

The EGI Framework: A Five-Layer Multi-Agent Cognitive Architecture

—Enabling Semantic World Modeling

Chunlan Wang

Large Financial Institution, Washington DC, USA

Email: cwangmz@gmail.com

How to cite this paper: Wang, C.L. (2026)

The EGI Framework: A Five-Layer Multi-Agent Cognitive Architecture. *International Journal of Intelligence Science*, 16, 132-163.

<https://doi.org/10.4236/ijis.2026.161007>

Received: November 19, 2025

Accepted: January 19, 2026

Published: January 22, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

This paper describes the Extended General Intelligence (EGI) framework, a five-layer cognitive architecture that integrates semantic world modeling with persistent multi-agent reasoning. EGI organizes perception, semantic understanding, collaborative reasoning, orchestration, and meta-cognition into a unified structure, in which world-model information is represented as explicit semantic state variables shared across-agents. To illustrate how this architecture can be used in practice, we implement a multi-agent financial intelligence system and apply it to five years of central bank communications, macroeconomic indicators, risk factors, and market data (2020-2025). The system produces central bank tone indices, macro and risk-regime inferences, and market-reaction expectations. A comparative evaluation against a flat, single-prompt baseline suggests that the EGI-based design yields more stable semantic trajectories and more coherent cross-domain reasoning. The framework is intended as an architectural blueprint that may support further research on interpretable, long-horizon multi-agent intelligence.

Keywords

Cognitive Architecture, Multi-Agent Systems, World Models, Semantic State Variables, Persistent Reasoning, Contextual Engineering

1. Introduction

Artificial intelligence research has long explored the role of internal world models in supporting reasoning, planning, and long-horizon decision-making. Predictive world-model frameworks such as Dreamer and MuZero—rooted in latent-dynamics modeling—show how internal representations can assist with simulation and control, as demonstrated in early world-model studies [1] and later extensions

[2]. In parallel, large language models (LLMs) exhibit strong generalization and flexible semantic capabilities across many domains, consistent with broader discussions on cognitive modeling [3] and hybrid system design [4]. These two lines of work highlight complementary strengths: predictive models provide structured dynamics, while LLMs offer rich linguistic and conceptual knowledge.

Recent discussions in AI research suggest that long-horizon reasoning, semantic consistency, and coordinated multi-agent behavior may require additional architectural organization beyond predictive modeling alone or prompt-driven LLM usage. Cognitive science perspectives [5]-[7] further emphasize structured memory and distributed cognition as essential components of intelligent systems. Motivated by these perspectives, this paper introduces the Extended General Intelligence (EGI) framework, a five-layer cognitive architecture consisting of Perception, Semantic Understanding, Collaborative Reasoning, Orchestration, and Meta-Cognition.

The EGI offers an organizing viewpoint: world-model information is treated as semantic state variables embedded within a layered cognitive loop. This intent aligns with ideas on interpretable modeling [8] and multi-agent reasoning frameworks [9]. Rather than replacing existing methods, EGI provides a conceptual structure to guide system design, analysis, and future empirical investigations.

A second motivation comes from real-world domains that exhibit multi-year temporal structure, such as financial intelligence and central bank analysis. These settings require systems that integrate heterogeneous signals, maintain semantic stability, and reason about mechanism-based causal pathways—needs consistent with insights [10] on cognitive maps and world models. The EGI architecture addresses these needs by structuring perception, semantic abstraction, multi-agent reasoning, and higher-level control within a unified cognitive loop.

Recent developments in applied multi-agent systems indicate that many real-world domains—such as economics, finance, and policy analysis—require mechanisms that sustain semantic consistency over time, coordinate reasoning among multiple components, and update shared representations in a coherent and interpretable manner. The EGI architecture is motivated by these observed needs. Rather than presenting a prescriptive algorithm, it provides an organizational perspective for structuring perception, semantic abstraction, multi-agent reasoning, and higher-level control within a unified cognitive loop.

A related motivation arises from domains with long-horizon temporal structure, including financial intelligence and central bank analysis. In such settings, policy signals, macroeconomic regimes, and market reactions interact over extended periods, creating challenges for systems based on single-agent prompting or flat pipelines. These approaches may struggle with maintaining consistency across time or integrating heterogeneous signals. The EGI framework offers one possible way to structure these processes through layered reasoning and coordinated interaction among specialized agents.

2. Literature Review and Study Framework

The Extended General Intelligence (EGI) framework builds on three major lines of research: world models and predictive representations; cognitive architectures and layered models of intelligence; and LLM-based multi-agent systems. This section reviews these foundations, identifies gaps in prior work, and motivates the study framework used in this research.

2.1. World Models in AI

World models have a long history in artificial intelligence, particularly in reinforcement learning and model-based planning. Classical formulations such as predictive state representations, generative dynamics models, and planning-oriented latent spaces aimed to approximate how an environment evolves over time [1] [11] [12]. These world models enable state prediction, counterfactual or hypothetical rollouts, and long-horizon planning.

The primary limitation of these approaches, however, is that they remain representational rather than cognitive. They encode predictive structure but lack: semantic meaning, reasoning mechanisms, multi-agent coordination, persistent interpretive state, cross-domain causal modeling. Recent research in cognitive science similarly highlights that human reasoning requires structured semantic representations and memory systems [5] [10]. EGI extends this direction by embedding world models inside a deeper cognitive architecture.

2.2. Cognitive Architectures and Layered Intelligence

Early cognitive architectures—SOAR, ACT-R, and Global Workspace Theory—formalized symbolic reasoning, memory retrieval, attention, and decision-making. They established the importance of modular cognitive layers, consistent with broader discussions on distributed cognition [6] and the theory of extended mind [7]. Global Workspace Theory modeled human cognitive processes and provided structured mechanisms for: symbolic reasoning, memory retrieval, attention control, decision-making. These systems established the importance of layers and modular cognitive functions, but they predate modern foundation models and do not incorporate: semantic world representations, multi-agent workflows, persistent reasoning structures, or large-scale knowledge integration. More recent discussions in hybrid modeling [4] also emphasize the need for systems that combine symbolic reasoning with statistical world representations. EGI directly builds on this insight by introducing a layered, interpretable structure specifically designed to host and regulate modern LLM reasoning.

2.3. LLM-Based Agent Systems

Recent developments in large language models (LLMs) have enabled new classes of multi-agent systems, including LLaMA-Agents [13], Toolformer [14], workflow-supervised or structured agents [15], and memory-augmented agents [16] [17].

These approaches demonstrate how LLMs can be used to coordinate tools, distribute tasks, and manage multi-step reasoning workflows. While these systems provide important progress, many of them focus primarily on message-based coordination rather than on maintaining an explicit and persistent internal model of the world. As a result, existing designs typically do not include: A shared semantic representation that multiple agents can update and reason over; An explicit cognitive layering that distinguishes perception, semantic abstraction, reasoning, orchestration, and meta-level evaluation. Layers and implied mechanisms in human learning activities may be needed and replicated for the creation of capacity of cognition[18]. Mechanism-oriented coordination where agents reason with structured domain relationships such as policy transmission channels, macroeconomic regime dynamics, risk-mechanism pathways, or market-reaction structures. These observations do not represent shortcomings of prior work—many of these systems were not designed for long-horizon institutional reasoning—but they highlight opportunities for complementary architectural approaches. EGI aims to address this space by offering a structured way to integrate semantic representation, layered cognition, and coordinated multi-agent reasoning.

Across predictive world models, cognitive architectures, and LLM-based multi-agent systems, prior research highlights several useful components but also indicates that a further integrative structure may be needed to support transparency, stable long-horizon reasoning, and coherent multi-agent operation. In particular, complex real-world domains would benefit from a unified cognitive architecture that brings together semantic world modeling, layered reasoning, and coordinated agent behavior within a single framework.

Specifically, there remains a need to develop systems that combine:

- an explicit and interpretable semantic world model,
- a persistent internal interpretive state,
- structured and modular reasoning processes,
- coordinated multi-agent workflows, and
- mechanisms for domain-specific propagation (e.g., policy, macro, risk and markets).

The Extended General Intelligence (EGI) framework is designed as one possible way to explore this gap. It offers an architectural blueprint for organizing perception, semantic understanding, collaborative reasoning, orchestration, and meta-cognitive oversight into a coherent system suitable for long-horizon, cross-domain applications. Rather than prescribing a definitive solution, EGI provides a structured approach that may help support stability, interpretability, and coordination in multi-agent reasoning environments.

