
From *Economic Agents* to *Agentic Economies*: A Systems Blueprint for Economic World Models

Jiale Han¹, Xiang Li², Jing Qian², Wenyuan Gu², Pin Gao², Ye Luo^{1,3},
Hongyuan Zha², Dacheng Tao⁴, Benyou Wang^{1,2}, Lin William Cong^{1,4*}

¹ Shenzhen Loop Area Institute

² School of Data Science, The Chinese University of Hong Kong, Shenzhen

³ University of Hong Kong ⁴ Nanyang Technological University

Abstract

Economic World Models require not only economic discipline but also scalable AI-system design. Building on Cong [2025], this paper develops an implementation-oriented perspective on how economic worlds can be built as computable AI systems. We study EconWM systems as generative engines in which heterogeneous agents act, interact, adapt, and co-evolve with markets and institutions, producing economic dynamics from inside the modeled world. We introduce a six-level capability ladder that traces how such systems become progressively more realistic: from fixed rule-based agent worlds, to adaptive and LLM-based agent worlds, to self-evolving agents, evolving institutional worlds, and finally self-correcting sim-to-real economic twins aligned with real observations. A systematic literature survey across these levels reveals a clear gap: existing work remains concentrated at lower levels, while higher-level systems with self-evolving agents, endogenous institutions, and persistent empirical alignment remain largely unrealized. The paper provides a CS/AI systems roadmap for implementing, evaluating, and scaling EconWMs as policy sandboxes, AI-agent training environments, planning engines, and safety testbeds. To support future research, we have released our curated paper list and related resources.²

1 Introduction

“If you didn’t grow it, you didn’t explain it.”

— Joshua M. Epstein

Economics has long advanced through a tradition of “observation and explanation”. Economists observe aggregate variables, collect micro-level evidence, estimate empirical relationships or causal effects, build theoretical and quantitative models, and ask what caused what [Haavelmo, 1944, Friedman et al., 1953, Heckman, 2001]. This approach is undeniably powerful to understand economy. Yet it also leaves open a deeper question of generative explanation. As Epstein [1999] famously argues, “If you didn’t grow it, you didn’t explain it.” Deriving an equilibrium, fitting parameters, or forecasting a turning point does not by itself explain an economic phenomenon unless the outcome can emerge from the model inside. This idea motivates the implementation problem studied here: **to make the Economic World Models agenda operational, we need AI systems in which economic outcomes can grow from agents, institutions, and rules inside the simulated world.**

This perspective connects naturally to the recent rise of *world models* in AI and reinforcement learning [Ha and Schmidhuber, 2018, Ding et al., 2025]. A world model is generally understood as an internal

* Send correspondence to Cong (will.cong@ntu.edu.sg) or Wang (wangbenyou@cuhk.edu.cn)

²<https://github.com/FreedomIntelligence/Awesome-Econ-World-Models/>

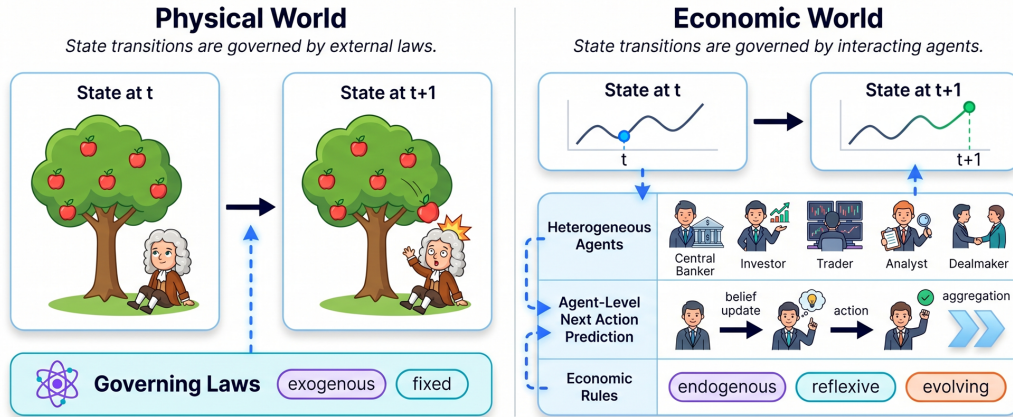


Figure 1: Physical versus economic worlds. Physical state transitions are governed by external laws, whereas economic state transitions are generated by heterogeneous agents interacting through markets, institutions, and evolving rules.

model of environment dynamics: given a current state and an action, it predicts the next states. Such models allow agents to imagine futures, evaluate possible actions, and plan before acting in the real world. In robotics [Wu et al., 2023a], autonomous driving [Wang et al., 2024a], and video generation [Wang et al., 2024b], world models are often built to capture the transition dynamics of physical or simulated environments. Their promise lies in turning prediction into simulation and simulation into decision support.

However, an economic world is not a physical world. As shown in Figure 1, in a physical world, state transitions are primarily governed by stable environmental laws, such as the laws of classical mechanics. Gravity, for example, constrains how objects move and does not vary according to the intentions of the objects it acts upon. An economic world is fundamentally different. Its state transitions are not merely imposed by an external environment; they are endogenously generated by the aggregated behaviors and interactions of the agents participating in it [Farmer and Foley, 2009, Arthur, 2021]. Prices, allocations, institutions, market structures, and expectations evolve as households, firms, banks, investors, and regulators observe, decide, trade, compete, coordinate, and adapt. Consequently, agents in the economic realm do not merely move through a pre-given environment; through their collective behavior, they produce the next state of the environment to which they later respond. This distinction implies that an EconWM system should natively accommodate four features: **heterogeneous agents**, because economic dynamics arise from diverse agents with different objectives, information, constraints, and decision rules; **endogenous outcomes**, because prices and allocations emerge from decentralized interaction rather than being imposed from outside; **reflexive feedback**, because beliefs about the economy reshape the economy and are reshaped by it in turn; and **evolving rules**, because the institutional environment itself may change over time.

Building on the EWM/DDGE framework of Cong [2025], we use the following AI-systems working definition for implementation purposes:

An EconWM system is a **generative engine for economic environments**: it predicts how an economy moves by modeling how the agents inside observe, reason, act, interact, and adapt under economic institutions and constraints.

This working definition highlights that **agents and world models are two sides of the same coin in economic modeling**. On one side, *agent-level next-action prediction* models how heterogeneous agents reason and choose actions under incentives, information, and constraints. On the other side, *economy-level next-state prediction* models how these actions interact through markets, institutions, and rules to generate prices, allocations, risks, institutional responses, and future states. The two sides cannot be separated: the economy moves through agents' actions, while agents can only be modeled meaningfully within the world that shapes their incentives, constraints, information, and feedback.

This coupling creates a co-improvement loop. More realistic EconWMs provide richer environments for training and planning agents under scarcity, incentives, and institutional constraints; more capable agents, in turn, improve EconWMs by producing more accurate action predictions, revealing strategic responses, and exposing edge cases. In short, **better worlds shape better agents, and better agents make better worlds.**

What EconWMs enable. With an EconWM, economic reasoning becomes operational and experimental. For human decision-makers, an EconWM functions as a high-fidelity sandbox for testing policies, market designs, firm strategies, systemic risks, and institutional reforms before they are deployed in the real world. For machine intelligence, it functions as an economic brain: a training environment where agents learn under scarcity, incentives, and institutional constraints; a planning engine for simulating downstream consequences before acting; and a safety testbed for detecting emergent risks such as manipulation, collusion, instability, and cascading failure. Over time, this creates a virtuous cycle: EconWMs train stronger economic agents, while stronger agents make the world model’s action predictions and state transitions more realistic.

Finance already provides several partial prototypes of the EconWM agenda. Cong et al. [2024] study an AI-mediated credit market in which borrower-side LLM adoption changes application text, lender-side multimodal models respond through retraining, and market outcomes depend on the resulting counterfactual equilibrium. Campello et al. [2025] develop a learned corporate environment and robust decision module for managerial policy search. Cong et al. [2021] show how financial decision agents can be trained directly on economic objectives through deep reinforcement learning, while Bini et al. [2025] study behavioral biases and correction mechanisms in LLM decision agents. These studies do not by themselves constitute complete EconWMs, but they provide important building blocks: learned financial environments, economic decision agents, robust control, behavioral calibration, and endogenous data-generation loops.

Relation to the EWM/DDGE framework. This paper is complementary to Cong [2025]. Cong introduces Economic World Models as an economic model class and develops Data-Driven Generative Equilibrium (DDGE) as the equilibrium discipline needed when behavior, beliefs, data generation, and learned environments are jointly endogenous. The present paper takes a CS/AI systems perspective: it asks how such economic worlds can be implemented, evaluated, and scaled as computational environments for agent training, policy sandboxes, planning engines, and sim-to-real economic twins. Accordingly, the capability levels below classify implementation maturity; they do not by themselves guarantee DDGE-style counterfactual consistency.

