claim is about the bare model used as the primary EWM, not about a properly-filtered subroutine — and in GNN such a subroutine is just another typed block whose connections are checked.

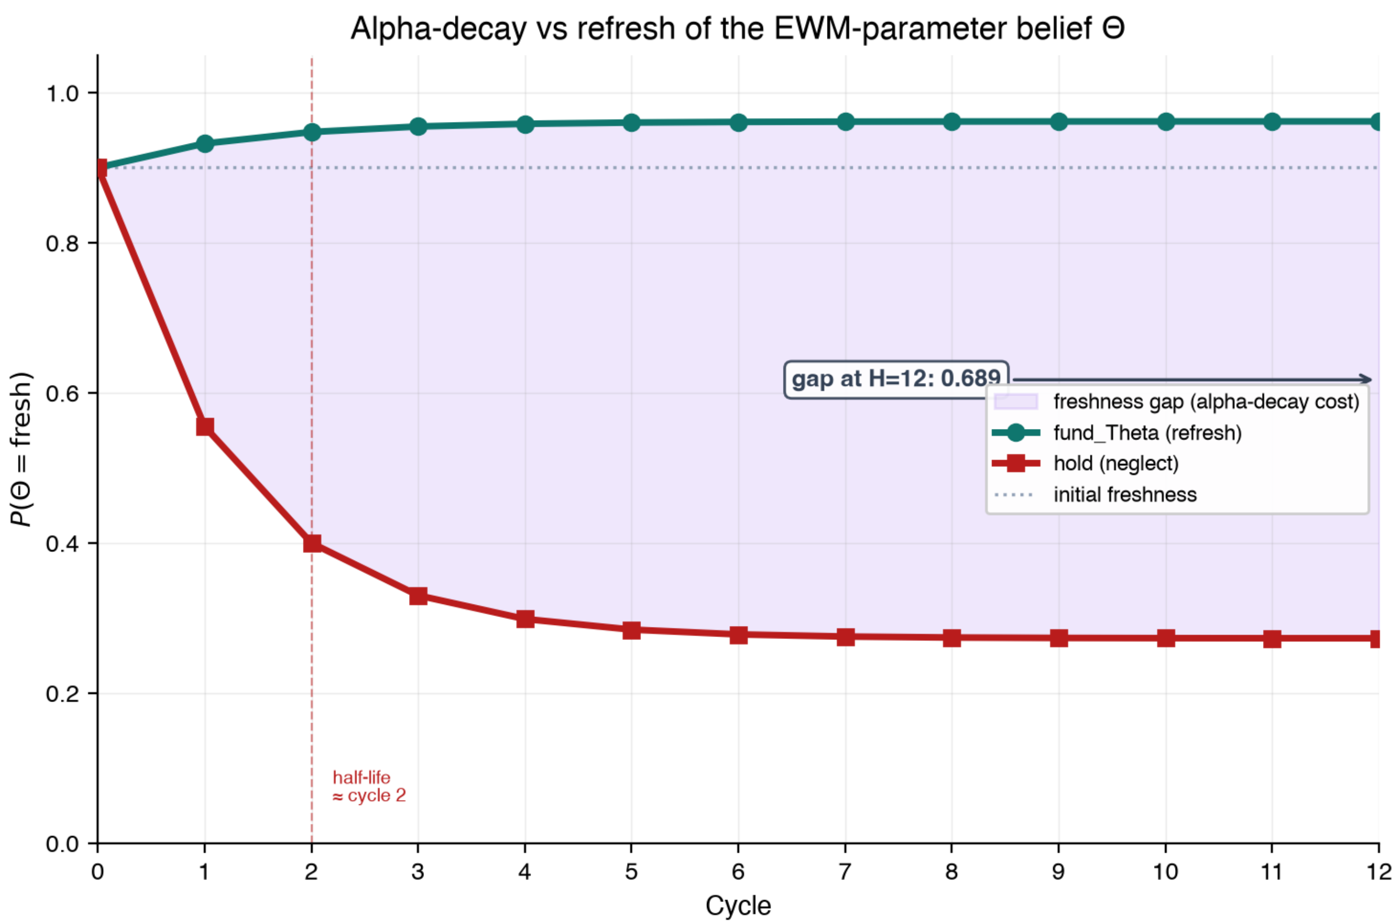

Figure 4: The B_Θ alpha-decay law: belief in fresh parameters traced over 12 cycles under `hold` (freshness decays via leak probability 0.4000) versus `fund_Theta` (freshness restored), with the transition read from `generative_model.default_model`.

6 Epistemic and pragmatic value, and t-RSI as the EFE certificate

6.1 Expected Free Energy is the marginal-return objective

At each cycle the controller scores a candidate policy π (a capital allocation) by its **Expected Free Energy**, which in Active Inference [3, 25] decomposes into two terms with opposite signs of intent:

$$G(\pi) = \underbrace{-\ \mathbb{E}_q\big[\ln P(o \mid C)\big]}_{\text{pragmatic cost}} - \underbrace{\mathbb{E}_q\big[D_{\mathrm{KL}}[\,q(s' \mid o, \pi) \,\|\, q(s' \mid \pi)\,]\big]}_{\text{epistemic value}}. \tag{4}$$

The controller **minimizes** G; equivalently it maximizes value $= -G$. The two parts answer two different questions about a marginal dollar:

- **Pragmatic value** — how much the action is expected to move outcomes toward preference. Here preference C is high realized log-equity reward and low predictive loss, so pragmatic value is **expected log-equity growth**: AlphaFund's cumulative objective J_t, and the create-side of its marginal-return vector. Investments and Actuators (I, U) are the pragmatic channels — they realize edge now, read off the broker ledger.
- **Epistemic value** — how much the action is expected to *reduce uncertainty about the hidden state*, in particular the EWM parameters Θ. This is the information gain a dollar buys, and it is exactly what AlphaFund's Sensors and R&D rows price: the data-scaling law (loss removed per decade of effective tokens) and the experiment-performance frontier (Sharpe gained per decade of experiments). Sensors and R&D (S, Z) are the epistemic channels — they sharpen the model that prices all future outcomes [3, 25].

In AlphaFund's continuous formalism [1] the marginal-return vector is $g_t = \partial J_t / \partial a_t$ and the optimum equates the same risk-adjusted shadow price of capital — AlphaFund's equimarginal identity $\hat{g}_t^k / \sigma_t^k = \lambda_{S,t}^*$. In AlphaCOGANT's reduced discrete action space, `free_energy.marginal_return_vector` computes the corresponding negative-EFE value for each of the 6 admissible funding moves; its argmax is the channel the portfolio optimizer funds this cycle. fig. 5 shows the full vector as a heatmap over all cycles of the greedy trajectory, making visible how the value landscape shifts as beliefs move from weak to strong.

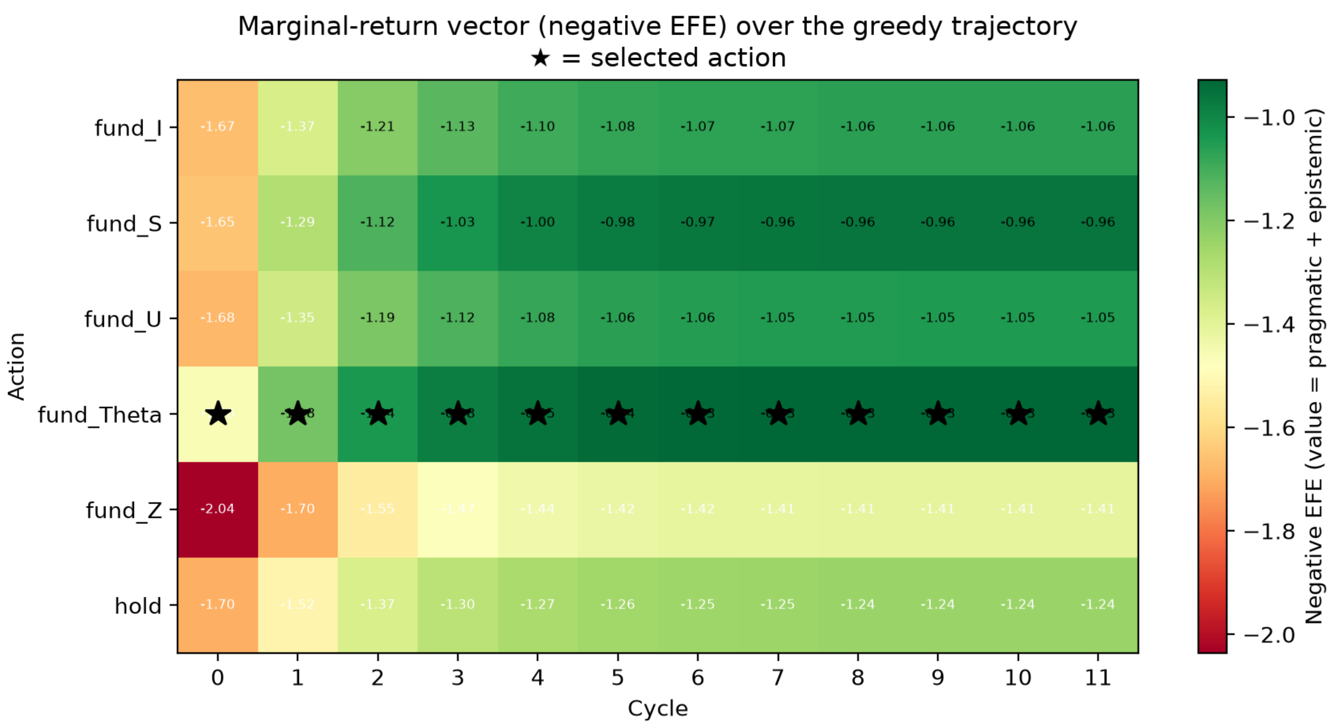

Figure 5: Marginal-return vector (negative EFE = pragmatic + epistemic) across all six actions and all cycles of the greedy trajectory from the self-improving point. The starred action is the greedy selection. Computed by `simulation.simulate_trajectory` and `free_energy.marginal_return_vector`.

fig. 6 decomposes the EFE itself — the objective the controller minimizes — into its pragmatic and epistemic components for all six actions at the IMPROVING operating point. The greedy action (lowest G, highest value) is highlighted; the diamond markers show total G = -(pragmatic + epistemic) for each action. The waterfall makes visible why the greedy policy funds Θ when it is stale: the epistemic value of that funding is small in absolute terms, but the pragmatic cost is the least negative of any action.