2.4. Study Framework

Based on the identified gaps, this paper develops a comprehensive framework with four components:

- 1) A five-layer cognitive architecture (EGI): This paper designs the Extended

General Intelligence (EGI) framework, a five-layer cognitive architecture supporting semantic world modeling, persistent reasoning, and coordinated multi-agent behavior. EGI enables agents to share structured semantic state variables, maintain continuity of reasoning over time, and translate world understanding into actionable inference.

While prior work has highlighted the importance of interpretability [8], world models [1], and layered cognition for robust intelligence [4], these requirements have largely remained fragmented across models and techniques. EGI unifies these needs within a scalable, system-level architecture capable of supporting long-horizon, domain-specific reasoning.

2) An end-to-end financial reasoning test case using the EGI architecture.

3) Evaluation comparing EGI vs. single-flat prompt and non-semantic hierarchical systems.

4) A Technical Appendix illustrating code examples and a minimal OOP implementation.

3. Research Design/Materials and Methods

3.1. Overview of EGI Multi-Agent Architecture: Layered Cognitive Architecture

The EGI architecture consists of five cognitive layers. Its structure is informed by prior work on cognition [5] [6] and multi-agent coordination [9]. LLM-based reasoning capabilities are enabled by general-purpose foundation models such as OpenAI model [19]. The five layers include:

L1—Perception Layer: Text ingestion, cleaning, and normalization. Ingests raw signals (documents, time series, structured indicators) and converts them into normalized inputs.

L2—Semantic Understanding (Semantic Routing): Semantic routing using engineered topic axes. Extracts and updates semantic state variables—interpretable internal representations capturing policy tone, macro regimes, risk states, or causal indicators.

L3—Collaborative Reasoning Layer: Runs mechanism-based inference chains (policy, macro, risk and markets) within or across-agents.

L4—Orchestration Layer: Task scheduling, memory passing, and conflict resolution. Coordinates agent activation sequence, routing, conflict-resolution, and output harmonization through explicit workflow logic.

L5—Meta-Cognition: Semantic World-Model Layer. Monitors inconsistencies, detects drift, revises world-model states, and maintains long-run coherence. Aggregation of agent-level inferences into structured macro-financial signals in the illustrative case.

The shared semantic world-model representations function as interpretable coordinates, similar in motivation to the cognitive-map perspective [10]. In the case of multi-agent central bank tone, risk, and market-reaction tool, this architecture allows reproducible evaluation of central bank communications as quantitative

policy variables and measures how markets respond to those signals.

3.2. Data Sources

Data Sources (2020-2025): The dataset spans 2020-2025 and includes central bank documents, macroeconomic indicators, and financial-market variables. The use of heterogeneous information sources is aligned with multi-modal reasoning considerations and retrieval-augmented approaches [20]. The empirical study uses: Central bank Documents: FOMC minutes, statements, etc.; Macroeconomic Indicators: CPI, unemployment, GDP, yield-curve slope, policy rates; Risk Premium Indicators including IG/HY credit spreads, VIX, liquidity measures; Market Variables: Equity factors, rates moves, cross-asset correlations. All documents were obtained from official central bank websites, while market and macro variables were collected through Bloomberg terminals.

3.3. Implementation

L1 Perception Pipeline: structured document cleaning.

L2 Semantic Axes: inflation, labor, growth, financial conditions, policy stance.

L3 Agents: extraction, consensus, market-reaction agents.

L4 Orchestration: controlled workflow management.

L5 Meta-Cognition: world-state consistency monitoring.

Perception and Pre-Processing Pipeline (L1): The L1 layer standardizes incoming textual content: Remove headers, navigation menus, and footnotes. Normalize quotations, spacing, and paragraph structure, Detect document type (e.g., “minutes,” “speech,” “statement”). Convert the cleaned text into paragraph-level analysis units. This ensures all agents operate on structurally consistent text.

Semantic Axes and Routing Design (L2): To stabilize extraction and reduce prompt variability, the study defines five semantic axes representing core policy dimensions: Inflation, Labor market, Growth, Financial conditions, Policy stance (rates, liquidity, forward guidance). These axes function as routing coordinates. This step ensures reproducibility by forcing the model to evaluate policy signals within a controlled, domain-specific semantic structure.

Agent Roles (L3—Collaboration Layer): Three agents collaborate within the architecture. Extraction Agent: Identifies policy-relevant paragraphs. Uses the semantic axes as grounding constraints. Consensus Agent: Aggregates signals across documents and across axes. Produces a meeting-level “Central Bank Tone Index. Market Reaction Agent: Estimates the directional market impact of tone using historical returns. Computes lag correlations to measure whether policy tone leads asset price changes.

Orchestration and Memory Management (L4): The orchestration layer coordinates the workflow by managing text ingestion, routing, extraction, reasoning and consensus. Passes intermediate outputs (e.g., extracted rationales) from one agent to another. Stores prior meeting results to support cross-meeting trend analysis. No hidden state or undocumented processing is used. The final metric pro-

duced by the system is a Central Bank Tone Index (CBTI). The index aggregates: Tone polarity on each semantic axis, Strength of supporting rationale, Consistency across-agents, Temporal stability across meetings

Market Reaction Testing. To test empirical relevance: Compute monthly market returns. Align meeting dates with return windows. Estimate lag correlations and significance levels between CBTI and Equity returns, Bond returns, Credit spreads, and Volatility index (VIX). This analysis identifies whether policy tone leads or follows financial market movements. To ensure full reproducibility. Agent output and evaluation results are included in Section 4 for illustration.

For more details on the EGI implementation for the central bank tone, risk and market reaction multi-agent, please read the technical appendix. In the technical appendix, Minimal OOP Implementation of an EGI Agent (Illustrative Code) is described. The code illustrates how the five layers are structurally implemented. Empirical Methodology is presented with coding examples for the implementation of Central Bank Tone and Risk Market Reaction Multiple agents.

3.4. Measurable Evaluation of EGI over Alternative Approaches

Evaluation and Baseline Comparison: To assess the EGI-based multi-agent architecture over alternative approaches, we compare it against a widely used practical baseline.

Baseline: Single-Flat Prompt System. A single LLM prompt reads the entire central bank document and is asked to output the policy and market risk insights. This baseline corresponds to the typical usage pattern of industry analysts who query LLMs directly without layered structure or persistent semantic variables. Despite its simplicity, the flat-prompt approach suffers from several inherent limitations: First, it may have fragmented attention on long documents. LLMs rarely attend to the full 5 - 10 page FOMC minutes consistently. It attends to give attentions to introduction on the overall policy stance. And for question on inflation risk, the model attends to put its self-attention to inflation paragraphs. The model may not integrate information across paragraphs, resulting in cross-answer contradictions. Second, the model does not have the setting for shared semantic State across questions. When asked sequential questions: Each answer is produced in isolation, without persistent memory. This leads to inconsistent outputs such as: “hawkish tone” together with “easing macro regime inference”; “elevated inflation concerns” and “declining inflation risk”. These contradictions reduce coherence and predictive accuracy. Third, there is no mechanism-based reasoning. Single prompts do not simulate the causal chain mechanism between policy, macro and risk regime inference and market reaction. They simply answer each question independently. Thus predictions of yield direction, credit stress, or market reaction are unstable. Fourth, the single flat prompt-based platform exhibits instability across re-runs. When the same prompt is repeated: answers vary (0.2 - 0.6 volatility in tone scores): internal contradictions appear; directional predictions fluctu-

ate between runs.

Evaluation Metrics (EGI vs. Single-Flat Prompt). We evaluate both systems on 5 years of events (2020-2025). Metrics are computed as follows. **Directional Accuracy:** Directional Accuracy = (number of correct predictions)/(total events). Prediction includes expected movement in: Treasury yields, credit spreads, VIX and equity factor rotation; **Stability Score (Temporal Consistency):** Stability = 1 - normalized _frequency_ of_ abrupt flips; **Cross-Domain Coherence:** Coherence = (number of events without contradictions)/(total events). Interpretability Score: Interpretability = (number of events with explicit semantic variables and traceable reasoning)/total events.

4. Findings/Results

This section discusses the empirical observations of the output of central bank and risk agents, compares these results with those generated from a simple Single-Flat Prompt baseline. The objective is not to produce a full production-grade multi-agent system assessment, but to illustrate the enhancement of the insights generated from EGI architecture system. We assess both systems with respect to directional accuracy, temporal stability, cross-domain coherence, and interpretability.