Why now? Large language models provide a substrate for modeling beliefs, memory, communication, reasoning, and context-sensitive decision-making [An et al., 2024, Xu et al., 2025, Ferrag et al., 2025, Guo et al., 2025]. Tool-using agents [Masterman et al., 2024, Shi et al., 2025] connect decisions to executable economic actions such as search, purchase, negotiation, trading, planning, and compliance. The rise of agentic AI further turns agents into actors that can plan, call tools, invoke reusable skills, and execute multi-step tasks through external harnesses [Zhou et al., 2026]. Multi-agent simulation [Chen et al., 2024b] provides an interaction substrate in which heterogeneous agents can respond to one another and jointly generate aggregate outcomes. Meanwhile, the growing availability of financial, textual, policy, administrative, and transaction data provides the empirical foundation for constructing, calibrating, and validating EconWMs. These ingredients make Economic World Models newly imaginable.

Contributions and organization. This paper makes four contributions. First, we provide an implementation-oriented working definition for EconWM systems. Section 2 develops the AI-systems view, identifies engineering desiderata for implementation, and introduces a six-level capability hierarchy. Second, we present an architectural blueprint for making EconWMs buildable. Section 3 develops this architecture. Third, we map the emerging field from both capability-based and historical perspectives. Section 4 uses the six-level hierarchy to assess how powerful current EconWM systems are, revealing a central gap: existing systems remain concentrated at lower levels, while co-evolving institutional worlds and empirically grounded economic digital twins are still largely missing. Section 5 takes a complementary historical view, tracing how AI for economics has moved through five waves: feature engineering, data engineering, prompt engineering, context engineering, and environment engineering. Fourth, we position EconWMs as infrastructure for decision intelligence and clarify their boundaries. Section 6 develops two application roles: sandboxes for human decision-

makers and economic brains for autonomous AI agents. Section 7 positions EconWM systems among adjacent paradigms. Section 8 closes with open challenges and future research directions.

The road ahead. Our aim in this paper is to translate the EWM agenda into a CS/AI systems roadmap: to articulate the implementation vision, operationalize the core concepts, build a capability ladder, and provide an architectural starting point. But building trustworthy, scalable, and empirically grounded EconWMs will require a community effort. We invite researchers across artificial intelligence, economics, management, social science, and policy to help build, evaluate, and govern the next generation of economic worlds. The goal is to make economic worlds buildable, testable, and alignable as AI systems, while preserving the economic discipline emphasized by the EWM/DDGE framework.

2 Economic World Models as AI Systems

2.1 An Implementation-Oriented Working Definition

World Models. Generally, a world model is a computable dynamical system that represents how the state of a world evolves over time, especially under actions or interventions. Its core role is to approximate the transition dynamics of an environment: given a current state and an action, it predicts the next state. Formally, a world model can be written as

$$\hat{s}_{t+1} = \mathcal{T}(s_t, a_t), \quad (1)$$

where s_t is the current state, a_t is an action or intervention, \hat{s}_{t+1} is the predicted next state, and \mathcal{T} denotes the transition operator.

Physical Worlds vs. Economic Worlds. Physical worlds are governed by stable laws: planets orbit, pendulums swing, and objects fall in ways that are consistent and predictable. Economic worlds are fundamentally different. They are shaped by strategic agents, institutional constraints, and feedback effects that are inherently adaptive and self-generating. Four features make economic worlds especially distinct.

- **Heterogeneous agents:** Unlike many physical systems that are governed by relatively uniform laws of motion, economic systems are populated by heterogeneous agents that differ in objectives, information, constraints, and decision rules.
- **Endogenous outcomes:** Unlike many physical systems in which outcomes are primarily determined by external forces and stable laws, core economic variables such as prices, allocations, and quantities are not externally imposed but emerge from decentralized interactions such as bidding, bargaining, trading, and market-clearing among agents.
- **Reflexive feedback:** Unlike most physical processes, where feedback does not typically depend on agents' beliefs about the system itself, economic dynamics are shaped by feedback between expectations and outcomes: agents form beliefs, act on them, thereby altering markets and macro conditions, which in turn reshape subsequent beliefs and decisions.
- **Evolving rules and mechanisms:** Unlike physical laws, which are relatively stable, economic institutions, policy rules, and behavioral regularities may themselves change over time as agents adapt, regulations shift, and new norms emerge, so that the transition dynamics co-evolve with the system.

Economic World Model Systems. For CS/AI implementation, we use EconWM to denote a computable dynamical system that captures the transition of an economy driven by heterogeneous and interacting agents. This working definition is narrower and more operational than the economic framework in Cong [2025]: it focuses on system components, capability levels, and implementation requirements rather than on equilibrium concepts.

As summarized in Figure 2, the state of the economy is given by

$$s_t = (x_t, \{z_t^i\}_{i=1}^{N_t}, \{b_t^i\}_{i=1}^{N_t}, \mathcal{I}_t), \quad (2)$$

where x_t denotes the macroeconomic economic state, such as prices, quantities, and macroeconomic conditions, $\{z_t^i\}_{i=1}^{N_t}$ are the individual states of heterogeneous agents, $\{b_t^i\}_{i=1}^{N_t}$ are agents' beliefs or




Figure 2: State transition in an Economic World Model.

expectations, and \mathcal{I}_t denotes the institutional environment and governing rules. Given an intervention a_t , the transition operator decomposes into the following stages:

$$b_{t+1}^i = B^i(b_t^i, z_t^i, x_t, a_t, \mathcal{I}_t), \quad i = 1, \dots, N_t \quad (\text{agent belief update}) \quad (3)$$

$$a_{t+1}^i = \pi^i(z_t^i, b_{t+1}^i, x_t, a_t, \mathcal{I}_t), \quad i = 1, \dots, N_t \quad (\text{next-action prediction of agents}) \quad (4)$$

$$z_{t+1}^i = \Omega^i(z_t^i, a_{t+1}^i), \quad i = 1, \dots, N_t \quad (\text{agent state update}) \quad (5)$$

$$x_{t+1} = G(\{z_{t+1}^i, a_{t+1}^i, b_{t+1}^i\}_{i=1}^{N_t}, x_t, \mathcal{I}_t) \quad (\text{next-state prediction of economic worlds}) \quad (6)$$

$$\mathcal{I}_{t+1} = \Phi(\mathcal{I}_t, x_{t+1}, \{b_{t+1}^i, a_{t+1}^i, z_{t+1}^i\}_{i=1}^{N_t}) \quad (\text{rules and mechanisms evolution}) \quad (7)$$

where B^i denotes the belief-update operator of agent i , mapping the agent's prior belief together with its current state, the aggregate economic state, the exogenous intervention, and the institutional environment into an updated subjective representation of the economy, π^i denotes the decision rule of agent i , mapping the agent's current state and updated belief to an action, Ω^i denotes the state transition function of agent i , updating the agent's internal state after the action is taken, such as wealth, inventory, balance-sheet position, or other micro-level variables, G denotes the endogenous aggregation and market-closure operator, which maps the collection of individual agent states into next-period aggregate outcomes such as prices, allocations, and macroeconomic conditions, and Φ denotes the evolution of institutions or governing mechanisms.

2.2 Engineering Desiderata for EconWM Systems

For implementation, a strong EconWM system should satisfy four engineering desiderata:

- D1. Endogenous closure.** *An EconWM system should produce key economic outcomes endogenously.* Prices, allocations, and macroeconomic conditions should emerge from decentralized interactions among heterogeneous agents, rather than be imposed from outside. This matters because these outcomes do not simply describe the state of the economy. They also shape what agents do next by feeding back into subsequent decisions and interactions. If they are treated as exogenous, the model can no longer represent the feedback process through which an economic world evolves.
- D2. Behavioral fidelity.** *Agents in an EconWM system should resemble real economic actors.* Agent behavior should reflect how households, firms, banks, and investors perceive, interpret, and respond to incentives, information, and constraints. Such fidelity may be realized in different ways, ranging from symbolic mechanisms that adapt rules and strategies in response to economic signals, to richer cognitive substrates that maintain explicit beliefs, expectations, and subjective representations of the economy. What matters is capturing heterogeneity, bounded rationality, and context-dependent adaptation, rather than reducing behavior to fixed or overly stylized rules. Behavioral fidelity is especially important when LLMs are used as economic agents. Recent evidence shows that LLMs can exhibit systematic behavioral patterns in economic and financial decision tasks, including biases in preference-based tasks and more rational responses in some belief-based tasks; prompting and correction procedures can mitigate some of these biases [Bini et al., 2025]. EconWM agents therefore require behavioral validation rather than persona prompts alone.
- D3. Evolving dynamics.** *In an EconWM system, both the agents and the governing rules of the economy should co-evolve over time.* Economic worlds are not static: strategies adapt, beliefs shift, agents may even acquire new cognitive capabilities and behavioral repertoires through

Table 1: Capability levels of EconWM systems as progressively stronger realizations of the engineering desiderata. Symbols indicate full, partial, and missing satisfaction: ✓, ○, and ✗. Agent capability improves from L1 to L4, rules remain fixed through L4 and evolve at L5, and real-time correction appears at L6.