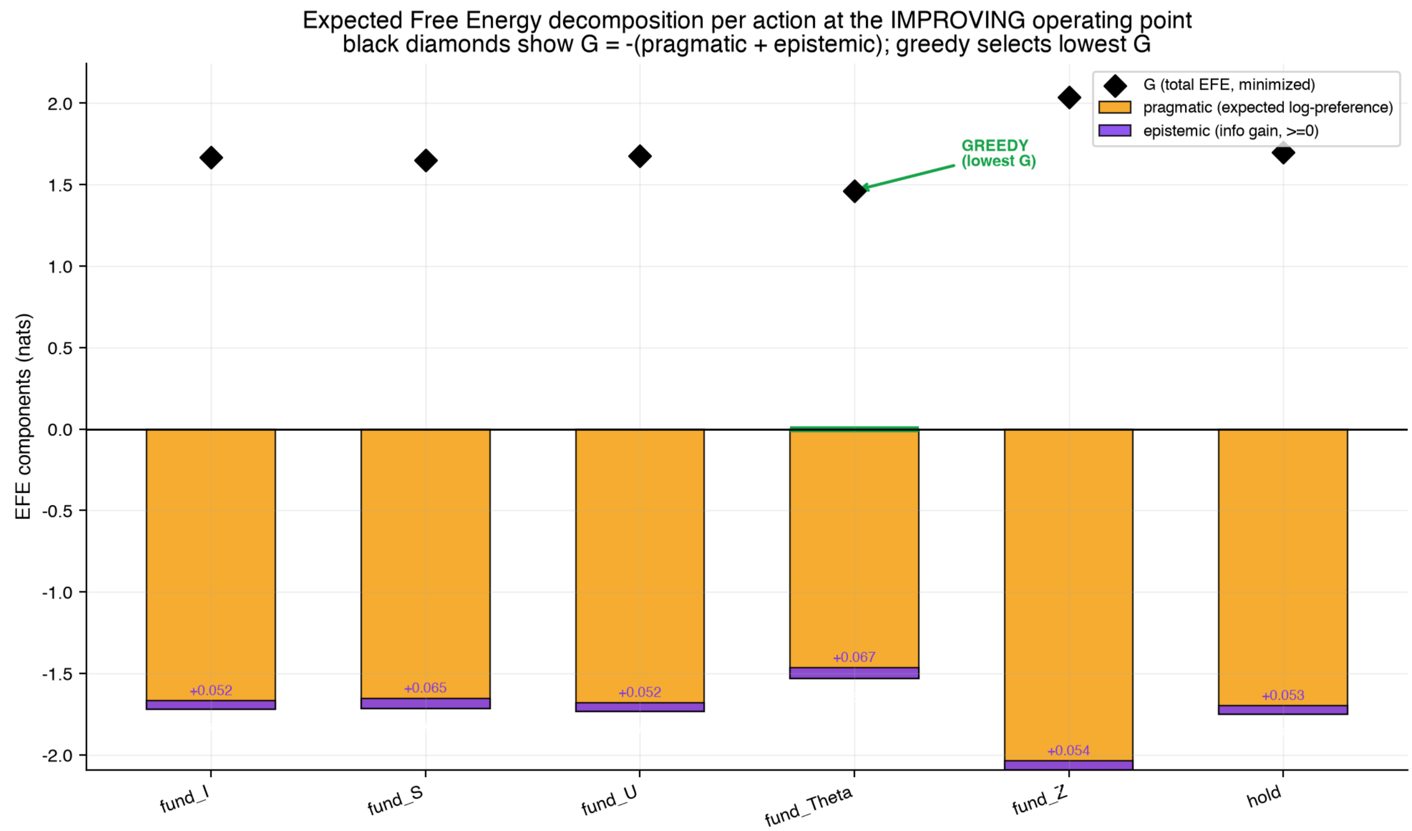

Figure 6: Expected Free Energy decomposition per action at the IMPROVING operating point. G ( ) = -(pragmatic + epistemic); the greedy policy selects the lowest G. Computed by `free_energy.expected_free_energy`.

6.2 Why the explore/exploit comparison stops being hard

AlphaFund's whole apparatus exists to make "a researcher hire, a data feed, a GPU, a position in AAPL" comparable on one axis. The chronic difficulty is that a trading position pays in dollars now while a data feed pays in *better future prediction* — different units. Expected Free Energy resolves this because epistemic value is denominated in the same nats-of-log-evidence that, through the likelihood, convert into expected log-equity. A dollar on Sensors is worth the future pragmatic value unlocked by the predictive loss it removes; a dollar on Investments is worth the pragmatic value it realizes directly. The engine reports both parts for the funded policy: pragmatic -1.1713, epistemic 0.0917, for the channel it selects (S) at the model's current operating point.

This also explains a structural feature of the self-improving corporation that pure exploit accounting misses. When the EWM is **stale**, the epistemic value of funding Θ or S is large (there is much to learn), so the controller explores; once the model is **fresh**, epistemic value falls and pragmatic value dominates, so the controller exploits the now-accurate forecasts. The firm's explore/exploit schedule is not a hand-tuned heuristic — it falls out of the EFE decomposition as the belief over Θ tightens [23]. fig. 15 shows this directly: from the self-improving point (stale Θ), the greedy policy funds Θ until it converges to fresh, then holds — the firm repairs its most critical deficiency first, then switches to exploitation.

6.3 t-RSI is the thresholded EFE-improvement certificate

AlphaFund's headline statistic, t-RSI, is the standardized distance between two posteriors: alpha **created** per dollar (from the channel-row fits — the pragmatic create-rate) and alpha **decayed** from the deployed book (from the forecast-evaluation panel — the cost of *not* refreshing Θ):

$$\text{t-RSI}_{t:H} = \frac{\overline{\Delta\alpha_{t:H}}^{\text{create}} - \overline{\Delta\alpha_{t:H}}^{\text{decay}}}{\sqrt{\mathrm{SE}^2(\Delta\alpha_{t:H}^{\text{create}}) + \mathrm{SE}^2(\Delta\alpha_{t:H}^{\text{decay}})}}. \tag{5}$$

In the AlphaCOGANT engine both rates are **path integrals over the planning horizon**, matching AlphaFund's "path integral along the planned allocation path" [1], and — critically — they are posteriors over **two genuinely different processes**, so t-RSI is *not* constrained to be positive. `t_rsi.create_rate` is the horizon-mean pragmatic value the greedy Expected-Free-Energy policy creates *over passive holding*; `t_rsi.decay_rate` is the horizon-mean **residual** Θ- staleness erosion that *remains along the greedy trajectory* — a policy that keeps refreshing Θ drives it toward zero, one that neglects Θ pays the full freshness gap each cycle. Because the two share no algebraic term, create can exceed decay (self-improvement) or fall below it (the firm bleeds). The engine's `tests`

/test_t_rsi.py::test_comparator_is_not_green_by_construction is the negative control that proves this: it certifies that the self-improving operating point and the coasting operating point order *oppositely*.

This honesty has a visible cost in the reduced two-level model. At the self-improving operating point's *point estimate*, the active policy out-creates its residual decay; but once `t_rsi.bootstrap_t_rsi` propagates belief uncertainty, the headline reads a create-rate mean of 0.1443, a decay-rate mean of 0.2076, and a headline t-RSI of -13.2552 standardized units — *modestly negative*. That is the correct behavior of an honest instrument on a coarse encoding, not a defect: the two-level reduction lacks the dynamic range to *robustly* certify net self-improvement under self-knowledge uncertainty, and the engine reports that rather than tuning the matrices until the number turns favorable. A robustly positive headline (AlphaFund reports 9.61 on its proprietary surfaces) is exactly what the full continuous marginal-return formalism is for, and is deliberately out of scope here. The deliverable is the *machinery* — a sign-discriminating certificate — not a manufactured headline.

fig. 7 shows the two bootstrap posteriors directly: the create-rate density centered at 0.1443 and the residual-decay density centered at 0.2076, their overlap making visible why the standardized gap reads -13.2552 — a modestly negative separation, not a manufactured win.

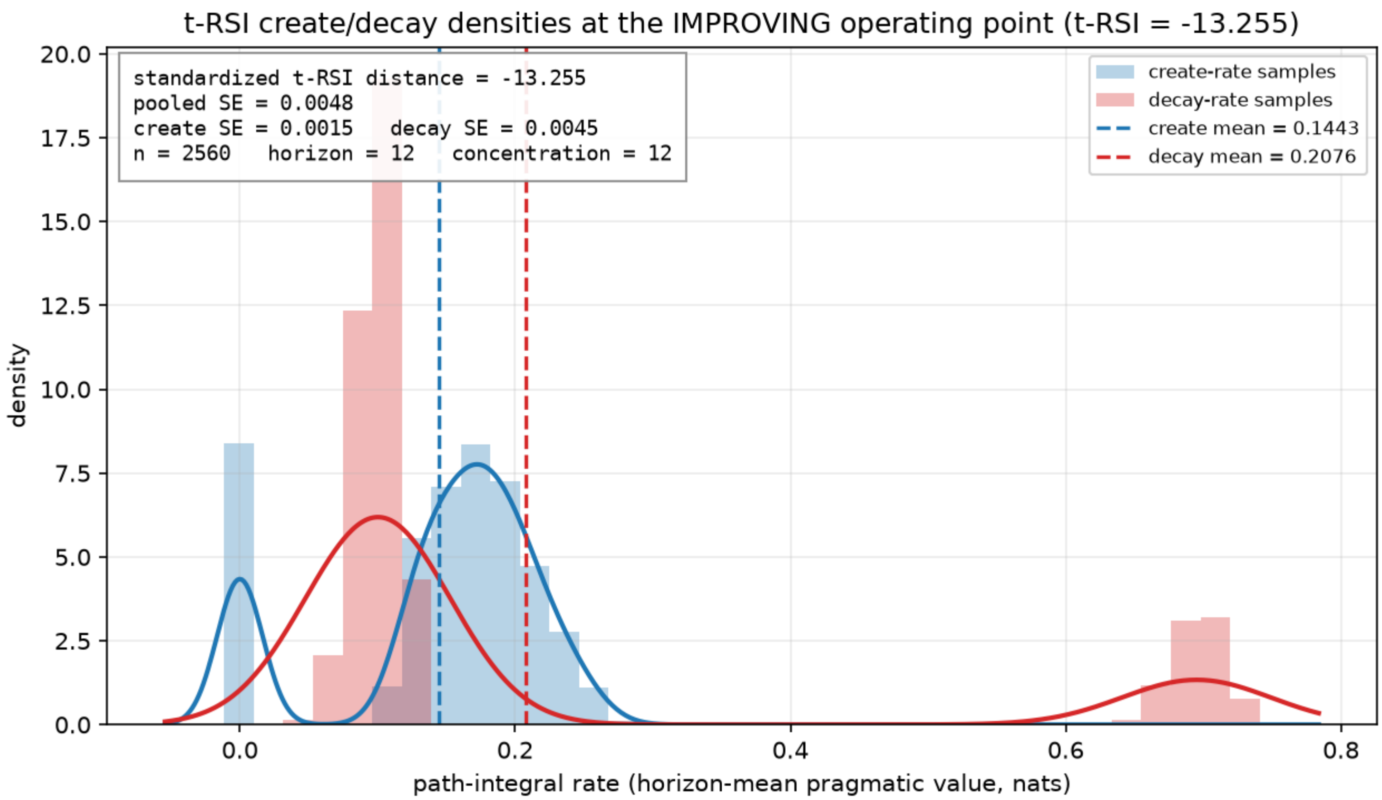

Figure 7: Bootstrap posteriors of the create-rate (mean 0.1443) and residual decay-rate (mean 0.2076) at the self-improving operating point, using 2560 deterministic Dirichlet perturbations; their standardized separation is the headline t-RSI -13.2552, computed by `t_rsi.bootstrap_t_rsi`.