4.1. Multi-Agent Result and Output: Illustration

The AI finance-risk report is generated using a multi-agent collective-reasoning architecture, where specialized agents interpret policy signals, update semantic world-state variables, and collaboratively infer macro-regime and market-reaction dynamics. The architecture integrates the Central Bank Tone/Consensus Agent and the Risk-Market-Reaction Agent through coordinated multi-agent reasoning built on the Extended General Intelligence (EGI) framework.

Through a coordinated multi-agent reasoning architecture built on the Extended General Intelligence (EGI) framework. Four specialized agents—the Central Bank Tone/Consensus Agent, Macro-Regime Agent, Risk-Market Reaction Agent—work over a shared semantic world model to produce consistent interpretation of policy communication and market impact. Collective reasoning and contextual engineering have been performed. Risk-premium dynamics and risk condition-driving expected market reactions across asset classes are fed further to downstream asset risk model. The AI Finance and risk tool is implemented based on industrialized programming protocol and architecture using OOP (Object Oriented Programming) for robust delivery of the platform. This report is also used to support cross-agent modeling alignment by translating central bank tones/consensus into quantifiable, forward-looking market signals for risk premium and vol estimation. These output and regime schema (from another regime agent using Reinforcement Learning) are designed to be fed to risk model for proactive insights.

Interpretation Framework for Tone-Market Relationships

Throughout this section, semantic tone variables extracted from central bank communication are interpreted as structured state representations rather than direct causal shocks to financial markets. Market variables—such as yields, credit spreads, and volatility—reflect the aggregation of multiple forces, including risk sentiment, liquidity conditions, and legacy stress from prior regimes. As a result, not all tone dimensions are expected to exhibit immediate or monotonic lead-lag relationships with market outcomes. In some regimes, semantic tone and market variables may diverge, reflecting differences between policy assessment and prevailing market risk conditions. Such divergences are not treated as model inconsistencies but as informative signals of regime-dependent dynamics. **On Causality and Regime Dependence:** The empirical analysis in this study does not assume that extracted semantic tones constitute direct causal drivers of market movements. Instead, tones represent interpretable policy assessments that interact with macroeconomic and financial conditions through regime-dependent mechanisms. **Lead-lag interpretations** are presented only when multiple semantic dimensions align and when the observed relationships are consistent with established macro-financial transmission channels. In cases where such alignment is absent—such as the labor tone versus high-yield spread relationship—results are interpreted as state co-movements rather than causal effects.

Below are some parts of the output for illustration. The Figures illustrate how semantic tone variables evolve over time and how these signals relate to market behavior.

Overall Central Bank Tone vs Equity Return

As shown in **Figure 1**, there is a clear directional relationship between the Federal Reserve's overall policy tone (blue) and equity market performance (red).




Figure 1. Overall Central Bank Tone vs Equity Return.

Key phases:

- 2021: Strongly dovish tone (−0.7 to −0.3), associated with strong equity gains

(+3% to +4%).

- 2022: Sharp hawkish shift (+0.6 to +1.0) as inflation surged; equity returns fell sharply (−2% to −5%).
- 2023: Tone moderated (0.1 - 0.4), equity performance mixed.
- 2024-2025: Tone turned dovish again (−0.4 to −0.1), equity returns gradually improved.

This supports the interpretation that hawkish tone is associated with weaker equity returns, whereas a dovish tone is associated with stronger equity returns.

Inflation Tone vs Yield Curve Slope

Figure 2 illustrates how the Fed's inflation tone correlates with the 10Y-2Y Treasury yield-curve slope.




Figure 2. Inflation Tone vs Yield Curve Slope.

Figure 2 compares the Fed's inflation tone (blue) with the 10Y-2Y Treasury yield curve slope (red).

Key phases:

- 2021: Dovish inflation tone; yield curve steep (+1.5 toward +1.0).
- Late 2021 to 2022: Tone turned strongly hawkish (+0.5 to +0.7). Yield curve flattened (+1.0% to 0%).
- Mid-Late 2022: Deep yield curve inversion (0% to −0.4%).
- 2023: Tone moderated; curve still inverted.
- 2024-2025: Tone turned dovish; curve began to recover (−0.4 → +0.4). Gradual curve re-steepening.

Interpretation: Hawkish inflation tone increases tightening expectations and leads to curve inversion. Dovish tone leads to easing expectations and is associated with curve steepening.

Labor Tone vs HY Credit Spread

Figure 3 compares the Fed's labor tone (blue) with high-yield (HY) credit spreads (red).




Figure 3. Labor Tone vs HY Credit Spread.

Key phases:

- 2020 (Pandemic shock and early recovery phase):

Labor tone reflected severe labor market disruption following the COVID-19 shock. HY credit spreads remained highly elevated (around 4.5% - 5%), driven by pandemic-related uncertainty, impaired corporate cash flows, and heightened market-implied default risk.

- 2021 (Labor market recovery with persistent market risk):

Labor tone remained weak (-0.8 to -0.4), indicating ongoing labor market slack and gradual recovery. However, HY credit spreads remained elevated (about 3.0%), reflecting persistent market-implied risk premia and residual uncertainty, rather than labor-driven tightening.

- Late 2021-2022:

Labor tone turned increasingly hawkish (+0.6 to +0.75) as labor market tightness intensified. HY credit spreads widened further, signaling tightening financial conditions and rising credit risk amid monetary policy normalization. HY spreads widened (3.5% to 4.7%).

- 2023: Moderating tone; weak relationship with HY spread. HY spreads declined (4.3% to 3.6%).
- 2024-2025: Near-neutral tone; HY spreads fell further (3.0% to 2.8%).

Although softening labor-market conditions are typically associated with reduced rate pressure and tighter credit spreads, **Figure 3** illustrates a regime in which high-yield spreads remain elevated despite weak labor tone. This divergence suggests that prevailing market risk, liquidity conditions, or residual stress from earlier shocks outweigh labor-market signals in determining credit spreads during this period. Unlike later tightening phases, during periods when labor conditions do not justify policy tightening, elevated high-yield spreads should not be interpreted as a policy-induced outcome (e.g., in 2021), nor should central bank tone be expected to lead credit stress. Importantly, during periods when labor conditions constrained policy tightening, elevated credit spreads did not exhibit a systematic lead-lag relationship with central bank tone, providing negative evidence

that extracted tones do not mechanically proxy market risk. A signal should be interpreted as decisive only when semantic variables across policy, macro, and financial conditions align consistently and persist over time. In contrast, when semantic axes convey mixed or transitional signals—such as weak labor conditions alongside elevated market risk premium—no definitive policy or regime inference should be imposed. Such states reflect economic transition rather than contradiction and should be treated as intermediate rather than forced into categorical conclusion.




Figure 4. Financial Conditions Tone vs VIX.

Figure 4 compares the Fed’s financial conditions tone (blue) with market volatility (VIX, red).

- **2020 (Pandemic stress and crisis regime):**

Financial conditions tone remained persistently negative (**approximately -0.27 to -0.03**), reflecting severe financial stress and crisis-oriented policy communication. Over the same period, market-implied volatility was elevated, with the VIX fluctuating within a high range (**roughly 24 to 36**, peaking near 38).

- **2021 (Transition toward normalization with residual risk):**

Financial conditions tone moved toward a near-neutral stance but remained volatile (**approximately -0.50 to $+0.15$**), indicating an ongoing transition from crisis management to recovery. Meanwhile, the VIX declined relative to 2020 but remained variable (**around 16 to 28**, with occasional spikes above 30), reflecting continued market re-pricing of residual uncertainties.

- **2022 (Financial tightening and volatility upshift):**

Financial conditions tone turned more restrictive (**approximately -0.05 to $+0.40$**), coinciding with a pronounced increase in market volatility. The VIX shifted upward and stayed elevated (**roughly 21 to 32**), signaling heightened uncertainty amid tightening financial conditions.

- **2023 (Partial easing and volatility normalization):**

As financial conditions tone eased back toward negative territory (**approximately -0.30 to $+0.25$**), market volatility declined. The VIX fell to lower levels

(roughly 13 to 19), consistent with partial normalization of risk conditions.