Desiderata	L1 Fixed-rule agent worlds	L2 Adaptive agent worlds	L3 LLM-based agent worlds	L4 Self-evol. agent worlds	L5 Evolving econ worlds	L6 Sim-to-real econ twins
Endogenous closure	✓	✓	✓	✓	✓	✓
Behavioral fidelity	○	○	○	✓	✓	✓
Evolving dynamics	✗	○	○	○	✓	✓
Reality alignment	✗	✗	✗	✗	✗	✓

	fixed	adaptivity	autonomy	self-evolution
Agent	fixed rules			
Rules	no online correction			rules evolve
Real world				online correction

experience, and institutional conditions evolve. A credible EconWM should therefore model not only state transitions, but also the evolution of strategies, cognition, and governing rules at each of these layers.

D4. Reality alignment. *An EconWM system should stay aligned with the real economy over time. It should not only generate plausible simulations, but also repeatedly compare them with newly observed outcomes and update its internal dynamics accordingly. In this way, the model can progressively narrow the gap between the simulated economy and the real one.*

2.3 Capability Levels of EconWM Systems

As shown in Table 1, EconWM systems can be organized into a hierarchy of six progressively stronger capability levels. These levels represent successive enrichments of the same underlying economic transition process: heterogeneous agents act under institutional conditions, their internal states evolve, and their interactions endogenously generate the next aggregate state. Moving up the hierarchy, the model incorporates increasingly rich forms of rule-based agent adaptation, LLM-driven autonomous reasoning, endogenous self-evolution of agent capabilities, institutional evolution, and reality-based self-correction. The hierarchy is an implementation taxonomy rather than an equilibrium concept: a system can be high on the capability ladder and still fail counterfactual consistency if learned components are not re-centered on the data generated by the counterfactual itself, the problem addressed by DDGE in Cong [2025].

Level 1: Fixed Rule-Based Agent Worlds. *This level models the economy as a closed system in which agents and the environment follow fixed, pre-specified rules.*

At the basic level, the economy is represented as a fixed rule-based system in which the agent decision rule π^i , the agent state transition Ω^i , and the aggregation operator G are all fixed ex ante. The cognitive operator B^i is absent and decisions are made directly from the agent’s state and the external environment; the institutional evolution operator Φ is also inactive, so the institutional environment remains unchanged over time. Even so, this level already captures the core feedback loop between heterogeneous agents and aggregate outcomes: aggregate states arise from agent interactions and in turn shape subsequent decisions. Such systems are transparent and interpretable, yet limited in flexibility and adaptivity.

Level 2: Adaptive Rule-Based Agent Worlds. *Adaptive Rule-Based Agent Worlds are economic worlds in which agents adapt their rule-based decision policies over time using symbolic or algorithmic mechanisms.*

At this level, agents are no longer fixed rule executors. The decision rule itself is allowed to vary with time, written as π_t^i , and is updated as a function of the agent’s history of interactions and aggregate outcomes generated within the simulated economy. The substrate underlying π_t^i remains symbolic or algorithmic: rule systems with adaptive rule selection, online reinforcement learners whose policies update during the rollout, evolutionary or genetic-algorithm agents, online learning agents, or adaptive heuristics. Agents at this level do not yet maintain rich subjective representations of the economy, so the cognitive operator B^i is still inactive; adaptation is captured through changes in the parameters

of π_t^i rather than through the formation and revision of beliefs and expectations. The institutional structure remains fixed, meaning that the institutional evolution operator Φ is not active.

Level 3: LLM-Based Autonomous Agent Worlds. *LLM-Based Autonomous Agent Worlds are economic worlds in which agents are driven by large language models, possess explicit cognitive states, and make autonomous decisions through belief-driven reasoning over time.*

At this level, agents are equipped with rich internal cognitive states—beliefs, expectations, and subjective representations of the economy—and use them to make decisions autonomously. Formally, the cognitive operator B^i becomes active and produces updated beliefs at each step, and the time-varying decision rule π_t^i is conditioned on these updated beliefs in addition to the external economic environment. Adaptation operates jointly through the cognitive layer and the strategy layer: as beliefs evolve with perceived signals, interaction history, and contextual information, the decision rule π_t^i responds in a context-sensitive way, allowing the model to capture expectation formation, heterogeneous perceptions, and behavioral patterns that respond flexibly to the same economic conditions in different ways. The cognitive substrate underlying B^i and π_t^i itself remains fixed during the rollout—its underlying reasoning capacity does not expand—but the cognitive states it produces are dynamic and history-dependent. The institutional structure of the world remains fixed at this level.

Level 4: Self-Evolving Agent Worlds. *Self-Evolving Agent Worlds are economic worlds in which agents not only update their beliefs and decisions but also expand their cognitive capabilities and behavioral repertoires over time.*

At this level, agent adaptation extends from beliefs and strategies to the cognitive substrate itself. The operators B^i and π_t^i are now time-varying not only in the states and actions they produce but also in their underlying capacity: agents acquire new strategies, new skills, new tools, or new behavioral routines during the simulated economic rollout, and these acquisitions persist and reshape the operators in subsequent steps. This corresponds to an endogenous capacity update of B^i and π_t^i themselves, driven by signals generated inside the model—reflection on past beliefs, observation of interaction outcomes, and learning from realized aggregate states—rather than by offline training. Agents may, for instance, develop new trading heuristics, learn new negotiation patterns, build new analytical procedures, or accumulate domain knowledge through reflection on past successes and failures, so that the agents acting in later periods are cognitively richer than those acting earlier. The institutional environment still remains fixed at this level, but the agent population is no longer cognitively static.

Level 5: Evolving Economic Worlds. *This level refers to economic worlds in which the rules and mechanisms governing the economy can evolve endogenously over time.*

At this level, adaptation extends beyond agents to the governing structure of the economy itself. In addition to agent-level adaptation, the institutional environment is allowed to evolve through the institutional evolution operator Φ , which may take the form of a simple update rule, an adaptive policy mechanism, a higher-level decision made by simulated policymakers with their own beliefs and objectives, or an emergent process that selects among candidate institutions based on aggregate outcomes and realized agent histories. The updated institution then becomes the environment within which agents subsequently act, so that policy regimes, market rules, contracts, governance structures, and broader institutional conditions co-evolve with agent behavior. Importantly, Level 5 only requires the institutional evolution operator Φ to be active alongside in-world agent adaptation; the agent cognitive substrate may correspond to any of Levels 2–4. What distinguishes this level is therefore not the sophistication of individual agents but the fact that the economic world itself is no longer a fixed stage on which agents act, though a Level 5 environment derived from Level 4 substrates is undeniably the most ideal state.

Level 6: Sim-to-Real Economic Twins. *At this level, the EconWM continuously closes the gap between simulation and reality, moving toward a digital twin of the economy.*

At the highest level, the EconWM no longer evolves only through its internal dynamics. It remains continuously tethered to the real economy. Whenever new real-world observations become available, the model compares its simulated outcomes against these observations and recalibrates itself to reduce the discrepancy. This creates a two-loop architecture: an inner loop in which the economic world endogenously evolves, and an outer loop in which empirical feedback drives repeated updates to the model. In this way, the model continuously learns from the real world, progressively narrowing the

sim-to-real gap. The highest-level EconWM therefore becomes a genuine self-correcting digital twin of the economy, enabling reliable counterfactual analysis.

3 Implementation Architecture for EconWM Systems

Several recent finance applications illustrate how these modules arise in practice. AlphaPortfolio is primarily an agent-side decision module: it trains a portfolio decision agent directly on investment objectives under market states [Cong et al., 2021]. AlphaManager adds a learned environment module and a robust-control decision layer for corporate policy search [Campello et al., 2025]. The online-lending application adds a richer agent–environment feedback loop: borrower-side generative behavior changes the data observed by lender-side discriminative models, and lender retraining changes future credit allocation [Cong et al., 2024]. These examples motivate the modular architecture below: agents, environments, co-evolution, and alignment should be designed as separable but interacting components.

As shown in Figure 3, an EconWM system can be implemented as a four-layer architecture: economic agents that make micro-level decisions, economic environments that constrain actions and return feedback, agent–environment co-evolution that generates adaptive dynamics, and real-world alignment that anchors the simulated world to reality.