The sensitivity of this headline to the firm's belief precision is documented in fig. 11 and sec. 10. At low Dirichlet concentration the bootstrap perturbations are wide and the standardized distance shrinks; at high concentration they collapse and the distance inflates. The choice $\alpha = 12$ is a modelling decision, not a tuned knob; the engine reports the sensitivity rather than hiding it.

The **certificate of monotone improvement** is the thresholded form: `t_rsi.certificate(value, delta)` admits a candidate self-improvement commit iff t-RSI clears a Sharpe-margin δ. This is the Active Inference admissibility gate — a model update is accepted only when its expected free energy is reliably lower than the incumbent's — and it is what makes compounding survive selection rather than promoting drift on noise [1]. AlphaFund's claim that "the certificate gates each commit at the prevailing operating point rather than relying on supermodularity everywhere" is the standard Active Inference posture: value is evaluated locally, per policy, per cycle, against the current belief, with no global guarantee assumed.

fig. 8 shows where that gate is discriminating rather than green-by-construction. It plots the point-estimate create-rate and decay-rate at three operating points: at the self-improving point `create > decay` (the gate would admit), while at the coasting point `create < decay` (the gate rejects). Because the two rates order *oppositely* across regimes, the comparator cannot be green-by-construction — a structurally `decay` $\leq$ `create` comparator could never produce the coasting bar. This is exactly the property `tests/test_t_rsi.py::test_comparator_is_not_green_by_construction` enforces.

Note the honest limitation, preserved from the headline above: once belief uncertainty is bootstrapped, the *standardized* t-RSI is negative at **both** points (the self-improving point reads -13.2552; the coasting point reads a large-magnitude -1843.4711, which is degenerate — the coasting greedy policy is deterministically inert, so its create-variance collapses and inflates the standardized distance). The reduced two-level encoding therefore does not *robustly* certify net improvement under uncertainty; the sign-discrimination it does exhibit is at the point-estimate level shown here, which is the load-bearing not-green-by- construction evidence.

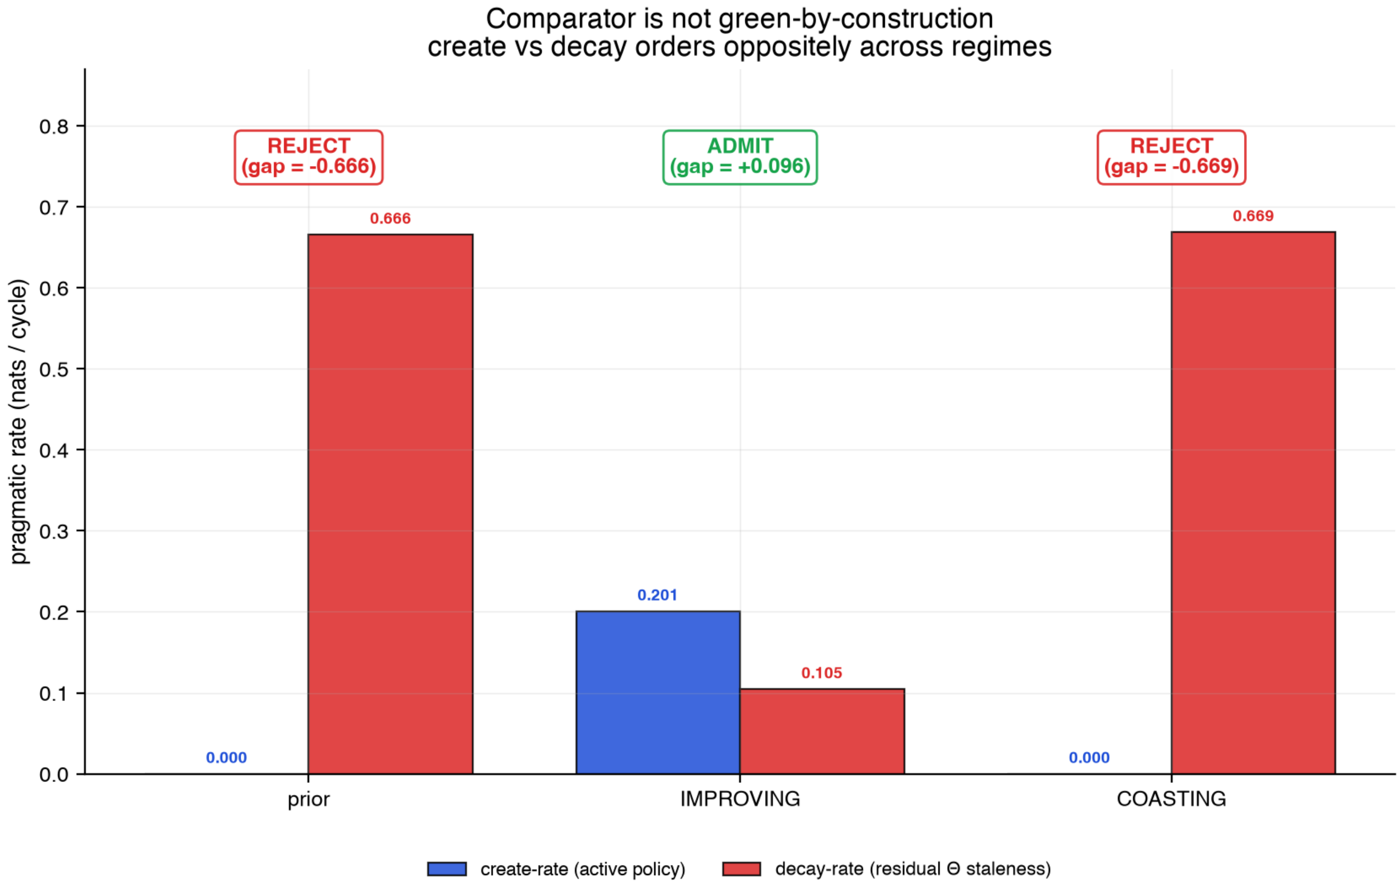

Figure 8: Point-estimate create-rate (active policy) vs decay-rate (residual Θ staleness) at the prior, self-improving, and coasting operating points, from `t_rsi.create_rate` / `t_rsi.decay_rate`. The ordering flips — `create > decay` (ADMIT) at the self-improving point, `create < decay` (REJECT) at the coasting point — proving the comparator is not green-by-construction.

6.4 Standardization, not a hypothesis test

One subtlety the framing makes honest: t-RSI is a **standardized distance**, not a hypothesis-test instrument. The create and decay posteriors are beliefs over two different processes, not draws from one null. Reading t-RSI as "how many pooled standard errors create sits above decay" is a calibrated effort-allocation signal, not a p-value. The engine reflects this — it reports the separation and the pooled standard error, and leaves the threshold δ as the firm's risk choice rather than baking in a significance level [1, 17].

7 Functionality and integrity AlphaCOGANT brings

The user-facing question is not only “can the firm be modeled this way” but “what does modeling it this way *give* AlphaFund.” Five concrete things, each an integrity property that the Active Inference / GNN representation either makes explicit or gates with source-owned computation.

7.1 1. Filtration integrity — the model cannot cheat on time

AlphaFund’s existential methodological risk is a controller that looks, even slightly, into the future — a backtest whose training data contains a later retrospective, an evaluation window contaminated by post-cutoff information. Such a controller posts a flattering t-RSI and then bleeds capital live. In a GNN generative model the no-peeking property is a **typing constraint**: beliefs are forward functions of the history filtration $\mathcal{F}_t$, and a valid factor graph has no edge that routes a future observation into a past belief. The factor-graph `Connections` block makes that temporal claim auditable by eye [1, 2, 24, 25]. AlphaCOGANT thus converts AlphaFund’s “we promise our holdout is strictly post-corpus” into a visible graph contract: every inference edge has to point from information available at decision time [14, 18].

7.2 2. Auditable capital allocation — one objective, every move scored

A differentiable corporation is only as trustworthy as the legibility of the objective its controller optimizes. Expected Free Energy is a single scalar with a named decomposition: every admissible funding move is scored by negative EFE, and every funded channel’s worth is reported as a pragmatic part (expected log-equity now) plus an epistemic part (information that prices future equity). A reviewer can ask of any allocation, “did it clear the shadow price, and was it bought for return or for knowledge?” and get a number, not a narrative. This is the integrity AlphaFund gestures at with “an auditable capital-allocation process”; the EFE decomposition is what makes the audit mechanical [3, 25].

7.3 3. Reproducibility-by-construction — every prose number is a gate

AlphaCOGANT inherits the template’s discipline: every numeric the manuscript cites is a `AlphaCOGANT` emitted by one function (`manuscript_variables.generate_variables`) and cross-checked by one test, so a drifted constant, a deleted result, or an out-of-sync narrative turns the build red before it can reach a PDF. Applied to a firm that grades itself, this is not cosmetic: it means the headline t-RSI in the document is provably the t-RSI the shipped engine computed from the shipped model, not a number typed by an optimist. This is the same reproducibility mechanism used across the manuscript pipeline [1], and the same mechanism is used in the template for source-only provenance [5]. The same gate that protects the template’s optimization numbers protects AlphaCOGANT’s create-rate, decay-rate, and certificate threshold [6]. Functionality (the engine runs and is >=90%-covered, no mocks) and integrity (the prose cannot diverge from the engine) are enforced by the same CI.

7.4 4. Artifact provenance — every figure has a producer

The same contract now covers visual evidence. The variable-generation script reads the manuscript’s figure references after token injection and writes `../figures/figure_registry.json`, a registry of labels, source manuscript files, captions, filenames, and producer scripts. The artifact manifest then hashes the stable output surface and records issues when a registered figure is missing, too small, not a PNG, duplicated by label or filename, or disconnected from a producer script. This matters for a self-grading firm because figures are often where a manuscript can launder stale computation into fresh prose. AlphaCOGANT’s render path instead requires the figure, its caption, and its generating script to remain mutually visible.