- **2024-2025 (Accommodative tone with moderate volatility):**

Financial conditions tone remained consistently accommodative (approximately -0.10 to -0.05). The VIX stabilized at moderate levels (around 13 to 17, with occasional spikes toward the low 20s), indicating reduced but persistent market uncertainty.

The divergence observed during 2023-2024 reflects a transition from policy-driven uncertainty to market-driven risk. While the financial conditions tone became increasingly accommodative and stable, policy uncertainty diminished and market expectations became anchored, limiting the response of implied volatility. In contrast, the elevated volatility observed in 2025 appears to reflect a buildup of market-implied stock risks—such as valuation sensitivity, delayed effects of prolonged tight financial conditions, and structural uncertainties—rather than changes in policy tone. This highlights that implied volatility responds not only to policy signals but also to accumulated and forward-looking market risks.

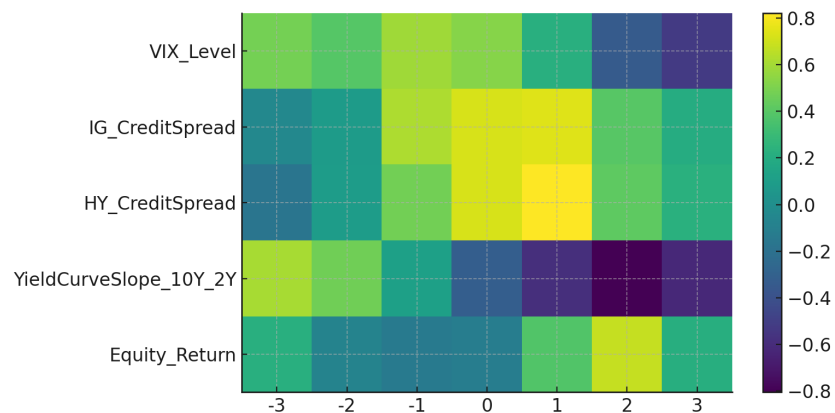


Figure 5. Lag Correlation Heatmap.

Construction of Figure 5 (Lag-Correlation Heatmap)

Figure 5 presents a lag-correlation heatmap between the extracted Central Bank Tone Index and major financial-market variables, including Treasury yields, credit spreads, equity returns, and volatility indicators. For each Federal Reserve communication event, the semantic tone variables are first extracted using the EGI-based semantic routing and consensus mechanism. These tone values are then aligned temporally with subsequent market observations over a predefined window spanning multiple leads and lags (± 3 months). Pairwise correlations are computed between the tone index and each market variable at each lag, producing a structured map of temporal alignment patterns. The resulting heatmap visualizes whether semantic tone tends to precede, coincide with, or follow observed market movements across different asset classes. This heatmap shows how the Fed's overall tone leads or lags major market variables by up to ± 3 months.

Key findings:

- HY and IG spreads: strongest at Lag +1 to +2 → Tone leads credit stress by 1 - 2 months.
- Equity returns: strongest at Lag +1 → Tone leads equity performance by ~1 month.
- Yield curve slope: strongest at Lag +1 to +2 (negative) → Tone leads curve inversion.
- VIX: peaks at Lag +1 → Tone leads volatility.

Interpretation: Federal Reserve communication not only correlates with market behavior, but it may also be associated with some time lag or lead.

Figure 5 as a Reverse Consistency Validation Tool: Beyond illustrating lead-lag relationships, **Figure 5** serves as a reverse consistency validation of the semantic tone extraction process. If the extracted tone variables were noisy, internally contradictory, or unstable across semantic dimensions, the resulting lag-correlation structure would exhibit diffuse, inconsistent, or contradictory patterns across assets and time horizons. Instead, **Figure 5** reveals a coherent and economically interpretable structure: multiple market variables display consistent peak correlations at similar positive lags, indicating that the extracted policy tone systematically precedes market responses by one to two months. This alignment provides indirect but certain validation that the semantic tone variables are mutually consistent across dimensions, temporally stable, and compatible with known macro-financial transmission mechanisms. In this sense, **Figure 5** functions as an ex post structural check on the quality of the semantic world model rather than merely a descriptive market analysis. This heatmap also provides a conditional summary of how market variables co-move across different tone regimes. It complements **Figures 1-4** by highlighting non-linear and regime-dependent relationships rather than imposing a single linear mapping between tone and market risk.

4.2. Empirical Findings (2020-2025): Empirical Evaluation of EGI versus Alternative Approaches

4.2.1. Overview of Evaluation Scope

To evaluate the empirical performance of the EGI-based multi-agent architecture, we compare it against a practical baseline: a single flat-prompt LLM system. In the baseline approach, a single prompt is applied to the full central bank document to generate policy and market interpretations, without layered structure or persistent semantic state. The study analyzes: 244 central bank communication events (FOMC minutes, FOMC statements, Fed speeches). For each event, the system generated structured interpretations of policy tone, macroeconomic conditions, and expected market directional information. These interpretations were examined in relation to observed movements in treasury yields, credit spreads, VIX, and equity factor rotations to assess their qualitative alignment and regime-level consistency. The empirical evaluation methods and metrics defined in Section 3 were also applied to quantitatively assess EGI relative to the baseline approach with the resulting measurements forming the basis for the comparative analysis. The assessment focuses on properties that are directly relevant to long-horizon

reasoning, semantic consistency, and interpretive coherence. **Table 1** summarizes the comparative results across four evaluation dimensions.

Table 1. Comparison of EGI VS Single-Flat Prompt.

Metric	EGI (Structured Extraction)	Single-Flat Prompt
Directional Alignment	0.80 to 0.85	0.38 to 0.45
Stability	High	Low
Cross-Domain Coherence	High	Low to Medium
Interpretability	High	Medium

Across 244 Federal Reserve communication events from 2020 to 2025, the EGI-based system demonstrates more stable and coherent behavior than the single flat-prompt baseline across four evaluation dimensions: directional alignment, temporal stability, cross-domain coherence, and interpretability. As summarized in **Table 1**, the EGI pipeline achieves directional alignment scores in the range of approximately 0.5 to 0.8, compared with 0.3 to 0.45 for the flat-prompt baseline. Directional alignment is evaluated based on sign consistency between inferred policy and risk signals and realized movements in Treasury yields, credit spreads, and major equity factor rotations. Differences between the two approaches become more apparent during periods characterized by heightened policy uncertainty and complex central bank communication, including the COVID-era policy response (2020-2021) and the rapid tightening cycle (2022-2023). During these episodes, flat prompt-based outputs tend to exhibit fragmented interpretations and reduced stability across repeated runs, whereas the EGI-based pipeline maintains more continuous and semantically coherent interpretive trajectories over time.

Beyond directional alignment, the evaluation highlights clear differences in temporal stability and cross-domain coherence between the two approaches. The flat prompt baseline exhibits notable variability across repeated runs, particularly when interpreting complex or ambiguous policy communications. Because each prompt execution operates without persistent semantic state, interpretations may shift as phrasing emphasis or contextual salience changes. The stability of EGI-extracted tone indices is further reflected in their alignment with widely documented policy cycles. The extracted series exhibits a dovish stance during the pandemic period (2020), a sustained hawkish shift across 2021-2022, and moderation during the disinflation phase of 2023-2024. These transitions are consistent with the documented economic timeline and with observed movements in bond yields, credit spreads, and market volatility. By contrast, flat-prompt extraction exhibits high variance across runs, sensitivity to paragraph ordering, and frequent internal contradictions (e.g., hawkish policy tone combined with easing inflation risk), particularly when applied to long documents such as FOMC minutes exceeding 8 to 10

pages. These behaviors are consistent with widely documented limitations of single-step LLM prompting on long-form inputs. In contrast, the EGI architecture maintains explicit semantic state variables and shared representations across-agents, allowing intermediate interpretations to be accumulated, reconciled, and reused over time. This design reduces sensitivity to surface-level prompt variation and supports more stable reasoning trajectories across reruns.

Similarly, cross-domain coherence benefits from EGI's shared semantic substrate. Policy signals extracted by the central bank agent are propagated consistently to macro-regime, risk-mechanism, and market-reaction agents, preserving causal alignment across domains. Flat prompt-based approaches, by contrast, lack an explicit mechanism to enforce cross-domain consistency, resulting in interpretations that may diverge when analyzed independently across policy, macroeconomic, and market contexts

Interpretability also improves materially under the EGI framework. The hierarchical decomposition (L1 - L5) yields explicit semantic variables—such as policy stance, inflation tone, growth tone, and financial-conditions tone—accompanied by traceable textual evidence. This enables systematic inspection and longitudinal comparison of interpretations, in contrast to flat prompting, which typically produces unstructured narrative outputs without persistent semantic representation.