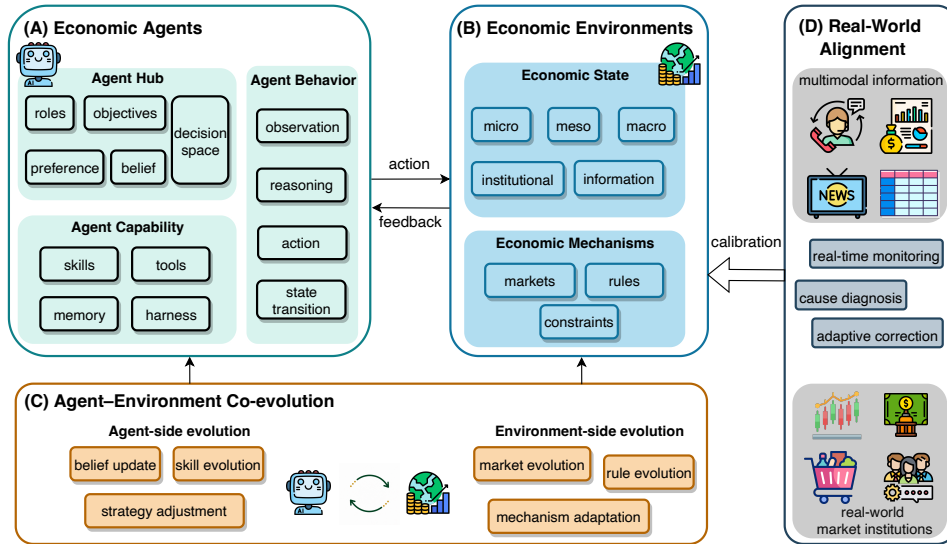


Figure 3: Core components of an Economic World Model.

3.1 Economic Agents

Agents are the micro-foundations of EconWMs. They represent economic actors such as households, firms, investors, banks, platforms, regulators, and governments. Unlike generic AI agents, economic agents must be defined by their roles, objectives, constraints, beliefs, feasible actions, and institutional positions. Their behavior is not merely text generation or rule execution, but economically meaningful decision-making under scarcity, incentives, and uncertainty. We organize economic agents into three components: the agent hub, the agent capability library, and the agent behavior workflow.

3.1.1 Agent Hub: Who Agents Are

The Agent Hub constructs and manages the population of agents in an EconWM. Its goal is to generate economically meaningful heterogeneity rather than superficial diversity. Each agent is specified by its role, objectives, constraints, information, and decision space.

Roles and institutional positions. Agents occupy distinct economic roles. Households participate in labor, consumption, credit, and housing markets; firms set prices, hire workers, manage inventories,

and invest; banks allocate credit and manage balance-sheet risk; governments and regulators design and enforce policy rules. These roles determine what agents can observe, what constraints they face, and what actions they can take.

Objectives and preferences. Agents differ in what they optimize or value. Households may care about consumption, savings, employment security, and risk. Firms may pursue profit, survival, market share, or long-run growth. Investors may optimize returns, liquidity, or risk exposure, while policymakers may target welfare, stability, inflation, employment, or fairness. Preferences such as risk aversion, time discounting, liquidity demand, and loss aversion shape how agents evaluate alternatives.

Endowments and constraints. Agents also differ in what they possess and what they are allowed or able to do. These include income, wealth, debt, inventory, productivity, capacity, credit access, leverage, capital requirements, information limits, and regulatory restrictions. Such constraints determine not only what agents want to do, but also what they can actually do.

Information, beliefs, and expectations. Agents do not observe the entire economy. They receive public signals, private information, local observations, prices, news, and messages from other agents. From these inputs, they form beliefs about current conditions, future states, risks, and others' behavior. This belief formation is central to economic reflexivity: agents act on beliefs, their actions change the economy, and the changed economy reshapes future beliefs.

Decision space. Each agent has a set of feasible actions. Households may consume, save, borrow, default, migrate, or search for jobs. Firms may set prices, produce, invest, hire, borrow, or exit. Investors may buy, sell, hedge, or rebalance. Policymakers may adjust taxes, transfers, interest rates, regulations, or enforcement intensity. The Agent Hub therefore defines who the agents are, what they know, what they want, what they can do, and how they are embedded in the economy.

3.1.2 Capability Hub: What Agents Can Use

The Capability Hub specifies the operational resources that allow agents to act effectively within an EconWM. While the Agent Hub defines who the agents are, the Capability Hub defines what agents can use to reason, act, learn, and comply with constraints. It consists of four core components: skills, tools, memory, and harnesses.

Skills. Skills are reusable capabilities that allow agents to solve recurring tasks through learned or programmed routines rather than ad hoc responses [Zhang et al., 2025a, Xu and Yan, 2026]. In economic settings, skills may include forecasting, pricing, negotiation, portfolio adjustment, credit assessment, inventory management, and policy evaluation. These skills give agents task-level competence and make their behavior more stable, interpretable, and reusable across scenarios.

Tools. Tools are callable external resources that extend an agent's reasoning beyond its internal model. They allow agents to search, calculate, retrieve, verify, optimize, and execute actions in external systems [Schick et al., 2023, Yao et al., 2023]. In economic settings, tools may include calculators, databases, search engines, optimization solvers, trading systems, accounting modules, credit-scoring models, and forecasting systems. Through tools, agents can query market information, compute financial quantities, check feasibility constraints, and connect language-based reasoning to executable economic operations.

Memory. Memory stores past observations, interactions, decisions, and outcomes, allowing agents to maintain persistence across time. Memory-augmented agents can retrieve prior experiences, track relationships, reflect on past behavior, and adapt over repeated interactions [Packer et al., 2023, Zhang et al., 2025b]. In economic settings, memory may record transaction histories, counterparties, fulfilled or broken contracts, past prices, credit outcomes, policy responses, and reputational signals. This enables history-dependent behavior, relationship formation, learning from past mistakes, and belief updates across economic cycles.

Harnesses. Harnesses are execution frameworks that coordinate the agent's reasoning and action process. They organize observation, planning, memory retrieval, tool use, skill invocation, feasibility checking, and action submission into a controlled loop [Zhou et al., 2026]. In economic settings, harnesses ensure that agents do not merely produce plausible language, but submit valid actions under market rules, institutional procedures, accounting identities, budget constraints, and regulatory

requirements. They are therefore the bridge between cognitive reasoning and economically valid execution.

Together, these components transform agents from static profiles into operational economic actors. Skills provide task competence, tools connect reasoning to external operations, memory creates persistence, and harnesses enforce valid execution. The Capability Hub therefore supports belief formation, feasible action generation, and experience accumulation over time.

3.1.3 Behavior Hub: How Agents Decide and Act

While the Agent Hub defines who the agents are and the Capability Hub defines what agents can use, the Behavior Hub defines how agents behave at each time step.

Observation. The agent receives signals from the environment, including market variables (e.g., prices and quantities), institutional signals (e.g., policies), and interaction-based information (e.g., messages or local outcomes). These signals are filtered by the agent's information set, reflecting partial and heterogeneous observability.

Reasoning. Given new observations, the agent interprets the current economic situation and evaluates possible courses of action. This process may involve updating expectations about future states, assessing risks, inferring other agents' behavior, retrieving relevant memories, invoking skills, and using external tools.

Action. Based on its observation, reasoning, objectives, and constraints, the agent makes an economically feasible action. The action space is role-dependent and may include consumption, production, pricing, hiring, trading, borrowing, or policy intervention. Feasibility is jointly determined by internal states and external institutional rules.

State transition. Executed actions induce changes in the agent's internal state, including wealth, inventory, employment, leverage, and belief configurations. These updated states define the agent's starting point for the next decision cycle.

3.2 Economic Environments

While agents provide the micro-foundations, the environment defines how individual actions are constrained, mediated, aggregated, and fed back as economic outcomes. We describe the environment along three dimensions: the multimodal state representation that captures economic conditions, the mechanisms and rules that govern interactions, and the update process that turns agent actions into the next economic state.

3.2.1 Economic State: What the World Contains

The economic state represents the current condition of the simulated economy. It is structured at multiple levels to capture both micro-level details and macro-level aggregates.

Micro-level state includes agent-specific variables such as income, wealth, debt, employment status, inventory, productivity, cash flow, leverage, beliefs, and portfolios.

Meso-level state captures intermediate structures, including sectoral distributions, supply chains, interbank exposures, regional markets, platform networks, and labor-market conditions.

Macro-level state aggregates the broader economy, encompassing prices, output, employment, inflation, interest rates, credit conditions, asset prices, liquidity, volatility, inequality, and fiscal conditions.