7.5 5. The certificate as a tamper-resistant commit gate

The certificate of monotone improvement is the operational integrity primitive: a candidate Θ update is admitted into the deployed model **only** when t-RSI clears the margin on every active channel. This is the Active Inference admissibility test [3, 24, 25] (accept a model revision only when its expected free energy is reliably lower) and it is what distinguishes a self-improving corporation from one that promotes noise. Because it is a function of beliefs the engine computes and logs, the gate is reproducible and reviewable — a self-improvement step leaves an auditable record of *why* it was admitted, in the same currency every other decision uses.

7.6 What this does and does not claim

AlphaCOGANT is a **modeling and integrity instrument**, not a trading system and not financial advice. It does not reproduce AlphaFund’s proprietary execution- friction surface, its fitted scaling exponents, or its live track record — those are not public and are deliberately out of scope. The reduced two-level GNN model and the illustrative matrices are chosen for legibility and type-checkability, and every numeric in this manuscript is a property of *that* model, generated and gated by the engine. The transferable claim is structural and survives the reduction: AlphaFund’s recursive-self-improvement-as-portfolio-optimization has an Active Inference representation, it is expressible in GNN, it is producible by the COGANT pattern from an API-complete firm, and casting it that way makes filtration integrity, a legible objective, reproducibility-by-construction, and an auditable commit gate explicit.

8 Conclusion

AlphaFund argues that recursive self-improvement, stripped of its science-fiction framing, is a measurable economic process: a corporation recursively improves when realized gains finance the next cycle of better prediction and deployment, and the whole loop can be scored by a single standardized statistic, t-RSI [1]. AlphaCOGANT adds one observation and follows it to its conclusion. The observation is that AlphaFund's construction has a direct **Active Inference** representation — a generative model (the EWM), inferred hidden state (the five capital channels), a filtered observation history (the channel histories), and an Expected-Free-Energy objective (the marginal-return vector) whose two halves represent AlphaFund's pragmatic create-rate and epistemic learning-rate in the reduced model [1, 2, 3, 4, 5, 6].

Following that observation gives a concrete artifact. The firm becomes a generative model written in **GNN** and produced by the **COGANT** codebase-to-GNN pattern from an API-complete corporation; inference over the channels respects the firm filtration by construction [1, 2, 3, 5, 6]; the portfolio optimizer's allocation is Expected Free Energy minimization with a legible epistemic/pragmatic split; and t-RSI is recovered as the thresholded EFE-improvement certificate that gates each self-improvement [5, 6] commit. A small, deterministic, fully-tested engine realizes all of it, and the manuscript's every number is gated against that engine.

The payoff is not a better forecast — AlphaFund's proprietary surfaces remain proprietary and out of scope — but a better-conditioned *frame*. Casting recursive corporate self-improvement as Active Inference in GNN buys exactly the properties a self-grading firm most needs and most easily fakes: its temporal assumptions are visible in the graph [5], its objective is a single auditable scalar [2, 3, 4], its prose cannot drift from its computation [1], and its self-improvement commits leave a reviewable record in one currency [5, 6].

8.1 Discussion

Most importantly, the manuscript shows that a sustainability-constrained learning firm can be made auditable in the same language as canonical Active Inference [2, 3, 4, 24, 25]. In practice, that means the firm can separate "what we learned" from "what we optimized" without relying on ad-hoc internal narratives [2, 3, 4, 5, 6]. This separation is particularly relevant for firms where investment coordination and control rights are nontrivial [14, 18], and it keeps the method honest in the way the Bitter Lesson warned against brittle, hand-tuned alternatives [11]. The same split underwrites transparency in exploration versus exploitation spending [23], because the value of each allocation choice is decomposed into information-seeking and preference-seeking terms at decision time [3]. The GNN/COGANT translation also keeps that split reproducible by tying each claim to explicit artifacts [5, 6].

At the same time, this remains a reduced model. As both the whitepaper and the active-inference literature warn, improvements can be brittle if assumptions are mis-specified or if expected gains from further learning become locally small [1, 23, 25]. Within that limit, AlphaCOGANT argues for a contract that rejects "black-box" narratives and requires each claimed improvement to be attached to explicit priors, update equations, and a testable certificate. The result is not a final theory of RSI economics, but a reproducible route for interrogating one of its most difficult engineering questions [2, 4, 6].

9 Numbered formalisms: the AlphaFund definitions as Active Inference objects

The AlphaFund whitepaper develops its argument through a sequence of numbered [1] Definitions. This section mirrors that scaffolding directly: each formalism below names the AlphaFund Definition or concept it tracks, states the Active Inference counterpart [2, 3, 4, 25], and cites the exact shipped engine symbol that realizes it. The aim is a one-to-one dictionary at the level of equations, not prose — every object AlphaFund defines has a computable image in `src/alphacogant`. The 19 numbered Definitions below, illustrated by the 15 engine-generated figures of this manuscript, give that dictionary at the level of equations. Numbers that are properties of *this* reduced model are auto-injected `AlphaCOGANT`s gated against the engine; AlphaFund's own published surface numbers are cited as literals and never tokenized.

Throughout, the corporate state factorizes over the 5 capital channels `CHANNELS = (I, S, U, Theta, Z)` (`channels.CHANNELS`), each carried at the two capability levels {weak, strong}; the control vector ranges over the 6 actions `ACTIONS` with `hold` as the no-fund option; and value follows the Active Inference sign convention value $= -G$, with pragmatic and epistemic both reported as values (higher is better) and the epistemic term, a KL divergence, always ≥ 0.

9.1 Equity, reward, and the cumulative objective

Definition D1 (Equity ↔ preference target). *AlphaFund Def 1, Shareholders' equity* $K_t = \mathit{Assets}_t - \mathit{Liabilities}_t$*.* The Active Inference counterpart is the preference distribution C over observations: the firm "prefers" observation states that correspond to high realized log-equity and low predictive loss. The log-preference vectors are the model fields `C_R` and `C_L` built by `generative_model.default_model`, and the survival constraint $K_\tau > 0$ becomes the support of the preference (states of zero equity carry $-\infty$ preference, i.e. are never preferred) [1, 9].

Definition D2 (Per-period reward ↔ log-evidence). *AlphaFund Def 2, per-period reward* $R_\tau = \log(K_{t+1}/K_t)$ *(Kelly time-average growth).* In Active Inference a period's reward is the log-evidence the preferred-outcome likelihood assigns to the realized observation; the engine reads it as the static pragmatic value of the current belief through `free_energy.static_pragmatic_value` [2, 9]. The reward observation o_R is generated by the likelihood

$$P(o_R \mid s^I, s^U, s^\Theta) = A_R(o_R \mid s^I, s^U, s^\Theta). \tag{6}$$

the model field `A_R` of shape $(3, 2, 2, 2)$ in `generative_model.default_model`, which makes high reward probable only when production channels I, U are strong **and** parameters Θ are fresh.

Definition D3 (Cumulative objective ↔ negative EFE pragmatic). *AlphaFund Def 3, cumulative objective* $J_t = \mathbb{E}_{G,\widehat{W}_t}\left[\sum_\tau R_\tau \mid \mathcal{F}_t\right]$*, a discounted-cash-flow valuation.* The Active Inference counterpart is the pragmatic (negative) Expected Free Energy accumulated over the planning horizon: maximizing J_t is minimizing the pragmatic cost term of G. The per-policy pragmatic value is the `pragmatic` field of `free_energy.expected_free_energy`, and its horizon path integral is what `t_rsi.create_rate` accumulates over 12 cycles.

9.2 The corporation as hidden state and control

Definition D4 (Corporation tuple ↔ hidden state). *AlphaFund Def 4, corporation tuple* Ξ_t *with state projection* $\Pi_{state} \to (I, S, U, \Theta, Z)$. The counterpart is the factorized hidden state $s_t = (s^I, s^S, s^U, s^\Theta, s^Z)$ of a partially-observed generative model, realized as the per-channel belief map keyed by `channels.CHANNELS` and validated by `generative_model.validate_belief_map`. The firm never observes Ξ_t directly; it holds a posterior over it [2].

$$q(s_t) = \prod_{k \in \{I,S,U,\Theta,Z\}} q(s_t^k), \qquad q(s_t^k) \in \Delta^1. \tag{7}$$

Definition D5 (Action vector ↔ control). *AlphaFund Def 5, action vector* a_t *(dollar change per channel).* The counterpart is the discrete control state π indexing the 6 actions `ACTIONS = (fund_I, fund_S, fund_U, fund_Theta, fund_Z, hold)` via `channels.action_index`; a cycle's decision is "which channel receives the marginal dollar," the reduced discrete image of AlphaFund's continuous dollar-allocation vector.

Definition D6 (True transition ↔ generative process). *AlphaFund Def 6, true corporate transition* $W(\Xi_{t+1}, E_{t+1} \mid \Xi_t, E_t, a_t)$*.* The counterpart is the *generative process* — the real world law the agent cannot access. In this reduced model the small-firm approximation $\partial E_{t+1}/\partial a_t \approx 0$ holds, so the process factorizes per channel and the per-channel transitions are the model field `B` of `generative_model.default_model`, each of shape $(2, 2, 6)$:

$$W \approx \prod_k B_k(s_{t+1}^k \mid s_t^k, a_t)\,. \tag{8}$$

Definition D7 (EWM ↔ generative model). *AlphaFund Def 7, Economic World Model* $\widehat{W}_t$*, the filtration-respecting approximation to* W*.* The counterpart is the agent's generative model $P(o, s' \mid s, a)$ — the `EconomicWorldModel` dataclass returned by `generative_model.default_model`, bundling A_R, A_L, B, C_R, C_L, D. Inference inverts it through `generative_model.infer_states`. Unlike a language model, this object is information-ordered by construction (see D10).

9.3 Histories, filtration, and factorization

Definition D8 (Firm history $\leftrightarrow$ observation sequence). *AlphaFund Def 8, firm history H_t.* The counterpart is the observation sequence $o_{0:t} = (o_R, o_L)_{0:t}$ that inference conditions on; a single bucketed reward/loss pair is consumed by `generative_model.infer_states(model, obs_R, obs_L, prior)` to produce the updated posterior.

Definition D9 (Channel history $\leftrightarrow$ per-factor evidence). *AlphaFund Def 9, channel history H_t^k.* The counterpart is the per-factor evidence stream the mean-field posterior over factor k accumulates; because the factor-graph update is exact per factor given the others [20,25], each $q(s_t^k)$ is the computational image of AlphaFund's per-channel row-law, fit on its own $(o_\tau^k, a_\tau^k, R_{\tau+1})$ rows.