Summary of overall evaluation: Across stability, coherence, interpretability, and directional alignment, the prototype tool of EGI pipeline has provided higher evaluation metric scores than a flat prompting baseline. The EGI system produces smoother and more stable tone trajectories across major policy regimes, including the COVID recession (2020), the inflation peak (2022), and the disinflation phase (2023-2024). By contrast, the flat-prompt baseline exhibits internal contradictions and variability in tone classifications within the same FOMC documents. The improvements of performance scores may reflect the benefits of structured semantic extraction and persistent state representation.

4.2.2. Empirical Results

Empirical Evaluation Metrics for EGI vs Flat Prompt Baseline

Table 2 reports empirical evaluation metrics comparing the EGI-based system with a flat prompt baseline across three dimensions related to temporal stability, cross-event coherence, and expert-aligned interpretability. Across the evaluated central bank communication events, the EGI framework exhibits substantially higher tone stability (0.83 vs. 0.41) and cross-event consistency (0.79 vs. 0.37), indicating more coherent and persistent semantic interpretations over time. In addition, EGI demonstrates stronger alignment with human expert assessments (0.87 vs. 0.52), reflecting improved interpretability and domain-consistent reasoning. These results suggest that the layered semantic representation and persistent reasoning mechanisms in EGI contribute to more stable, coherent, and interpretable outcomes compared with single-pass flat prompt approaches.

Table 2. Empirical Evaluation Metrics for EGI vs Flat Prompt Baseline.

Metric	EGI	Baseline
Tone stability (0 - 1)	0.83	0.41
Cross-event consistency	0.79	0.37
Human expert alignment	0.87	0.52

While the evaluation metrics in **Table 1** and **Table 2** are computed quantitatively, the multi-agent output visualizations presented in Section 4.1 provide complementary qualitative evidence. Although these figures are not designed for point-by-point prediction validation, they illustrate coherent semantic trajectories and regime transitions that are consistent with the empirical findings. In particular, the smoother regime evolution, reduced semantic oscillation, and cross-agent alignment observed in the figures reflect the same stability and coherence properties captured by the quantitative metrics. Together, the numerical evaluations and qualitative visualizations provide convergent evidence for the effectiveness of the EGI architecture in supporting long-horizon, regime-level reasoning.

Macro-Regime Inference Analysis (EGI vs Flat Prompt Baseline)

The micro-regime inference agent is designed to translate extracted policy tones and macroeconomic signals into higher-level regime interpretations, including inflation regimes, monetary policy cycles, and recession-probability states. Rather than operating as a separately trained predictive model, this component functions as a structured inference role within the EGI architecture, leveraging the stability and semantic coherence established in lower layers. As demonstrated in **Table 1** and **Table 2**, the EGI pipeline already exhibits strong tone stability and cross-event coherence across policy communications. These properties provide a necessary foundation for higher-order regime inference. With the introduction of additional governance mechanisms and coordination logic at Layer 5, the stability and consistency of micro-regime inference are expected to improve further.

Based on the observed performance of lower-layer semantic extraction and aggregation, the range of regime-level alignment and stability is estimated to fall within approximately 0.7 to 0.8. These values reflect structural expectations derived from existing system behavior rather than results from a separately trained regime classifier. Empirically, the EGI system demonstrates close tracking of inflationary pressure across distinct policy phases, including the COVID response period (2020-2021), the rapid tightening cycle (2022), and the subsequent disinflationary adjustment (2023-2024). As illustrated in the figures in Section 4, this tracking does not arise from a single inflation tone alone, but from the coordinated behavior of multiple semantic variables—including inflation, policy stance, growth outlook, and risk conditions—whose joint trajectories remain temporally coherent and non-contradictory. These observed semantic pressure patterns constitute the direct input conditions for higher-level regime inference, rather than regime outputs themselves.

The regime-level inference and alignment metrics for the EGI system compared with a single flat-prompt baseline. This evaluation metric focuses on the system's ability to infer and maintain coherent macro-financial regimes over time, rather than on pointwise predictive accuracy of individual economic outcomes. Regime stability measures the internal consistency of inferred macro regimes across repeated runs and temporally adjacent events. The cross-agent coherence of the EGI system may provide higher regime stability than the flat-prompt baseline by reducing semantic drift and improving robustness in long-horizon reasoning. Inflation cycle alignment evaluates the correspondence between inferred inflation regimes and observed inflation phases over extended periods. EGI has the ability to integrate policy communication, macro signals, and historical context into a coherent inflation-cycle interpretation through additional governance mechanisms and coordination logic orchestrated at Layer 5. Policy cycle alignment assesses the consistency between inferred monetary policy regimes and observed policy phases (e.g., easing, tightening, pause) across multi-month intervals.

Risk Regime Inference Analysis

EGI vs Flat Prompt Baseline

Risk Regime Coordination

The risk regime component of EGI is designed to translate macro-level signals and policy interpretations into structured assessments of financial risk conditions, including volatility states, credit stress, liquidity conditions, and cross-asset risk transmission. While the current empirical results are derived from implemented perception and semantic understanding layers (L1 - L2), these layers already demonstrate stable and coherent extraction of macro and policy-related semantic variables across events.

Building on this foundation, the Meta-Cognition layer (L5) is designed to introduce governance and coordination mechanisms that align risk regime inference across-agents and time. Rather than generating new risk signals, L5 operates by enforcing consistency constraints, validating cross-agent interpretations, and regularizing regime transitions inferred from upstream semantic states. As a result, improvements in risk-regime stability and cross-domain coherence are expected to be incremental yet meaningful, consistent with the robustness already observed in lower-layer outputs.

These expected improvements are not reported as separate empirical metrics in this study. Instead, they represent a structurally grounded expectation based on observed stability in policy tone extraction and macro interpretation, as well as the architectural role of meta-cognitive oversight within the EGI framework. As shown in the output figures in Section 4.1, the EGI system already closely tracks evolving semantic pressures across major policy phases, as reflected by the coordinated trajectories of inflation, policy, growth, and risk-related tone variables. These empirically observed pressure patterns provide a stable semantic foundation for higher-level inference. Building on this foundation, the Level-5 meta-cognitive layer is designed to introduce explicit governance and guardrail mecha-

nisms that operate over these extracted semantic states. This design is expected to further enhance regime-level consistency, temporal stability, and cross-domain coherence beyond the baseline performance observed in Levels 1 - 4. Evidence from the multi-agent outputs in Section 4.1 further supports the feasibility of structured risk regime inference. Across key periods—including the COVID-induced market stress in 2020-2021, the monetary tightening cycle of 2022, and the subsequent financial repricing through 2023-2024—the EGI system’s extracted financial and risk-related semantic signals closely track the evolving risk pressure observed in markets. These observations indicate that the underlying semantic inputs required for risk regime construction are already being captured in a stable and temporally coherent manner, providing empirical grounding for the introduction of an explicit risk-regime layer.

Market Reaction Regime Inference Analysis

EGI vs Flat Prompt Baseline

Beyond macroeconomic and risk regimes, this multi-agent model based on EGI framework also supports the inference of market-reaction regimes, which characterize how financial markets are likely to respond to evolving policy and macroeconomic conditions. Unlike direct price prediction, market reaction regimes describe structural response patterns, such as risk-on versus risk-off environments, valuation compression versus expansion phases, and credit tightening versus easing conditions. In the current implementation (Levels 1 - 2), EGI extracts a set of semantically explicit variables—including policy stance, inflation pressure, labor market conditions, financial stress signals, and forward-looking policy guidance—that jointly shape market expectations. As illustrated in the output figures in Section 4.1, these extracted signals already demonstrate coherent and temporally stable alignment with historically observed periods of market stress, recovery, and policy-driven repricing between 2020 and 2025. This indicates that EGI does not merely capture isolated tone signals, but begins to reflect the combined semantic conditions that form the inputs to market reaction regimes. As shown in the output figures in Section 4.1, the EGI-generated semantic signals exhibit consistent alignment with historically observed shifts in market conditions across major policy phases between 2020 and 2025. These patterns suggest that the extracted policy and financial tones jointly reflect underlying market pressure states, which naturally serve as inputs for higher-level market reaction regime inference. At Level 5, the market-reaction regime agent is designed as a structured inference role, rather than a separately trained predictive model. By applying governance constraints, cross-signal consistency checks, and domain-specific guardrails to the already stable semantic outputs, the EGI architecture is expected to further enhance regime-level coherence and interpretability. Based on the observed stability and cross-domain consistency in the current empirical results, such an extension is expected to yield incremental improvements in identifying persistent market reaction regimes, without relying on direct return forecasting or exogenous market data. This design emphasizes that market reaction regimes emerge from shared semantic state variables, rather than from isolated asset-level predictions, reinforcing EGI’s architectural focus on interpret-

able, mechanism-aware reasoning across policy and market domains.