Institutional-level state represents the prevailing rule system, including tax schedules, regulatory regimes, policy stances, market protocols, enforcement mechanisms, and institutional constraints.

Information state contains unstructured and multimodal information, such as news, policy speeches, earnings calls, contracts, social media, market narratives, images, audio signals, and inter-agent messages. These signals shape agents' beliefs, expectations, coordination, and market sentiment, making them essential to realistic economic dynamics.

3.2.2 Economic Mechanisms: How Actions Are Mediated

Economic mechanisms and rules define how interactions are organized and how outcomes are generated. They constitute the operating system of the economic world.

A complete environment may include **market mechanisms**, such as auctions, order books, bargaining protocols, matching processes, posted-price markets, and market-clearing rules. It may also include **pricing and allocation rules**, which determine how prices, quantities, contracts, and resources are formed and distributed.

The environment must also encode **accounting and feasibility constraints**. Budget constraints, balance-sheet identities, production capacity, inventory limits, leverage restrictions, resource constraints, and regulatory requirements ensure that simulated outcomes remain economically coherent.

In addition, the environment includes **contracting, enforcement, and institutional rules**. Loans, labor contracts, supply agreements, collateral arrangements, tax rules, monetary-policy regimes, regulatory constraints, and legal restrictions all shape how agent actions translate into outcomes. At lower levels of EconWMs, these rules may be fixed. At higher levels, they may evolve in response to agent behavior and aggregate outcomes.

3.2.3 Environment Update: How the World Moves Forward

The environment state is updated by processing agent actions through economic mechanisms. This process can be broken down into five key steps:

Action collection. The environment collects actions from agents, such as consumption, production, pricing, trading, and policy interventions.

Feasibility check. Actions are verified for validity under existing rules and constraints. Invalid actions are either rejected or adjusted.

Interaction execution. Valid actions are executed through markets and institutions, where buyers and sellers interact, firms meet workers, and policies affect agents.

Outcome aggregation. The environment aggregates individual actions into macroeconomic variables, updating prices, output, employment, credit flows, and fiscal balances.

State publication. The updated state is shared with agents according to their information sets, with some signals being public, others private or delayed.

This update mechanism ensures that the economic state evolves endogenously from agent interactions within a rule-governed environment, rather than being externally imposed.

3.3 Agent–Environment Co-evolution

A higher-level EconWM should not assume fixed agents acting in a fixed environment. In real economies, agents adapt to the environment, and the environment changes in response to agents. This mutual adaptation is the source of co-evolution.

On the **agent side**, beliefs, strategies, decision rules, memory, tool use, skills, and even objectives may evolve over time. Agents learn from realized outcomes, imitate successful strategies, adjust to new constraints, and develop new routines. A firm may revise its pricing strategy after repeated demand shocks; an investor may change risk exposure after a crisis; a household may alter saving behavior after unemployment; a regulator may change its policy rule after observing systemic instability.

On the **environment side**, economic states, market structures, networks, policies, rules, and institutions may also evolve. Firms enter and exit, supply chains reorganize, financial networks rewire, platforms redesign matching rules, regulators adjust enforcement, and governments revise policy regimes. The rules of the game are therefore not always fixed; they may become endogenous objects of the model.

This co-evolution can be implemented through several adaptation strategies. Rule-based adaptation updates agents or rules according to predefined thresholds, triggers, or state-contingent conditions. Selection-based adaptation changes the population composition by allowing successful strategies, firms, or institutions to expand, while unsuccessful ones shrink, exit, or are replaced. Learning-based

adaptation allows agents or institutions to update policies, beliefs, or skills through reinforcement learning, imitation learning, evolutionary search, or belief-based updating.

3.4 Real-World Alignment

An EconWM should not only generate internally coherent economic dynamics. It must also remain aligned with the real economy. Real-world alignment refers to the process of anchoring simulated agents, mechanisms, states, and outcomes to empirical evidence. Without this alignment, an EconWM may produce plausible narratives or stable internal dynamics while drifting away from observed economic behavior.

Alignment operates at several levels. At the state level, simulated macroeconomic and financial variables should be compared with real-world observations, such as output, employment, inflation, interest rates, credit conditions, asset prices, volatility, liquidity, and fiscal balances. At the agent level, simulated behavior should be checked against observed patterns in consumption, saving, borrowing, investment, trading, hiring, pricing, default, and policy response. At the mechanism level, market-clearing rules, matching processes, pricing protocols, credit allocation, and institutional constraints should reflect the way real economic systems operate.

Real-world alignment can be implemented as an online correction loop supported by several complementary techniques. The loop performs three functions, namely real-time monitoring, cause diagnosis, and adaptive correction. Real-time monitoring continuously compares simulated behaviors, outcomes, and reasoning traces with real-world evidence, using techniques such as LLM-as-judge evaluation [Gu et al., 2024, Li et al., 2024a], expert-in-the-loop assessment [Natarajan et al., 2025], and statistical calibration. Cause diagnosis then attributes detected deviations to possible sources, including agent behavior, market mechanisms, or state-transition processes. Adaptive correction updates the corresponding components of the EconWM. Specifically, agent alignment adapts agent parameters, skills, memories, or policies using observed economic behavior and evaluator feedback. Mechanism correction and rule recalibration revise market mechanisms, institutional constraints, and transition functions.

This alignment loop is especially important at the highest levels of EconWMs. Lower-level systems may only be calibrated once before simulation. Higher-level systems should continuously absorb new data, compare simulated trajectories with observed economic dynamics, and revise the world accordingly. In this sense, real-world alignment turns an EconWM from a closed artificial economy into a living empirical model. It anchors imagination to evidence, ensuring that counterfactual worlds remain useful for real-world reasoning, planning, and safety verification.

4 Representative EconWM Systems Across the Six Capability Levels

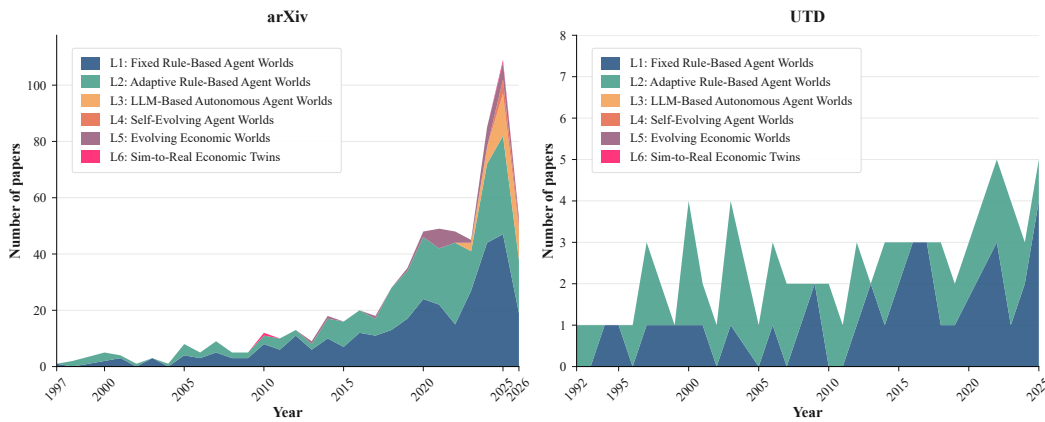


Figure 4: A literature landscape of Economic World Model research across the six capability levels.

To obtain a high-level picture of how existing research relates to our six-level hierarchy, we collected papers from top computer science conferences and leading business and economics journals using

a broad set of queries related to agent-based economics, macroeconomic simulation, digital twins, reinforcement learning in economic environments, and language-model-based economic agents. We then applied a three-stage classification pipeline. The first stage performs keyword-based pre-filtering on titles and abstracts to assemble a broad pool of candidate papers. The second stage applies a recall-oriented LLM screening prompt that retains papers exhibiting EconWM-like structure—economic agents, endogenous outcomes from agent interactions, and dynamic feedback—while discarding those with only generic economic or AI terminology. The third stage applies a precision-oriented full-PDF classification prompt that judges each paper against the EconWM engineering desiderata and, for papers identified as EconWMs, assigns one of the six levels; a subset of borderline cases is further verified through manual review. The detailed operational procedures for the above steps can be found in Appendix A. Figure 4 provides an illustrative view of the resulting literature distribution across the six levels over time. This process yields a final set of 737 papers.

Key Findings. The literature landscape reveals several important patterns.

- **Agent intelligence has become the main frontier of recent progress.** The arXiv literature grows rapidly with the rise of LLM-based and learning-based agents, with most papers concentrated in L1–L3. This suggests that recent progress is driven mainly by richer agent intelligence rather than by fully evolving economic worlds.
- **Business and economics journals remain concentrated at lower levels.** Papers from leading business and economics journals are fewer in number and remain concentrated in lower levels, reflecting the longer tradition of rule-based agent-based modeling.
- **Higher-level EconWMs remain rare.** L4 systems with genuinely self-evolving agents are rare, L5 systems with endogenous institutional evolution are even less common, and L6 sim-to-real economic twins remain largely an open frontier.