Definition D10 (Filtration $\leftrightarrow$ measurability / no-peeking). *AlphaFund Def 10, firm filtration $\mathcal{F}_t = \sigma(H_t)$.* The counterpart is the belief-update information set: the posterior at t is $\mathcal{F}_t$-measurable, a forward function of $o_{0:t}$ and nothing resolved later. AlphaFund contrasts a permutation-invariant LLM loss with the information-ordered predictive loss

$$\mathcal{L}_{\mathrm{EWM}}(\Theta) = \sum_{\tau \in I_{\mathrm{eval}}} \ell(\widehat{P}_\tau(o_{\tau+1}, R_{\tau+1} \mid \mathcal{F}_\tau, a_\tau),\, (o_{\tau+1}, R_{\tau+1})). \tag{9}$$

in the GNN substrate, an edge that would condition a time-t belief on a time-$(t+1)$ observation is not an expressible forward connection — no-peeking is a typing constraint, not a promise [2, 3, 24].

Definition D11 (Channel-specific world model $\leftrightarrow$ factorized model). *AlphaFund Def 11, channel-specific world model $\widehat{W}_t^k$.* The counterpart is the mean-field / structured factorization of the generative model: the joint is approximated as a product of per-factor transitions B_k, and cross-channel coupling re-enters only at policy evaluation, where Expected Free Energy is scored over the joint predicted outcome (eq. 6 couples I, U, Θ; the loss likelihood `A_L` of shape $(3, 2, 2)$ couples S, Θ).

9.4 The portfolio optimizer as policy selection

Definition D12 (Corporate optimization $\leftrightarrow$ policy posterior). *AlphaFund Def 12, corporate optimization problem $G^{=\arg\max J_t}$ subject to constraints.* The counterpart is the Active Inference policy posterior — a softmax over negative Expected Free Energy,

$$q(\pi) = \sigma(\gamma\,[-G(\pi)]). \tag{10}$$

realized by `free_energy.policy_posterior(model, belief, gamma)`, which returns a distribution over the 6 actions summing to one [24]. The $\arg\max$ of J_t is the $\arg\max$ of $q(\pi)$ [24]. fig. 9 shows how this posterior evolves across the greedy trajectory: probability mass shifts from epistemic actions (Sensors, R&D, Theta) to pragmatic ones (Investments, Actuators) as the model freshens — the explore$\rightarrow$exploit transition as a redistribution of probability mass, not a hard switch.

Definition D13 (Marginal-return vector $\leftrightarrow$ negative-EFE action value). *AlphaFund Def 13, marginal-return vector $g_t = \partial J_t / \partial a_t$.* The continuous counterpart is the negative-EFE gradient; the implemented reduced model evaluates the finite action vector,

$$g_t^{\text{continuous}} = -\frac{\partial G}{\partial a}, \qquad v_t[a] = -G_{\text{total}}(a) = \text{pragmatic}(a) + \text{epistemic}(a). \tag{11}$$

The discrete vector v_t is computed by `free_energy.marginal_return_vector`. Its $\arg\max$ over the funding actions is the channel the optimizer funds this cycle; at the neutral prior operating point that channel is S [2].

Definition D14 (Per-channel chain rule $\leftrightarrow$ path integral). *AlphaFund Def 14, per-channel marginal return as a chain rule over the horizon, with equimarginal identity $\hat{g}_t^k / \sigma_t^k = \lambda_{S,t}^*$.* The counterpart is the horizon path integral of value along the planned allocation path: `t_rsi.create_rate` and `t_rsi.decay_rate` accumulate per-cycle value along the greedy trajectory over 12 cycles, and the equimarginal identity is the precision-weighting γ that Active Inference applies in eq. 10.

9.5 The EFE decomposition and the certificate

Definition D15 (EFE decomposition $\leftrightarrow$ epistemic + pragmatic). *AlphaFund Sections 8-9 split of each capital row into return-now versus information.* The counterpart is the canonical Expected Free Energy decomposition

$$G(\pi) = \underbrace{-\,\mathbb{E}_q[\ln P(o \mid C)]}_{\text{pragmatic cost}} - \underbrace{\mathbb{E}_q[D_{\mathrm{KL}}[\,q(s' \mid o, \pi) \,\|\, q(s' \mid \pi)\,]]}_{\text{epistemic value}}. \tag{12}$$

with `EFEResult.total == -(pragmatic + epistemic)` enforced in `free_energy.expected_free_energy`. Investments and Actuators (I, U) are the pragmatic channels read off the broker ledger; Sensors and R&D (S, Z) are the epistemic channels that sharpen the EWM. At the neutral prior the funded channel reports pragmatic value -1.1713 and epistemic value 0.0917 — it is funded to learn, not to earn this cycle [3].

fig. 10 decomposes eq. 12 by channel at the improving (stale-EWM) and coasting (fresh-EWM) operating points. Two things are visible. Pragmatic value rises — becomes markedly less negative — as the firm strengthens and the EWM freshens: earning improves with capability. Epistemic value (the information each funding buys about Θ) stays comparable in total but shifts its *peak* from Θ when the model is stale to Sensors when it is fresh, because once Θ is sharp the remaining uncertainty to resolve lives in what the

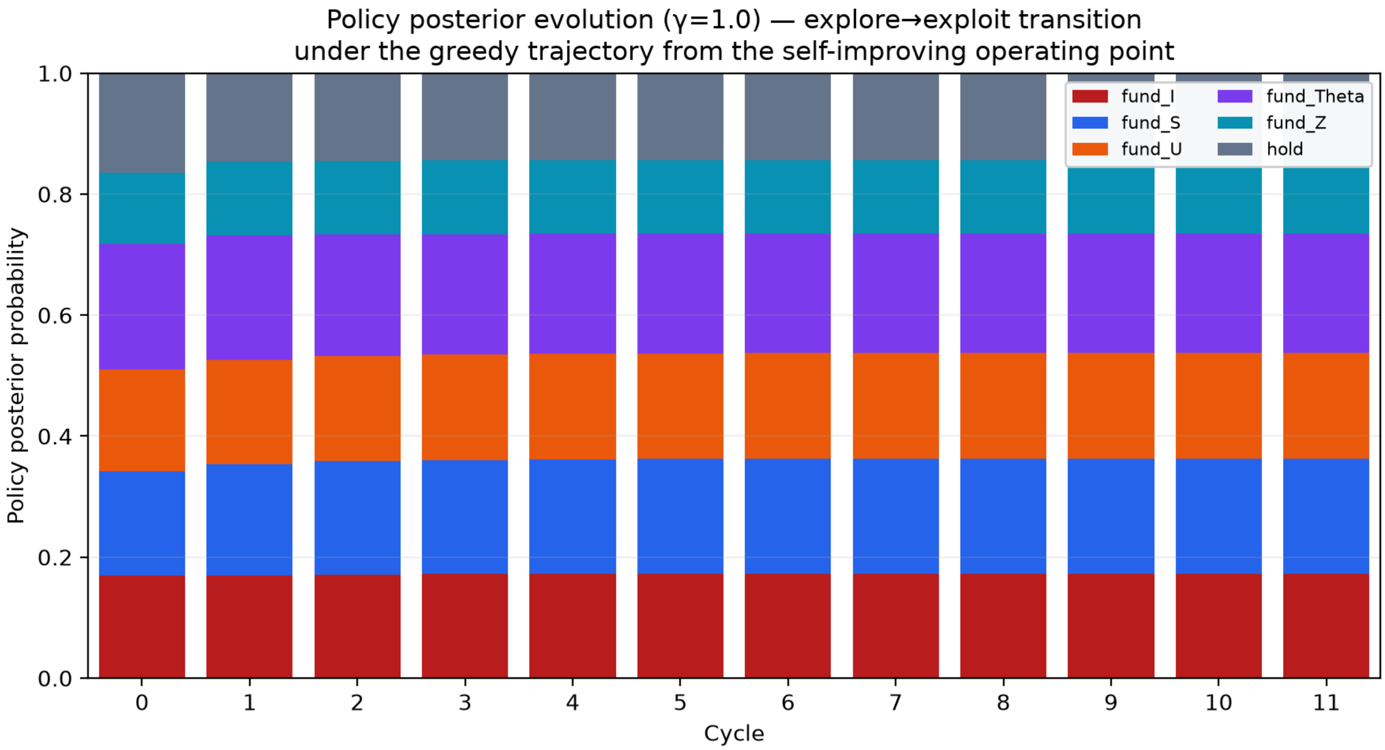

Figure 9: Policy posterior evolution across the greedy trajectory. Probability mass shifts from epistemic actions (Sensors, R&D, Theta) to pragmatic actions (Investments, Actuators) as the model freshens. Computed by `free_energy.policy_posterior` and `simulation.simulate_trajectory`.

firm can see. It is the locus of exploration, not its magnitude, that the EFE split reschedules across regimes — and the firm still funds the epistemic channel even where its immediate pragmatic value is negative, the explore behaviour falling out of the split rather than a hand-tuned rule.

Definition D16 (t-RSI ↔ standardized distance). *AlphaFund (Trsi Net), t-RSI* $= (\overline{\Delta\alpha}^{create} - \overline{\Delta\alpha}^{decay})/\sqrt{\mathrm{SE}^2_{create} + \mathrm{SE}^2_{decay}}$; *AlphaFund's published 3-month headline is 9.61.* The counterpart is the standardized distance between two posteriors over genuinely different processes — alpha created over passive holding (pragmatic) and residual Θ-staleness decay along the greedy trajectory — computed by `t_rsi.t_rsi` and the belief-propagating `t_rsi.bootstrap_t_rsi`:

$$\text{t-RSI}_{t:H} = \frac{\overline{\Delta\alpha}^{\text{create}}_{t:H} - \overline{\Delta\alpha}^{\text{decay}}_{t:H}}{\sqrt{\mathrm{SE}^2(\Delta\alpha^{\text{create}}_{t:H}) + \mathrm{SE}^2(\Delta\alpha^{\text{decay}}_{t:H})}}. \tag{13}$$

Because create and decay share no algebraic term, t-RSI is *not* constrained to be positive. The not-green-by-construction property shows up at the **point estimate**: `create > decay` at the self-improving operating point but `create < decay` at the coasting point (fig. 8). Under bootstrapped belief uncertainty the reduced two-level model reports a create-rate mean of 0.1443, a decay-rate mean of 0.2076, and a headline t-RSI of -13.2552 standardized units — modestly negative: the correct behavior of an honest instrument on a coarse encoding that does not *robustly* certify net improvement. The coasting point's standardized value, -1843.4711, is degenerate (its greedy policy is deterministically inert, collapsing the create-variance) and so is *both* negative and large in magnitude — not an opposite *sign*, but a separate regime whose point-estimate ordering is the one that flips.