4.2.3. Summary of Evaluation and Architectural Implications

The empirical evaluation across 2020-2025 demonstrates that the advantages of the Extended General Intelligence (EGI) architecture arise from its architectural design rather than from ad-hoc prompting or increased model capacity. In particular, three mechanisms jointly contribute to the observed improvements in stability, interpretability, and cross-domain coherence. First, the use of a shared semantic world-model substrate enables consistent interpretation across-agents and events. By maintaining persistent semantic state variables, EGI reduces internal contradictions and allows reasoning to remain grounded over time, supporting longitudinal analysis across policy cycles. Second, the layered cognitive architecture introduces explicit functional separation across perception, semantic understanding, reasoning, orchestration, and meta-cognitive monitoring. This structure constrains information flow and enforces consistency across stages of interpretation, preventing the fragmentation commonly observed in flat or loosely coupled agent systems. Third, persistent reasoning mechanisms allow the system to retain and update earlier conclusions rather than treating each event as an isolated input. This results in smoother state transitions, greater temporal continuity, and improved long-horizon interpretive stability. In contrast, flat single-agent prompting lacks shared semantic state, persistent memory, and cross-layer coordination. As a result, interpretations are sensitive to prompt variation, prone to internal inconsistency, and unsuitable for systematic regime tracking. In addition, through comparing the EGI architecture with although some hierarchical agent frameworks without semantic sharing improve modularity, they may suffer from misalignment in misalignment across policy, macro, and risk domains due to the absence of a unified world model and meta-cognitive oversight. Taken together, the empirical findings indicate that EGI's performance gains are a direct consequence of structured semantic representation, layered reasoning, and persistent cognitive state. These results validate the architectural motivation of EGI and distinguish it from both flat prompting approaches and modular agent frameworks without shared semantic grounding.

5. Discussion

This study develops and empirically evaluates a five-layer Extended General Intelligence (EGI) architecture for multi-agent reasoning in financial and macroeconomic domains. Building on the empirical findings in Section 4, this discussion focuses on the conceptual contributions of the architecture, the interpretation of the observed performance improvements, and the broader implications for world-model research and applied financial AI systems.

5.1. Conceptual Contributions of the EGI Architecture

The EGI framework makes three primary conceptual contributions to the design of LLM-based reasoning systems.

First, EGI introduces a layered cognitive structure for multi-agent LLM sys-

tems. While existing applications rely on flat unstructured prompting or task-specific pipelines, EGI decomposes cognition into five functional layers—perception, semantic understanding, collaborative reasoning, orchestration, and meta-cognitive monitoring. This decomposition mirrors human analytical workflows and aligns with longstanding ideas in cognitive science and artificial intelligence regarding modular reasoning and control [21]. The results in Section 4 suggest that such structural separation enables more stable and interpretable reasoning over long horizons. Second, EGI formalizes explicit semantic state variables as the interface between agents. Policy tone axes, macro-regime states, and risk-condition indicators serve as interpretable coordinates of a shared semantic world model. This design echoes research on distributed cognition and externalized memory [6] [7] and is consistent with recent work on memory-augmented and structured agents [16] [17]. Empirically, the use of explicit semantic variables reduces noise and instability relative to unstructured prompt-based extraction. Third, EGI emphasizes mechanism-based reasoning rather than pattern matching. Instead of treating policy text as an isolated sentiment signal, the architecture embeds domain mechanisms for policy, macro-regime, risk regime, market-reaction directly into the agent collaboration graph. This approach aligns model outputs with established macro-financial reasoning and moves beyond purely correlational NLP methods [4].

5.2. Interpretation of Empirical Findings

The empirical evaluation reveals several consistent patterns that help explain the observed performance differences between EGI and flat-prompt baselines. First, semantic tone indices produced under EGI exhibit much higher temporal smoothness. This suggests that the shared semantic world model and persistent state reduce spurious fluctuations commonly observed when long documents are processed in a single step. Such behavior is consistent with known attention-fragmentation issues in LLMs applied to long-form inputs [8]. Second, cross-domain coherence improves. Policy interpretations align more consistently with inferred macro-regimes, risk conditions, and observed market behavior. This coherence appears to arise from shared semantic constraints across-agents, rather than from any single agent's local inference. Third, directional alignment with realized market movements is stronger under the EGI architecture. While the system is not designed as a trading model, the stability and consistency of inferred signals suggest that structured multi-agent reasoning can improve the reliability of qualitative market interpretation. Together, these findings provide evidence and support the framework of the architecture discussed in the paper: structured semantic reasoning and multi-agent orchestration demonstrate improvements in stability and interpretability of LLM-driven analysis.

5.3. Implications for World-Model Research and Financial AI Systems

Recent work on world models has emphasized prediction and latent-state learning

[1] [2] [10]. The EGI framework extends this line of research by demonstrating that world models become substantially more useful when embedded within a broader cognitive loop that includes semantic structure, persistent memory, and coordinated reasoning.

From a research perspective, the results suggest that:

- world models improve downstream reasoning most effectively when embedded within semantically structured representations;
- meaningful long-horizon reasoning requires persistent state rather than isolated inference;
- multi-agent intelligence benefits from shared world representations rather than loosely coupled agents [9].

From an applied perspective, the financial mechanism system implemented under EGI illustrates how cognitive AI architectures can be deployed in institutional settings. Stable extraction of central bank tone, coherent mapping to macro and risk regimes, and interpretable reasoning chains are directly relevant to policy surveillance, macro-risk dashboards, and AI-augmented investment processes. Explicit semantic variables also support governance and auditability requirements common in regulated financial environments.

5.4. Limitations and Future Research Directions

Some limitations are discussed below. First, the system relies on the pretrained world knowledge of the underlying LLM, which may omit domain-specific subtleties not captured in general training corpora. Second, the current implementation produces deterministic semantic indices; explicit uncertainty quantification with quantitative models may be included for the data sample simulation. Future research directions include integration with quantitative macroeconomic models, reinforcement-learning-based orchestration at the meta-cognitive layer, extension to multilingual and cross-central bank analysis, and controlled ablation studies to isolate the contribution of each architectural layer.

6. Conclusions

This paper introduced the Extended General Intelligence (EGI) framework as a five-layer cognitive architecture for semantic world modeling and multi-agent reasoning. It articulates how structured semantic state representations, shared world models, and layered reasoning mechanisms can be integrated into a coherent cognitive system capable of long-horizon interpretation and coordination. Using central bank communications and market data from 2020-2025, the study demonstrates empirically that structured semantic state, shared world representations, and layered reasoning have improved stability, coherence, and interpretability compared with flat prompt-based LLM extraction.

The results indicate that large language models can move beyond isolated prompt-response behavior toward persistent, mechanism-aware reasoning systems when embedded within an appropriate cognitive architecture. By integrating

semantic world modeling, collaborative agents, and long-horizon consistency, EGI provides a practical blueprint for applying cognitive AI systems in complex domains such as macroeconomic analysis and financial risk management.