4.1 Level 1: Fixed Rule-Based Agent Worlds

The representative works at this level construct a fixed rule-based economic world in which prices, allocations, and macroeconomic patterns arise endogenously from decentralized interaction, but the behavioral rules of agents, the environment dynamics, and the governing mechanisms are all specified *ex ante*.

A first cluster develops microscopic market models in which heterogeneous boundedly rational traders generate endogenous price dynamics, volume, and return distributions from simple local rules [Iori, 2000, Sznajd-Weron and Weron, 2000, Solomon and Richmond, 2001, Degryse et al., 2009]. A second cluster extends rule-based interaction to monetary, network-resource, and platform-level economies, where money values, bandwidth prices, token dynamics, or Treasury spreads emerge from interactions among agents subject to fixed optimization or institutional constraints [Bornholdt and Wagner, 2001, Rasmusson and Aurell, 2001, Cong et al., 2022, He et al., 2022]. A third cluster studies endogenous strategic and asset-pricing dynamics in general-equilibrium settings, where prices, volatility, and entrepreneurial rents emerge from heterogeneous agents with fixed decision rules [Hugonnier and Prieto, 2015, Keyhani et al., 2015].

Across these works, the economic world is governed by fixed rules and the feedback loop between agents and aggregate outcomes is genuine, but agents do not adapt their decision rules in-world, no cognitive substrate maintains explicit beliefs or expectations, and neither institutions nor the model itself evolves through experience.

4.2 Level 2: Adaptive Rule-Based Agent Worlds

What separates this level from fixed rule-based worlds is a single but consequential capability: agents now revise their own strategies in response to what unfolds inside the economy. The cognitive machinery driving these revisions, however, remains symbolic—reinforcement learning, evolutionary search, score-based strategy switching, Bayesian updating over a fixed parameter set—rather than language-based reasoning over rich subjective representations.

The earliest works in this strand emerge from physics-inspired financial modeling. Minority-game and Ising-type traders update strategy scores or opinions from realized market outcomes, producing endogenous volatility clustering, fat-tailed returns, and bubble–crash dynamics [Bak et al., 2000,

Sherrington et al., 2000, Corcos et al., 2001, Rothenstein and Pawelzik, 2002]. Vidal and Durfee [2000] formalizes this picture from a different angle, deriving difference equations that predict how reinforcement-learning agents converge—or fail to converge—in repeated market settings.

A parallel literature in industrial organization and corporate finance turns the same idea toward firm-level decisions. Firms search over pricing and development strategies as demand evolves [Adner, 2002], accumulate or lose knowledge as rivals diffuse innovations [Knott, 2003], and learn capital-structure choices through evolutionary selection over realized payoffs [Noe et al., 2003]. More recent work pushes adaptive rule-based agents into operational and intermediation settings: Castillo et al. [2022] simulates crowdsourced delivery drivers whose acceptance probabilities track tips, mileage, and opportunity costs, while Lee and Neuhann [2023] embeds adaptive borrower effort inside a dynamic lending economy where collateral quality, liquidity, and market freezes co-evolve.

What is missing across all of these works is internal cognition. Agents adapt, but they do so through update rules over numerical states; they do not maintain beliefs about the economy in any rich sense, do not reason over context, and operate inside institutional structures that themselves remain fixed throughout the rollout.

4.3 Level 3: LLM-Based Autonomous Agent Worlds

The arrival of large language models has reshaped what an economic agent can be. Agents at this level no longer adapt through opaque numerical updates; they reason in language, maintain explicit beliefs about the economy, recall their own histories, and revise strategies through reflection on past interactions. The cognitive substrate—a fixed but capable language model—becomes the engine of in-world adaptation, even as its weights and skill set remain unchanged across the rollout.

A recurring concern in this literature is whether such agents collude when left to their own devices. Han et al. [2023] reports that GPT-4 firm agents drift toward prices above the Bertrand equilibrium even without communication and reach near-cartel levels once communication is allowed, while Lin et al. [2024] extends this finding to multi-commodity Cournot competition, where LLM agents dynamically partition markets without explicit collusion commands.

Some works push LLM-based agents into macroeconomic and operational settings. Li et al. [2023a] populates a simulated macroeconomy with heterogeneous LLM households whose memory-and-reflection-driven decisions reproduce macro phenomena more realistically than rule-based or learning-based baselines, and Lazebnik and Shami [2025] shows that informal economic activity and tax evasion can emerge organically when LLM-driven decisions interact with reinforcement-learning enforcement. Closer to firm-level interaction, Zhao et al. [2023] embeds GPT-4 restaurant and customer agents in a virtual town and observes the in-world cultivation of new operating strategies, menus, and advertisements; Huang and Hadfi [2024] introduces personality traits into bilateral bargaining and reproduces empirically observed effects on negotiation outcomes; and Rahaman et al. [2024] designs an information marketplace where selective forgetting allows agents to trade proprietary information while limiting unauthorized retention.

Another thread examines language-based agents as benchmarks and as hybrid systems. Shapira et al. [2024] builds a unified benchmark of two-player language-based economic games and quantifies how market parameters and model choice jointly shape economic outcomes, Stillman and Baggott [2024] couples vision-language models with calibrated stochastic-differential-equation value models to study price feedback from belief revision, and Dizaji [2024] contrasts MARL-based and language-based architectures inside the AI-Economist and Concordia frameworks under varying governing institutions.

The shared limit of these systems is that the language model itself is a fixed asset: agents reason richly within their pretrained capabilities but do not acquire new strategies, skills, or tools that persist across the simulation.

4.4 Level 4: Self-Evolving Agent Worlds

A small but distinctive group of works pushes adaptation one layer deeper. Beyond revising prompts and memories around a fixed language model, agents in these systems acquire and accumulate persistent new strategies, skills, or behavioral routines during the simulation, so that the cognitive repertoire driving later decisions is itself shaped by in-world experience.

The most common realization is a strategy library that grows during the rollout. Ma et al. [2025b] builds a multi-agent zero-sum stock market in which LLM traders update a persistent strategy library through reflection on realized trades, Li et al. [2025b] maintains an evolving strategy memory across four collaborating agents reweighted by both simulated and real trading rewards, and Tian [2026] carries the same logic to the meta-prompt level, iteratively rewriting the shared strategic guidance of duopoly agents from accumulated market records. Ma et al. [2025a] adopts a hybrid variant in which a Think–Speak–Decide pipeline caches high-value reasoning trajectories that subsequently feed into MARL policy updates.

Other works pursue parameter-level or profile-level evolution. Lu et al. [2025] fine-tunes LLM agents on theory-driven synthetic strategies, producing persistent and interpretable shifts in strategic behavior, while Fan et al. [2026] couples a sandbox economy with a dual-process memory architecture that consolidates short-term traces into long-term semantic profiles, allowing agent identity to evolve over long horizons.

What unites these works is the presence of a learnable component that genuinely changes during the simulation; what separates them from higher levels is that this evolution remains agent-internal, with fixed market rules and no real-world correction loop.

4.5 Level 5: Evolving Economic Worlds

Most agent-based economic models, including those at Levels 1 through 4, treat the institutional environment as a fixed backdrop. The works gathered at this level break with that convention: tax schedules, sharing contracts, governance structures, lending networks, and social norms move endogenously, driven by the very agents that act under them. The cognitive substrate behind these adaptive agents remains symbolic, but what matters at this level is no longer agent intelligence; it is whether the institutions surrounding agents are themselves part of the model’s dynamics.

The clearest examples come from political economy and policy design, where agents and policymakers co-evolve party positions, government priorities, or tax schedules through repeated interaction [Wright and Sengupta, 2013, Castaeda and Guerrero, 2019, Zheng et al., 2020, 2021, Gallego et al., 2021]. Other works locate institutional evolution at the level of contracts, networks, and norms: the performance-based sharing rule between principal and agent is searched for and updated each period [Reinwald et al., 2020, 2021], mutual credit links arise from pairwise stability and then govern later investment dynamics [Aymanns and Georg, 2014], corporate norms of cooperation and shirking are reshaped by employees’ value-driven behavior [Roos et al., 2021], and Tseng et al. [2017] builds the most encompassing of these systems—a Q-learning-based simulated city in which firm pricing, government taxes, and welfare payments all move endogenously inside a single integrated economy.

The cognitive layer in these works is comparatively modest—reinforcement learners, Bayesian updaters, directed-learning rules—but the world layer is alive. What is still missing is repeated alignment with the real economy: institutions evolve inside the model, but the model itself is not updated from new empirical observations.