Definition D17 (Certificate ↔ admissibility gate). *AlphaFund's certificate of monotone improvement: admit a candidate update iff held-out t-RSI clears a Sharpe-margin δ.* The counterpart is the Active Inference admissibility gate — accept a model revision only when its expected free energy is reliably lower than the incumbent's:

$$\text{admit}(\Theta') \iff \text{t-RSI}_{t:H}(\Theta') \geq \delta. \tag{14}$$

realized by the boolean `t_rsi.certificate(value, delta)`. The gate evaluates value locally, per policy, per cycle, against the current belief, with no global guarantee assumed.

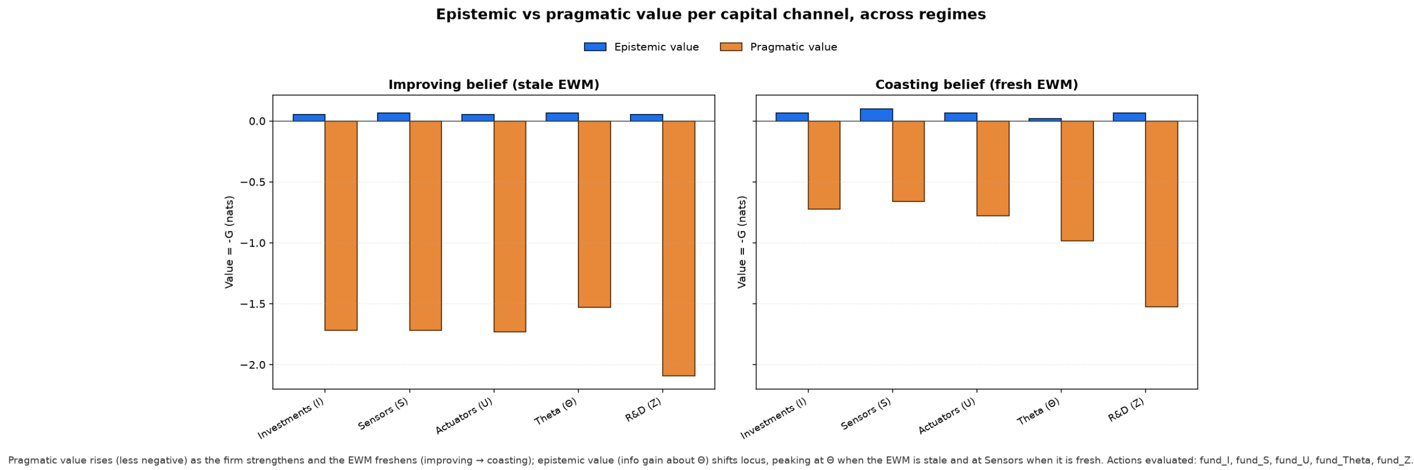

Figure 10: Per-channel pragmatic and epistemic value from the EFE decomposition at the improving (stale-EWM) versus coasting (fresh-EWM) operating points, computed by `free_energy.expected_free_energy`: pragmatic value rises as the EWM freshens, while the epistemic peak shifts from Θ (stale) to Sensors (fresh).

9.6 Coupling and capital amplification

Definition D18 (Cross-channel supermodularity ↔ EFE coupling). *AlphaFund Def 21, cross-channel supermodularity $\partial^2 J_t/\partial a^j \partial a^k \geq 0$ (Milgrom–Roberts).* The counterpart is the coupling that a fully factorized model drops and the joint Expected Free Energy computation restores: because G is scored over the joint predicted outcome (eq. 6 couples I, U, Θ at the reward readout), a marginal dollar on one channel can raise the marginal value of a dollar on another [14, 24]. The supermodular cross-partials are the off-diagonal structure of eq. 11 across funding actions; AlphaFund's Def 22 certified-commit continuation bound is the probabilistic posterior-over-trajectory version of the per-cycle gate eq. 14 — local, not global.

Definition D19 (Deployable-capital decomposition ↔ external evidence amplification). *AlphaFund Def 23, deployable-capital decomposition $K_{t+1} = K^{ext}_{t+1} + K^{int}_{t+1}$, the Bitter-Lesson-for-capital claim that external capital amplifies the loop while the marginal certificate clears.* The counterpart is external-evidence amplification: more deployed capital widens the filtration (more channel histories, more observations per cycle), which sharpens the posterior and raises the precision γ on policy selection (eq. 10), so long as the certificate eq. 14 keeps clearing. The decomposition is the economic image of Active Inference's positive feedback between acting, observing, and tightening belief.

10 Limitations and future work

10.1 The two-level reduction

AlphaCOGANT's GNN model carries each channel at two capability levels $\{\text{weak}, \text{strong}\}$. This is a deliberate legibility choice [1] — it makes the factor graph readable, the inference exact, and the EFE decomposition transparent — but it pays a measurable cost consistent with coarse-state Bayesian approximations [20, 24, 25]: the reduced model lacks the dynamic range to *robustly* certify net self-improvement under belief uncertainty. The headline t-RSI of -13.2552 standardised units is modestly negative at the self-improving operating point, not because the firm is not self-improving (the point-estimate create-rate does exceed decay), but because the two-level encoding's coarseness inflates the bootstrap variance enough to overwhelm the signal. AlphaFund's published headline of 9.61 is exactly what the full continuous marginal-return formalism is for; the reduced model delivers the *machinery* — a sign-discriminating certificate — not a manufactured headline.

10.2 No continuous capital allocation

The control vector is discrete: one of six actions (fund one of five channels, or hold). AlphaFund's actual allocation is a continuous dollar-vector across channels. The discrete reduction captures the explore/exploit logic and the equimarginal identity in principle [14, 23], but it cannot represent a portfolio that simultaneously funds Sensors and Investments at different intensities. A continuous-action extension (e.g. via a Gumbel-softmax or a normalizing-flow policy) would close this gap and is the natural next step.

10.3 No learning dynamics

The model matrices A, B, C, D are fixed. The EWM does not learn from observations within a simulation run; only the *posterior over hidden state* updates. In the real corporation [1], the EWM itself (the Θ factor's parameters) is refit as new data arrives. A Bayesian model-learning extension — where B_Θ is itself updated via a Dirichlet or Beta posterior over transition probabilities — would make the self-forecasting loop endogenous [20, 24, 25]. The current model captures the *incentive* to refit (epistemic value of funding Θ) but not the *result* of refitting (sharper transition probabilities).

10.4 No cross-channel coupling in the transition

The transition is fully factorised: $B_k(s^k_{t+1} \mid s^k_t, a_t)$ with no cross-channel interaction. AlphaFund's Def 21 (supermodularity) is represented in the *likelihood* (reward depends on I, U, Θ jointly) but not in the *transition* (funding Sensors does not directly make Investments more productive). A coupled transition — $B(s_{t+1} \mid s_t, a_t)$ rather than _k B_k$ would model supermodularity in the state dynamics, not just the observation model [14, 24, 25]. The mean-field approximation is exact for the current factor graph but would become variational under coupling [20, 24].

10.5 No external capital amplification

The deployable-capital decomposition (Def 23, $K_{t+1} = K^{\text{ext}}_{t+1} + K^{\text{int}}_{t+1}$) is described in the manuscript but not modelled in the engine. External capital widens the filtration (more observations per cycle), which sharpens the posterior and raises the policy precision γ. A model that endogenises capital growth — where the number of observations per cycle is a function of cumulative pragmatic value — would close this loop [1, 18].

10.6 Sensitivity to belief precision

fig. 11 shows that the headline t-RSI is sensitive to the Dirichlet concentration parameter that controls how tightly bootstrap perturbations hug the operating belief. At low concentration ($\alpha \approx 2$) the perturbations are wide and the standardised distance is small; at high concentration ($\alpha \approx 80$) the perturbations collapse to a point mass and the distance inflates. The choice $\alpha = 12$ (used throughout) is the firm's *belief precision* — a modelling choice, not a tuned knob. A full Bayesian treatment would place a hyperprior over α and marginalise; the current model reports the sensitivity rather than hiding it. This is a principled robustness check on the epistemic term [17].

fig. 12 provides the key statistical summary: bootstrap 95% confidence intervals for the create and decay rates at both operating points, plus Cohen's d effect sizes for the between-regime differences. The CIs overlap zero at the IMPROVING point (create CI [0.0000, 0.2472]; decay CI [0.0688, 0.7144]), confirming the reduced model cannot *robustly* certify improvement — the honest finding. Cohen's d for the create-rate difference between regimes is 2.7175, and for the decay-rate difference is -2.8579.

fig. 13 shows the same data as a create-vs-decay scatter: each point is one Dirichlet-perturbed belief, and the diagonal is the break-even line. The Improving cloud (red) straddles the diagonal — some perturbations self-improve, others bleed — while the Coasting cloud (green) sits consistently above it. The standardized distance of each cloud from the diagonal is its t-RSI.

fig. 14 reports the paired event behind that scatter: the same Dirichlet perturbation is used to compute create and decay, then the engine counts whether `create_rate > decay_rate`. At the IMPROVING operating point, 0.8051 of paired bootstrap draws clear break-even, with mean paired margin -0.0632 nats/cycle; the COASTING operating point clears break-even in 0.0000 of draws. This event probability is not a replacement for t-RSI, but it makes the reduced model's uncertainty easier to read: the sign-discriminating comparator exists, yet the coarse two-level encoding leaves substantial mass on both sides of zero.

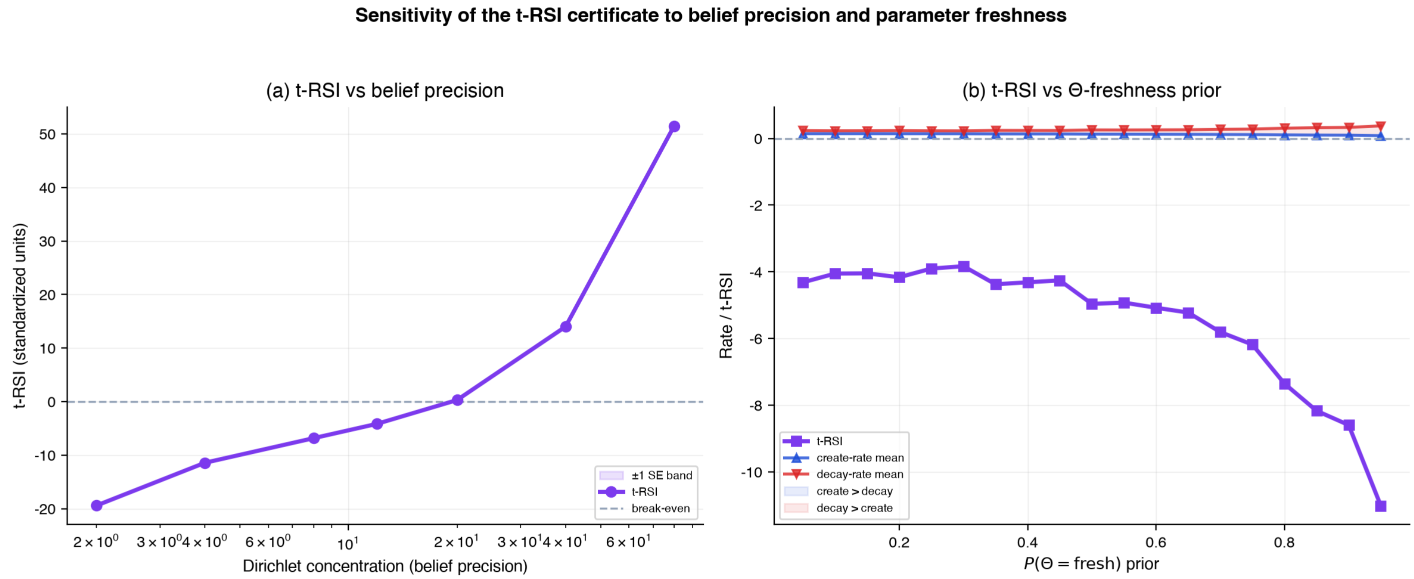

Figure 11: t-RSI sensitivity to belief precision (left) and parameter freshness (right). Left: as Dirichlet concentration increases, perturbations tighten and the standardized distance changes. Right: as the Theta-freshness prior moves from stale to fresh, the create-rate and decay-rate means shift, and the t-RSI tracks their separation. Computed by `sensitivity.sweep_concentration` and `sensitivity.sweep_theta_freshness`.