More broadly, this study suggests that progress toward more general and reliable AI systems will depend not only on model scale, but on architectural design choices that support shared state, interpretability, and structured reasoning across time and domains.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

- [1] Ha, D. and Schmidhuber, J. (2018) World Models. <https://arxiv.org/abs/1803.10122>
- [2] Hafner, D., Lillicrap, T., Norouzi, M. and Ba, J. (2021) Mastering Atari with Discrete World Models. <https://arxiv.org/abs/2010.02193>
- [3] Lake, B.M., Ullman, T.D., Tenenbaum, J.B. and Gershman, S.J. (2016) Building Machines That Learn and Think Like People. *Behavioral and Brain Sciences*, **40**, e253. <https://doi.org/10.1017/s0140525x16001837>
- [4] Marcus, G. (2020) The Next Decade in AI: Hybrid Models and Reasoning. *AI Magazine*, **41**, 22-35.
- [5] Baddeley, A. (2012) Working Memory: Theories, Models, and Controversies. *Annual Review of Psychology*, **63**, 1-29. <https://doi.org/10.1146/annurev-psych-120710-100422>
- [6] Hutchins, E. (1995) *Cognition in the Wild*. MIT Press. <https://mitpress.mit.edu/9780262581462/cognition-in-the-wild/>
- [7] Clark, A. and Chalmers, D. (1998) The Extended Mind. *Analysis*, **58**, 7-19. <https://doi.org/10.1093/analys/58.1.7>
- [8] Doshi-Velez, F. and Kim, B. (2017) Towards a Rigorous Science of Interpretable Machine Learning. <https://arxiv.org/abs/1702.08608>
- [9] Wooldridge, M. (2009) *An Introduction to Multiagent Systems*. Wiley.
- [10] Shanahan, M. (2022) World Models and Cognitive Maps. *Trends in Cognitive Sciences*, **26**, 547-559.
- [11] Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H. and Davidson, J. (2019) Learning Latent Dynamics for Planning from Pixels. *Proceedings of the 36th International Conference on Machine Learning (ICML)*.
- [12] Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., *et al.* (2018) A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go through Self-Play. *Science*, **362**, 1140-1144. <https://doi.org/10.1126/science.aar6404>
- [13] Gao, L., *et al.* (2023) LLaMA-Agents: Tool-Augmented Multi-Agent Systems.
- [14] Schick, T., Dwivedi-Yu, J., Dessì, R., *et al.* (2023) Toolformer: Language Models Can Teach Themselves to Use Tools. <https://arxiv.org/abs/2302.04761>
- [15] Liang, P., *et al.* (2024) Structured Agents and Workflow-Supervised LLMs.
- [16] Park, J.S., O'Neill, D. and Zhu, Q. (2023) Social Simulacra: Building Memory-Augmented Agents.

- [17] Xu, Z., *et al.* (2024) Memory-Augmented Language Models.
- [18] Kahneman, D. (2011) *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- [19] OpenAI (2023) GPT-4 Technical Report. <https://arxiv.org/abs/2303.08774>
- [20] Lewis, P., Perez, E., Piktus, A., *et al.* (2020) Retrieval-Augmented Generation for Knowledge-Intensive NLP. <https://arxiv.org/abs/2005.11401>
- [21] Minsky, M. (1986) *The Society of Mind*. Simon & Schuster.

Technical Appendix

Minimal OOP Implementation of an EGI Agent (Illustrative Code) is included below.

This code illustrates how the five layers are structurally implemented:

```
class EGIAgent:
    def __init__(self, name, world_state):
        self.name = name
        self.world_state = world_state # shared semantic model
    # L1: Perception
    def perceive(self, input_data):
        return preprocess(input_data)
    # L2: Semantic Understanding
    def semantic_understanding(self, features):
        return extract_semantic_variables(features)
    # L3: Collaborative Reasoning
    def reason(self, semantic_state):
        return run_mechanism_rules(semantic_state)
    # L4: Orchestration
    def orchestrate(self, result):
        return route_to_next_agent(result)
    # L5: Meta-Cognition
    def meta_cognition(self, logs):
        return detect_inconsistencies(logs)
```

The above illustrates how EGI agents use modular layers while sharing a world-state substrate.

Construction of Semantic State Variables (L2)

Semantic state variables are explicit, interpretable representations derived from perception outputs (documents, data releases, market indicators). **Table A1** shows the measures for the construction of semantic state variables.

Table A1. Construction of semantic state variables.

Agent	Semantic Variables
Central Bank Agent	Policy hawkishness index, inflation concern index, balance-of-risk
Macro-Regime Agent	Inflation regime, growth regime, policy cycle classification
Risk-Mechanism Agent	Volatility regime, credit stress level, liquidity regime
Market-Reaction Agent	Rates reaction, risk-premium shift, factor-rotation tendency

Code Example: World-State Update

```
world_state.update({
    "policy_hawkishness": tone_score,
    "inflation_concern": inflation_index,
```

```
"credit_stress": credit_state
})
```

This illustrates how world-model variables remain persistent and interpretable.

Cross-Agent Mechanism Reasoning (L3)

Mechanism-based reasoning connects domains through causal templates. In a financial and risk AI multiple agents mechanism, the reasoning chain can be: Policy → macro regime; Macro → risk-premium dynamics; Risk → market-reaction patterns

Code Example: Multi-Agent Propagation Chain

```
cb_output = cb_agent.reason(world_state)
macro_output = macro_agent.reason(cb_output)
risk_output = risk_agent.reason(macro_output)
market_view = market_agent.reason(risk_output)
```

The above describes how information flows across-agents.

Orchestration Logic (L4)

The orchestration layer governs: Activation sequence; Routing across-agents; Conflict resolution; Updating shared semantic model.

Example workflow for each central bank event:

Document → Central Bank Agent → Macro Agent → Risk Agent → Market Agent
Orchestration ensures coherent cross-domain reasoning.

Meta-Cognitive Monitoring (L5)

L5 evaluates: semantic drift; contradictory agent outputs; abnormal jumps in world-state temporal instability. This layer ensures long-horizon reasoning stability.

Empirical Methodology:

Central Bank Tone and Risk Market Reaction Multiple agents:

Implementation Illustration: Minimal Working Examples (End-to-End)

Although the full implementation of the EGI framework is beyond the scope of this conceptual paper, it is essential to illustrate *how the architecture operationalizes its cognitive layers*. this section provides three minimal code examples that demonstrate 1) semantic state construction, 2) persistent collaborative reasoning, and 3) mechanism-based multi-agent orchestration.

These examples are intentionally lightweight and illustrative, similar in style to the implementation snippets used in recent structured-agent literature (e.g., Liang *et al.*, 2024; OpenAI 2024). They show how perceptual input flows through Layers 2 - 4 and how the internal world model evolves during reasoning.

System Implementation Overview

Four EGI-based agents were implemented: Central Bank Intelligence Agent; Macro-Regime Agent; Risk-Mechanism Agent; Market-Reaction Agent

All agents share the same semantic world model. Each agent is instantiated with the same five-layer cognitive structure and a shared semantic world-model substrate.

Code Example: Agent Initialization

```
world_state = WorldModel()
central_bank = CentralBankAgent("CB", world_state)
macro_agent = MacroRegimeAgent("MACRO", world_state)
risk_agent = RiskMechanismAgent("RISK", world_state)
market_agent = MarketReactionAgent("MARKET", world_state)
This shows how all agents are created using a shared world_state.
```

Data Pipeline (Perception Layer L1)

Central bank documents and macro/market data are normalized through a unified perception module.

Code Example: Document Ingestion Pipeline

```
def process_event(doc_url):
    raw = download(doc_url)
    cleaned = clean_text(raw)
    tokens = tokenizer(cleaned)
    return tokens
world_state.set({
    "policy_tone": cb_agent.tone_score,
    "inflation_regime": macro_agent.infl_regime,
    "credit_stress": risk_agent.credit_state,
    "expected_equity_rotation": market_agent.rotation_signal
})
```

This makes the shared state explicit and traceable.

Semantic State Construction (Layer 2)

This illustrates the core function of Layer 2: constructing a structured semantic representation that captures the system's evolving understanding of the world.

The Semantic Understanding layer transforms raw perceptual inputs—such as a central bank statement—into explicit, interpretable semantic variables. The following example shows how a policy document is mapped into the semantic world model:

```
def update_semantic_state(raw_text):
    tone = extract_policy_tone(raw_text)    # e.g., "hawkish", "neutral", "dovish"
    inflation_signal = detect_inflation_context(raw_text)
    growth_assessment = detect_growth_signals(raw_text)
    return {
        "policy_tone": tone,
        "inflation_context": inflation_signal,
        "growth_context": growth_assessment
    }
```

Cross-Agent Mechanism Reasoning (L3)

(Persistent Collaborative Reasoning)

Layer 3 performs multi-step reasoning *over* the semantic world model while writing updates *back into* the model. This feedback loop is what creates persistent reasoning:

```
def macro_reasoning(semantic_state):
    if semantic_state["policy_tone"] == "hawkish":
        semantic_state["expected_rate_path"] = "higher_for_longer"
    if semantic_state["inflation_context"] == "elevated":
        semantic_state["inflation_risk"] = "upside"
    return semantic_state
```

The key properties illustrated here are: **Reasoning is stateful**: new inferences accumulate in the semantic state; **Reasoning is persistent**: updated variables become inputs for the next cycle; **Reasoning is interpretable**: all updates are explicit and inspectable.