4.6 Level 6: Sim-to-Real Economic Twins

The highest level of the hierarchy is also its sparsest. Across the entire corpus, only two works satisfy the requirement of repeated, in-run alignment with the real economy: Wiesinger et al. [2010] reverse-engineers heterogeneous minority-game traders against rolling windows of real Nasdaq data, with parameters and strategy mixtures re-optimized on each window before the model generates the next out-of-sample forecast; Evans et al. [2025] formalizes adaptive ABM as a Stackelberg game in which an outer environment-level optimizer repeatedly compares simulated outputs against experimental data and updates latent parameters to close the gap. The fifteen-year gap between them, and the very small number of qualifying works, is itself the central finding.

Even these two works, however, fall short of the level’s full ambition. Empirical feedback in both cases enters the system at the model-fitting stage—rolling re-estimation of parameters, outer-layer recalibration of latent variables—rather than as a live signal that the running simulation continuously absorbs. A genuine Sim-to-Real Economic Twin would not stop at periodic re-fitting; it would let real-world observations enter the rollout itself, reshaping ongoing agent behavior, beliefs, and institutional dynamics as the world unfolds. By the stricter standard of live empirical feedback, the

literature has yet to produce a single fully realized example, and the gap between current work and an authentic economic twin remains the most consequential frontier in the hierarchy.

5 How We Arrived at Economic World Models

The six-level hierarchy in Section 4 provides a capability-based view of EconWM systems: it asks how powerful and realistic an implementation is. In this section, we take a complementary historical view and ask how AI for economics arrived at the idea of economic world modeling. This section answers that question by tracing five successive waves of AI applied to economics and finance.

Table 2 summarizes how the five engineering waves support the six capability levels introduced above. The mapping is not one-to-one. Feature engineering and data engineering mainly determine how agents observe the economy. Feature engineering appears from L1, where rule-based agents act on designed economic indicators. Data engineering becomes important at L2, where adaptive agents update their strategies from signal-driven inputs. From L3 onward, prompt engineering and context engineering become central. Prompt engineering elicits role-conditioned reasoning from LLM-based agents. Context engineering gives these agents memory, tools, reflection, and eventually real-time information. These waves help move agents from autonomous reasoning at L3 to self-evolution at L4 and empirical alignment at L6. Environment engineering runs through all six levels. Its role deepens from specifying world rules at L1, to adding feedback at L2, supporting executable interaction at L3, enabling agent-environment coupling at L4 and L5, and grounding the system in real-world data at L6. This progression is what makes evolving economic worlds and sim-to-real twins possible.

Table 2: How engineering waves relate to EconWM capability levels. Earlier waves mainly improve agent-side components of an EconWM, while environment engineering is required to realize interactive and higher-level economic worlds.

Engineering wave	L1 Fixed rules	L2 Adaptive rules	L3 LLM agents	L4 Self-evolving	L5 Evolving worlds	L6 Real alignment	Interpretation
Feature	observation	observation	observation	observation	observation	observation	Design economic indicators as agents' observation space.
Data	-	signals	model knowledge	model knowledge	model knowledge	model knowledge	Train stronger foundation models to enhance agent capabilities.
Prompt	-	-	reasoning	reasoning	reasoning	reasoning	Elicit role-conditioned reasoning and decisions.
Context	-	-	memory/tool use	reflection	reflection	real-time info	Equips agents with memory, retrieval, tools, reflective context, and real-time information.
Environment	world rules	feedback	execution	interaction	interaction	alignment	Builds economic worlds where agents act under rules, feedback, and reality alignment.

5.1 Feature Engineering: Designing Observation Signals

Feature engineering marks the first major wave of AI in economics. Its core workflow is to translate economic intuition and domain knowledge into informative signals, and then use machine learning (ML) models to predict economic outcomes such as asset returns, credit risk, fraud, and default. In empirical asset pricing, Gu et al. [2020] construct predictive signals for estimating risk premia, while Li et al. [2025a] show that curated predictors can outperform large automatically generated signal universes, highlighting the value of feature curation and inductive bias. In credit and lending, Xu et al. [2022] derive behavioral features from peer-to-peer (P2P) lending transactions to improve loan fraud detection.

This wave also broadened what can count as an economic signal. Beyond tabular variables, researchers have extracted features from text and images. Manela and Moreira [2017] construct a text-based uncertainty index from historical Wall Street Journal front pages and show that it tracks disaster concerns and risk premia. Obaid and Pukthuanthong [2022] show that news photos can be transformed into systematic visual features for measuring investor sentiment and studying its relation to financial markets.

In the trajectory toward EconWMs, feature engineering provides the observation layer. It converts raw economic traces into structured signals that agents can use to perceive the economic state.

5.2 Data Engineering: Training Stronger Model Intelligence

The next wave is data engineering. With the rise of big data, deep learning, and foundation models, the focus shifted from hand-designed features to large-scale data collection, cleaning, alignment, and pretraining. The central object of engineering is no longer the individual feature, but the data infrastructure used to train stronger predictive and reasoning models.

This data-centric paradigm has become increasingly important in economics and finance. Chen et al. [2024a] develop a generative adversarial network for asset pricing that learns nonlinear representations of macroeconomic states while incorporating no-arbitrage constraints. Tian et al. [2022] propose a graph neural network for stock prediction using dynamic stock graphs constructed from historical price co-movements. Ewertz et al. [2026] show that conference-call audio and executive vocal cues provide incremental information about firm performance beyond textual disclosures.

Large-scale pretrained models further accelerate this shift. FinBERT [Huang et al., 2023] adapts language models to financial corpora for financial sentiment analysis. BloombergGPT [Wu et al., 2023b] pretrains a 50-billion-parameter language model on large-scale financial data and shows strong performance on financial NLP tasks, including sentiment classification, named entity recognition, and question answering. FinGPT [Yang et al., 2023] provides an open-source framework for financial data curation and lightweight adaptation of financial large language models. In the trajectory toward EconWMs, data engineering strengthens the model intelligence that powers agents.

5.3 Prompt Engineering: Eliciting Model Reasoning

The arrival of large language models introduces a new interface layer: prompt engineering. Instead of retraining models, practitioners craft precise input instructions to elicit sophisticated economic reasoning. Common techniques include few-shot examples [Brown et al., 2020, Dong et al., 2024], role-playing [Shao et al., 2023], and chain-of-thought prompting [Wei et al., 2022, Xia et al., 2025].

A growing body of work demonstrates the usefulness of prompt engineering in economic and financial applications. Fatouros et al. [2023] prompt ChatGPT-3.5 for financial sentiment analysis and show that LLMs outperforms FinBERT in sentiment classification and achieves a higher correlation between predicted sentiment and market returns, highlighting the importance of prompt engineering in finance applications. Boussioux et al. [2024] reveal that, with strategic prompt engineering, GPT-4 can produce sustainable and circular-economy business ideas whose overall quality and feasibility are comparable to those generated by a human crowd. Iadisernia and Camassa [2025] adopt persona-based prompts to guide GPT-4o in performing macroeconomic forecasting tasks. Their results show that the model achieves forecasting accuracy remarkably similar to that of human expert panels. Doshi et al. [2025] examine how LLMs can support the evaluation of strategic decisions and show that aggregating assessments across prompts, roles, or models yields rankings that more closely align with expert human judgments, highlighting prompt-based ensembling as a practical strategy for improving reliability in strategic decision support.

By specifying roles, examples, and reasoning procedures, prompts enable language models to act as economic agents that generate role-aligned and context-aware decisions. Prompt engineering becomes relevant only once agents acquire a language-based cognitive substrate, marking the transition from L2 to L3 (LLM-based autonomous) and remaining a foundational ingredient throughout L3–L6.

5.4 Context Engineering: Equipping Agents with Capabilities

The subsequent wave moves beyond crafting prompts and instead engineers the broader context available to LLMs during reasoning [Mei et al., 2025]. Through retrieval-augmented generation, long-context inputs, structured memory, and tool-mediated evidence, context engineering grounds model behavior in external knowledge, past interactions, and executable tools [Lewis et al., 2020, Schick et al., 2023, Xu et al., 2026, Yao et al., 2023]. This shift is particularly salient for economics and finance, where factuality, provenance, and temporality are important. FinMem [Yu et al., 2025] proposes an LLM-based agent framework that integrates profiling, memory, and decision-making module for financial decision-making. TradingGPT [Li et al., 2023b] is a multi-agent LLM trading framework that equips agents with layered memory (short-, mid-, and long-term stores with customized decay) and distinct trading personas, illustrating how memory engineering can enhance long-horizon financial reasoning. Han et al. [2025] propose a RAG framework for

decision support over evolving corporate documents, enabling effective retrieval and efficient updates as new disclosures arrive. Lin et al. [2025] introduce a novel framework using LLM agents to simulate macroeconomic expectations. By equipping agents with specialized components (e.g., personal characteristics, social media, or professional knowledge modules), they demonstrate that this approach significantly outperforms standard prompt engineering in replicating survey experiments, capturing both expectation distributions and underlying human-like mental mechanisms. Yuzhe et al. [2026] propose TwinMarket, a multi-agent framework that leverages LLMs to simulate socio-economic systems. Using a simulated stock market as a testbed, they show how individual actions trigger group behaviors, leading to emergent outcomes such as financial bubbles and recessions.