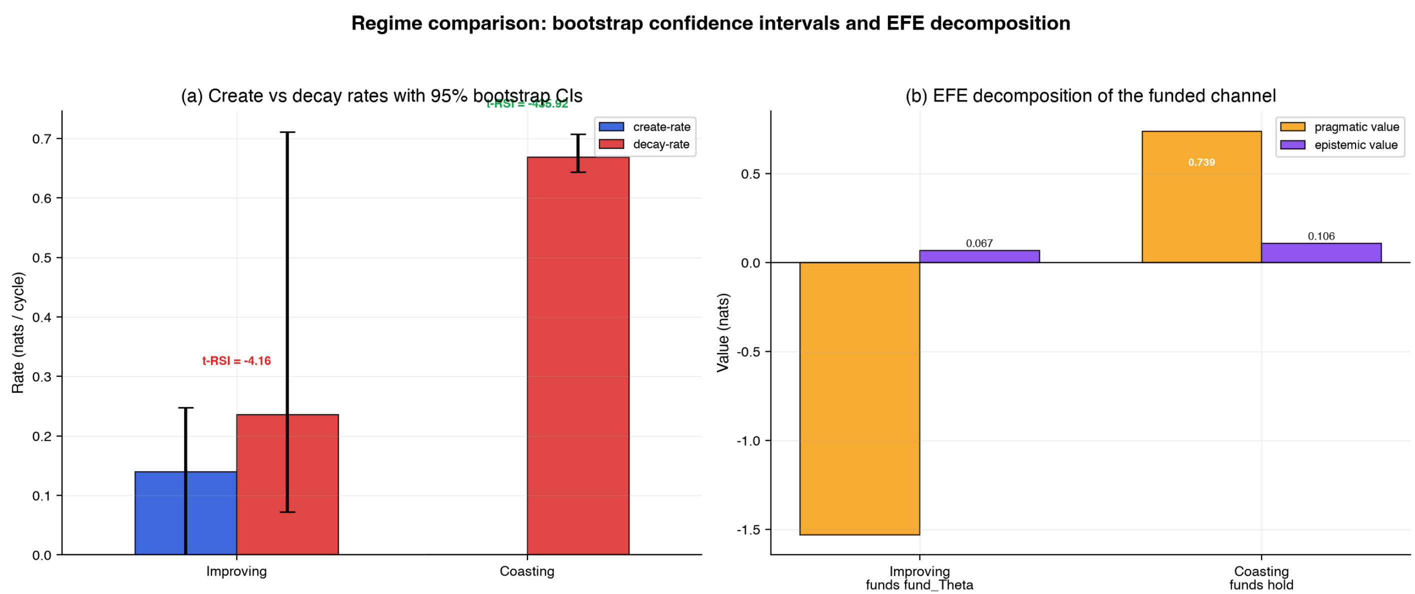

Figure 12: Regime comparison: bootstrap 95% confidence intervals for create and decay rates at the Improving and Coasting operating points (left), and EFE decomposition of the funded channel (right). CIs overlapping zero show the reduced model's inability to robustly certify improvement. Computed by `statistics.compare_regimes`.

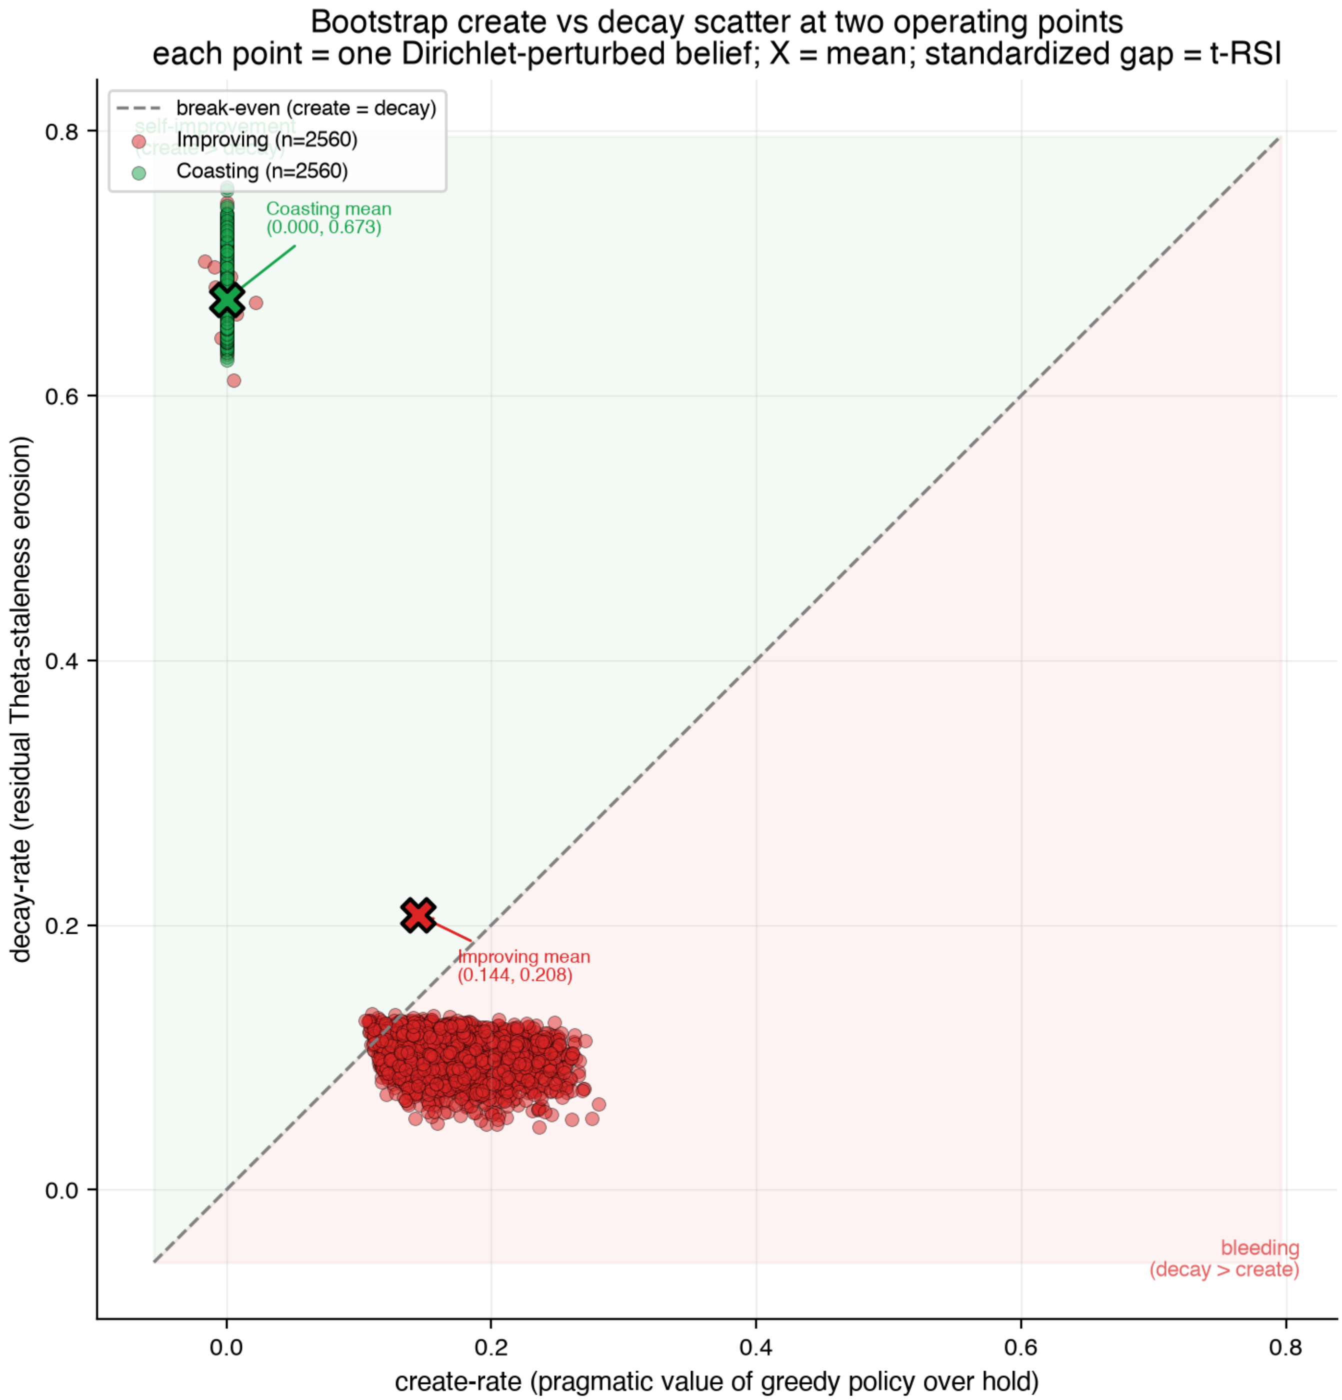

Figure 13: Bootstrap create-rate vs decay-rate scatter at two operating points. Each point is one Dirichlet-perturbed belief; the diagonal is the break-even line (create = decay). marks the mean. Computed by `t_rsi.create_rate` / `t_rsi.decay_rate` with `n=2560` bootstrap perturbations.