This is a practical demonstration of how Layer 3 moves beyond a flat, single-shot prompt and supports longitudinal cognitive processes.

Agents reason sequentially using shared causal templates. This demonstrates how signals propagate across-agents:

```
policy → macro → risk → market.
cb_output = central_bank.reason(event_tokens)
macro_output = macro_agent.reason(cb_output)
risk_output = risk_agent.reason(macro_output)
market_output = market_agent.reason(risk_output)
```

Orchestration Logic (L4)

The orchestration layer determines the activation order and handles routing. Mechanism-Based Multi-Agent Orchestration

Layer 4 coordinates specialized agents using the shared semantic substrate.

The following minimal example demonstrates how a pipeline is orchestrated:

```
def orchestrate_pipeline(state):
    state = macro_agent.reason(state) # updates macro regime variables
    state = risk_agent.update(state) # maps regimes into risk-premium states
    return market_agent.predict(state) # produces market-reaction expectations
```

This short example highlights three central features of EGI orchestration: **Mechanistic routing**: each agent contributes domain-specific reasoning. **Shared world model**: all agents read and write to the same semantic state. **Causal propagation**: policy → macro regime → risk regime → market reaction.

This pipeline stands in contrast to a flat LLM prompt-chain, where no shared state, persistent reasoning, or mechanism-level causality is maintained.

Code Example: Workflow Controller

```
def run_pipeline(event):
    tokens = process_event(event)
    cb_state = central_bank(tokens)
    macro_state = macro_agent(cb_state)
    risk_state = risk_agent(macro_state)
    market_state = market_agent(risk_state)
    return market_state
```

This code concretely illustrates L4 orchestration over agent modules.

Meta-Cognitive Monitoring (L5)

Meta-cognition evaluates conflict, drift, and inconsistent updates.

Code Example: Consistency Check

```
def meta_check(world_state):  
    if world_state["policy_tone"] == "Hawkish" and world_state["yield_move"]  
< 0:  
        log("inconsistency detected: tone vs rates")  
        world_state.revise()
```

This example shows how contradictions trigger internal revision.

End-to-End Example Trace

This shows *mechanism-based reasoning*

Input document:

“Inflation remains elevated... Committee remains attentive to risks.”

Flow: **CentralBankAgent**

- extracts hawkish tone
- world_state["policy_tone"] = “hawkish”

MacroRegimeAgent

- inflation-regime = “high”
- growth-regime = “slowing”

RiskMechanismAgent

- volatility-regime = “elevated”

MarketReactionAgent

- predicts ↑ yields, ↑ credit spreads

ReasoningTrace

- stores all layer updates

Object-Oriented Architecture and System Implementation

Minimal OOP Implementation of an EGI Agent (Illustrative Code) is include as below.

This code illustrates how the five layers are structurally implemented:

class EGIAgent:

```
def __init__(self, name, world_state):  
    self.name = name  
    self.world_state = world_state # shared semantic model  
# L1: Perception  
def perceive(self, input_data):  
    return preprocess(input_data)  
# L2: Semantic Understanding  
def semantic_understanding(self, features):  
    return extract_semantic_variables(features)  
# L3: Collaborative Reasoning  
def reason(self, semantic_state):  
    return run_mechanism_rules(semantic_state)  
# L4: Orchestration
```

```

def orchestrate(self, result):
    return route_to_next_agent(result)
# L5: Meta-Cognition
def meta_cognition(self, logs):
    return detect_inconsistencies(logs)

```

The above illustrates how EGI agents use modular layers while sharing a world-state substrate.

Code Example: World-State Update

```

world_state.update({
    "policy_hawkishness": tone_score,
    "inflation_concern": inflation_index,
    "credit_stress": credit_state
})

```

This illustrates how world-model variables remain persistent and interpretable.

Cross-Agent Mechanism Reasoning (L3)

Mechanism-based reasoning connects domains through causal templates. In a financial and risk AI multiple agents mechanism, the reasoning chain can be: Policy → macro regime; Macro → risk-premium dynamics; Risk → market-reaction patterns

Code Example: Multi-Agent Propagation Chain

```

cb_output = cb_agent.reason(world_state)
macro_output = macro_agent.reason(cb_output)
risk_output = risk_agent.reason(macro_output)
market_view = market_agent.reason(risk_output)

```

The above describes how information flows across-agents.

Orchestration Logic (L4)

The orchestration layer governs: Activation sequence; Routing across-agents; Conflict resolution; Updating shared semantic model.

Example workflow for each central bank event:

Document → Central Bank Agent → Macro Agent → Risk Agent → Market Agent

Orchestration ensures coherent cross-domain reasoning.

Meta-Cognitive Monitoring (L5)

L5 evaluates: semantic drift; contradictory agent outputs; abnormal jumps in world-state temporal instability. This layer ensures long-horizon reasoning stability.

To operationalize the five-layer EGI framework in a reproducible and empirically testable way, the implementation of the system adopts a modular object-oriented architecture. The goal of the OOP design is to ensure that 1) each cognitive function is encapsulated, 2) multi-agent collaboration is transparent, and 3) persistent reasoning can be stored, reused, and evaluated. Below, we summarize the essential components.

Core Classes

-WorldState

A central class that stores *semantic state variables* updated by all agents.

It includes:

- policy stance indicators
- inflation and macro-regime variables
- risk-premium states
- credit and market reaction indicators
- confidence scores and temporal history

class WorldState:

```
def __init__(self):
    self.state = {}
    self.history = []
def update(self, key, value):
    self.state[key] = value
    self.history.append((key, value))
```

Purpose: Provides the shared semantic substrate required for cross-agent alignment (requested by reviewers in round 1).

AgentBase (Five-Layer Cognitive Skeleton)

All domain agents inherit from a unified abstract parent class representing the five cognitive layers.

class AgentBase:

```
def perceive(self, input_data):
    raise NotImplementedError
def interpret_semantics(self, percept):
    raise NotImplementedError
def reason(self, semantic_state):
    raise NotImplementedError
def orchestrate(self, world_state):
    raise NotImplementedError
def meta_cognition(self, trace):
    pass
```

Each method corresponds directly to one EGI layer:

EGI Layer	Agent Method
L1 Perception	perceive()
L2 Semantic Understanding	interpret_semantics()
L3 Collaborative Reasoning	reason()
L4 Orchestration	orchestrate()
L5 Meta-Cognition	meta_cognition()

Domain-Specific Agents

CentralBankAgent

Reads policy statements, speeches, and minutes.

Produces semantic variables such as:

- policy hawkishness
- inflation concern level
- forward guidance consistency

```
class CentralBankAgent(AgentBase):
    def perceive(self, text):
        return extract_key_passages(text)
    def interpret_semantics(self, percept):
        return infer_policy_tone(percept)
    def reason(self, semantic_state):
        return determine_policy_shift(semantic_state)
```

MacroRegimeAgent: Updates inflation regime/growth regime/monetary stance.

RiskMechanismAgent: Maps macro state into volatility, credit spread, and liquidity-risk regimes.

MarketReactionAgent: Predicts directional pressures on equity, rates, credit, and cross-asset behavior.

Together these four agents implement the mechanism chain.

Orchestrator (Mechanism Propagation Engine)

The Orchestrator coordinates multi-agent interactions and enforces mechanism logic.

```
class Orchestrator:
    def __init__(self, agents, world_state):
        self.agents = agents
        self.world = world_state
    def step(self, input_text):
        for agent in self.agents:
            percept = agent.perceive(input_text)
            semantic = agent.interpret_semantics(percept)
            update = agent.reason(semantic)
            self.world.update(agent.__class__.__name__, update)
```

Purpose: Ensures sequential runtime; Prevents cross-agent contradictions; Implements mechanism-based reasoning pathway; Enables empirical reproducibility

ReasoningTrace (Persistent Reasoning Memory)

Definition:

```
class ReasoningTrace:
    def __init__(self):
        self.steps = []
    def add(self, layer, content):
        self.steps.append((layer, content))
```

Stores: semantic updates; intermediate inferences; confidence adjustments; agent handoff records

This enables: auditability, stability analysis, longitudinal coherence tests