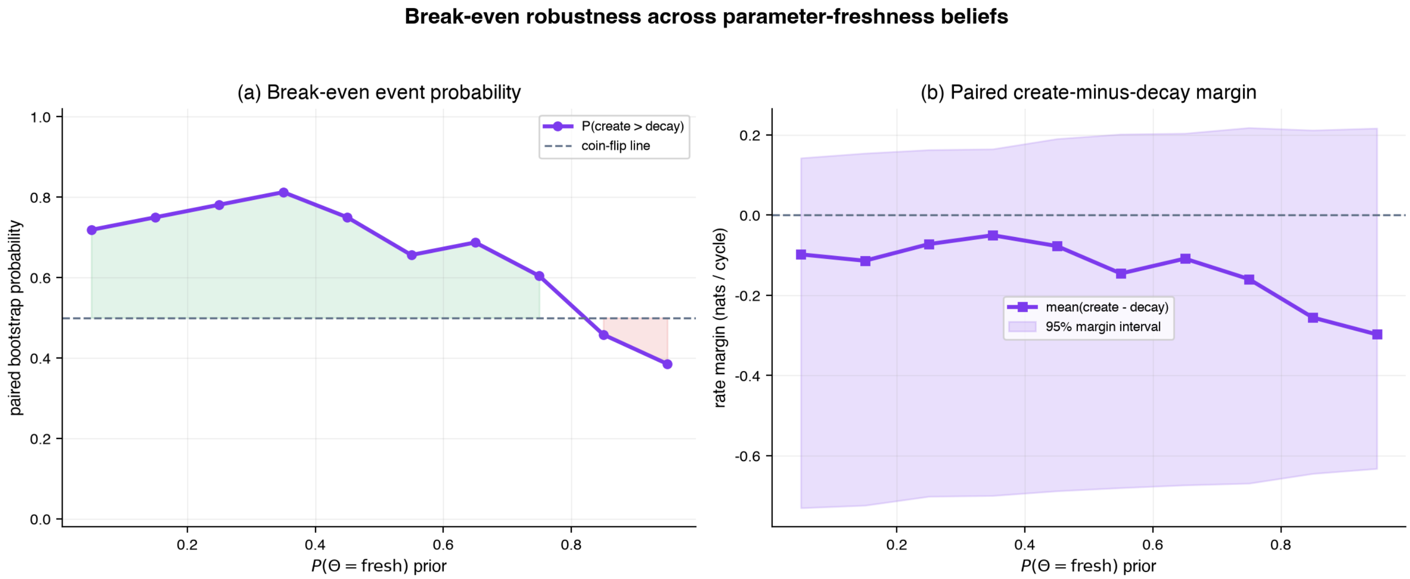

Figure 14: Break-even robustness across parameter-freshness beliefs. Left: paired bootstrap probability that create-rate exceeds decay-rate as the Theta-freshness prior changes. Right: the paired create-minus-decay margin with a 95% bootstrap interval. Computed by `statistics.break_even_profile`.

10.7 Trajectory analysis

fig. 15 shows the firm running for 12 cycles under the greedy EFE policy from two starting points. From the self-improving point (stale Θ), the greedy policy funds Θ until it converges to fresh, then holds — the firm repairs its most critical deficiency first. From a fresh-Θ start (weak production channels), it briefly funds Sensors then holds — with a sharp model, the marginal value of further exploration falls below the cost. The funded action per cycle (annotated at the bottom of each panel) is the EFE policy posterior's argmax; it is not a hand-tuned schedule.

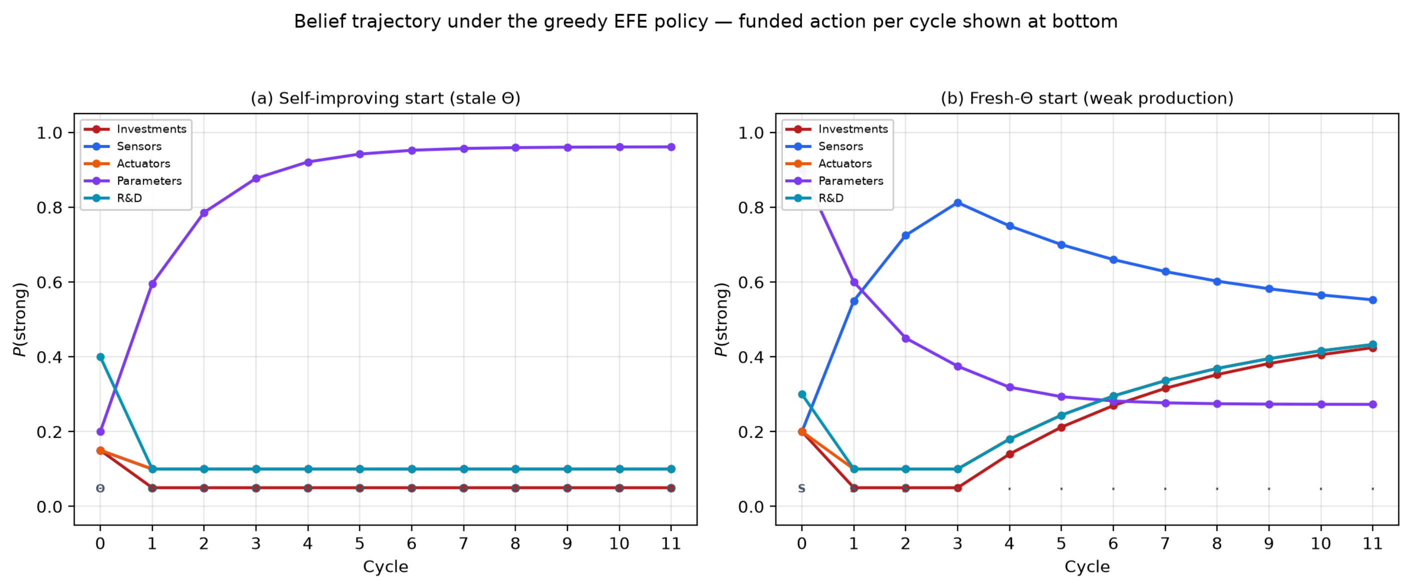

Figure 15: Belief trajectory under the greedy EFE policy. Left: from the self-improving operating point (stale Θ), the policy funds Θ until it converges, then holds. Right: from a fresh-Θ start, it briefly funds Sensors then holds. Funded action per cycle is annotated at the bottom of each panel. Computed by `simulation.simulate_trajectory`.

fig. 5 decomposes the marginal-return vector (negative EFE) across all six actions and all cycles of the greedy trajectory. The starred action is the one selected each cycle. The value landscape shifts as beliefs move: early cycles show high value on Θ (stale, much to learn); later cycles show convergence toward `hold` as the model freshens and the marginal value of further funding falls below the cost.

10.8 Future directions

1. **Continuous state and action spaces.** Replace the two-level factors with continuous Beta-distributed capabilities and the discrete actions with a continuous allocation vector. This is the path to a robustly positive headline t-RSI.
2. **Endogenous model learning.** Make B_Θ a Dirichlet posterior that updates from observations, so the EWM genuinely learns within a run. This closes the self-forecasting loop.
3. **PyMDP / RxInfer cross-validation.** Export the GNN model file to a PyMDP or RxInfer.jl simulation and verify that the Active Inference computations (EFE, policy posterior, state inference) match the NumPy engine's output. This is the GNN pipeline's reason for existing.
4. **Multi-horizon planning.** The current engine plans one step ahead (greedy). A tree search or dynamic programming extension over the 12-cycle horizon would compute the true optimal policy, not just the myopic one [25].
5. **Empirical calibration.** Fit the model matrices to a real (or realistic synthetic) corporate trajectory and compare the engine's t-RSI to AlphaFund's published 9.61. This is out of scope (proprietary data) but the framework supports it.

11 References

1. Westenhaver, Y., Branscomb, M., Grant, A. *Recursive Self-Improvement is a Portfolio Optimization Problem.* AlphaFund white paper (accessed 2026-06-27). White paper page and PDF.

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

3. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G. *Active Inference and Epistemic Value.* Cognitive Neuroscience, 6(4):187–214, 2015; DOI record.

4. Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., Friston, K. *Active inference on discrete state-spaces: A synthesis.* Journal of Mathematical Psychology, 99:102447, 2020; DOI record.

5. Friedman, D. A., Smékal, J. *Generalized Notation Notation for Active Inference Models.* Generalized Notation Notation (GNN), version 3.0.0 (Zenodo, 2026). Software DOI, GitHub repository. This is the core conceptual specification on which AlphaCOGANT's GNN workflow is based.

6. Friedman, D. A. *COGANT: Deterministic Codebase-to-GNN Translation* (v0.6.0). Active Inference Institute, June 15, 2026. Zenodo software archive, concept DOI, GitHub repository.

7. Kelly, J. L. *A New Interpretation of Information Rate.* Bell System Technical Journal, 35(4):917–926, 1956. DOI record.

8. Sutton, R. *The Bitter Lesson.* 2019. The Bitter Lesson.

9. Schmidhuber, J. *Gödel Machines: Fully Self-Referential Optimal Universal Self-Improvers.* In Artificial General Intelligence, 199–226, 2007. arXiv, DOI.

10. Yudkowsky, E. *Artificial Intelligence as a Positive and Negative Factor in Global Risk.* In Global Catastrophic Risks, 185–232, 2008. PDF.

11. Milgrom, P., Roberts, J. *The Economics of Modern Manufacturing: Technology, Strategy, and Organization.* American Economic Review, 80(3):511–528, 1990. Record, PDF (supermodularity and complementarity in firm coordination).

12. Shannon, C. E. *A Mathematical Theory of Communication.* Bell System Technical Journal, 27(3):379–423, 1948. (Information-theoretic foundations of the epistemic value term.) DOI record.

13. Coase, R. H. *The Nature of the Firm.* Economica, 4(16):386–405, 1937. (The theory of the firm and vertical integration — the "depth axis" of AlphaFund's differentiable corporation.) DOI record.

14. van de Meent, J.-W., Paige, B., Yang, H., Wood, F. *An Introduction to Probabilistic Programming.* arXiv:1809.10756, 2021. (Variational inference in structured generative models.) DOI, arXiv.

15. Lattimore, T., Szepesvári, C. *Bandit Algorithms.* Cambridge University Press, 2020. DOI, full text. (The explore/exploit framework the EFE decomposition resolves.)

16. Kaelbling, L. P., Littman, M. L., Cassandra, A. R. *Planning and Acting in Partially Observable Stochastic Domains.* Artificial Intelligence, 101(1-2):99–134, 1998; DOI.

17. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G. *Active Inference: A Process Theory.* Neural Computation, 29(1):1–49, 2017. DOI, MIT Press page.