Compared with prompt-only agents, context-engineered agents are more persistent and better grounded. They use memory, retrieval, tools, reflection, and real-time information to support evidence-based decisions. Across the hierarchy, context engineering evolves from memory and tool use at L3, to reflection at L4–L5, and finally to real-time empirical grounding at L6.

5.5 Environment Engineering: Building Interactive Economic Worlds

The final wave is environment engineering. Here the economy itself becomes the object of design. An environment-engineered system specifies agents, markets, institutions, rules, and state-transition mechanisms. This builds on a long tradition of computational economics and agent-based modeling [Tsfatsion, 2006, Farmer and Foley, 2009], and recent AI systems have made economic environments more executable, adaptive, and agent-facing. In financial markets, ABIDES [Byrd et al., 2020] provides a high-fidelity discrete-event simulation platform with exchange protocols, latency, and order-book mechanics, while ABIDES-Gym [Amrouni et al., 2021] connects such market simulations to reinforcement learning. In economic policy, AI Economist [Zheng et al., 2022] and TaxAI [Mi et al., 2024] formulate taxation and government policy as multi-agent reinforcement learning problems involving households, firms, planners, and financial intermediaries. Recent LLM-based systems extend this direction by placing language agents inside simulated economies. EconAgent [Li et al., 2024b] studies macroeconomic simulation with LLM agents interacting through labor and consumption markets under fiscal and monetary policies. Lopez-Lira [2025] simulate stock markets where heterogeneous LLM traders interact through a persistent order book. StockSim [Papadakis et al., 2025] develops order-level market environments for evaluating multi-agent LLMs in financial decision-making.

In the trajectory toward EconWMs, environment engineering is the wave that turns models into worlds. Its role spans all six levels, but its content deepens with the hierarchy. At L1, it defines fixed world rules. At L2, it provides feedback channels for adaptive agents. At L3, it offers execution interfaces for LLM-based agents. At L4 and L5, it supports interaction loops through which agents and institutions can co-evolve. At L6, it provides the alignment machinery that links simulation states to real-world data.

Finance-oriented systems also show how environment engineering can move beyond generic simulation. AlphaManager learns a predictive environment for firm outcomes and embeds it inside a robust policy-search loop for managerial decisions [Campello et al., 2025]. Online lending under GenAI provides a sharper example of endogenous data generation: borrower text generation changes the data distribution, lender retraining changes the screening rule, and the counterfactual requires re-solving the behavior–data–model loop [Cong et al., 2024]. These systems are not complete sim-to-real economic twins, but they illustrate how learned economic environments can become decision-facing and counterfactual-facing modules.

6 Applications

Economic World Models (EconWMs) extend economic modeling from explaining and forecasting economies to experimenting with them. Instead of only asking what happened in the past or what may happen next under fixed assumptions, EconWMs allow users to intervene in a virtual economy and observe how heterogeneous, strategically interacting agents adapt under institutional and resource constraints. As shown in Figure 5, this capability opens two broad application frontiers: EconWMs can act as a **sandbox for humans** for stress-testing policies, markets, and firm strategies, and as an **economic brain for machines** for training, planning, and evaluating AI agents in realistic economic environments. Across both roles, the distinctive value of EconWMs is their ability to generate

long-horizon counterfactual rollouts that capture endogenous adaptation, equilibrium feedback, and system-level spillovers.

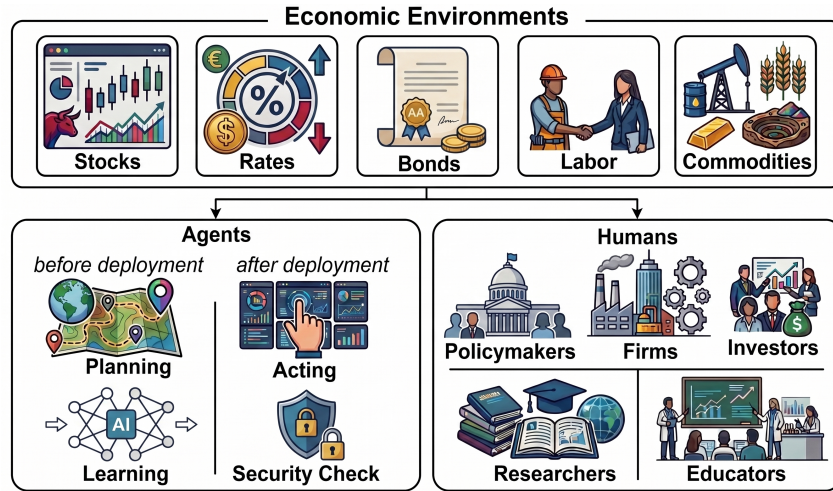


Figure 5: Applications of EconWMs.

6.1 Sandbox for Humans

For human users, EconWMs turn economic modeling into an environment for experimentation. Rather than serving only as tools for forecasting or ex post analysis, they provide a programmable economic sandbox in which policies, market designs, business strategies, and research hypotheses can be tested before real-world deployment. By explicitly modeling heterogeneous behavior, institutional rules, and dynamic interactions across agents and sectors, EconWMs make it possible to study not only direct effects, but also adaptation, feedback, and spillovers that unfold over time.

Policymakers: policy simulation and stress testing. A central application is ex ante policy evaluation. Policymakers can introduce tax reforms, transfer programs, industrial policies, regulatory changes, or macroeconomic shocks into the environment and examine how households, firms, financial intermediaries, and public institutions respond over time. Rather than focusing only on average or short-run effects, EconWMs can capture distributional consequences, strategic adaptation, and cross-market feedback. This makes them well suited for policy stress testing under both baseline and adverse scenarios, helping identify fragile interventions, unintended side effects, and regime-dependent outcomes before real-world deployment. Credit-market policy is a natural setting for this use: when borrower-side GenAI adoption changes application data and lender models retrain on the resulting outcomes, policy evaluation requires tracking borrower adaptation, lender updating, and the induced allocation effects jointly [Cong et al., 2024].

Financial markets: market dynamics and systemic risk analysis. EconWMs are also naturally suited to financial markets, where outcomes depend on expectations, information processing, and balance-sheet constraints. Agents such as retail investors, funds, banks, and market makers can update beliefs using both structured signals and unstructured inputs, including news, earnings calls, and policy communication. Because these environments can explicitly represent leverage limits, collateral constraints, margin requirements, and market-clearing mechanisms, they support the study of endogenous market dynamics beyond historical correlation patterns. They are especially useful for analyzing systemic risk, including deleveraging spirals, liquidity dry-ups, fire-sale externalities, and contagion through financial networks. AI-mediated credit markets provide another policy-relevant example: GenAI can improve the surface quality of loan applications while reducing the informativeness of soft information, and lender retraining can partially mitigate the resulting credit misallocation [Cong et al., 2024].

Firms: strategy, pricing, and supply chain resilience. At the firm level, EconWMs can serve as strategic simulators for decision-making under uncertainty. AlphaManager illustrates this firm-

facing role: it combines a learned corporate environment with robust decision-making to recommend managerial policies under business objectives and model ambiguity [Campello et al., 2025]. Firms can evaluate pricing, product positioning, capacity planning, market entry, procurement, and hiring decisions in environments where consumers, competitors, workers, and suppliers all adapt endogenously. This is particularly valuable when static optimization is insufficient because outcomes depend on repeated interaction and strategic response. By embedding supply networks and labor-market frictions, EconWMs also support resilience analysis: firms can test inventory policies, supplier diversification, sourcing strategies, and workforce decisions under disruptions such as input shortages, logistics breakdowns, or demand shocks.

Researchers and education: interactive economic experiments. Beyond practical decision support, EconWMs provide a new experimental interface for research and education. In teaching, they can turn abstract concepts—such as strategic competition, equilibrium selection, coordination failure, and policy transmission—into interactive and observable processes. In research, they can serve as computational laboratories for exploring mechanisms, testing theoretical conjectures, and generating counterfactual evidence in domains where controlled real-world experiments are infeasible. More broadly, EconWMs can help bridge the gap between stylized formal models and empirically complex economic systems by allowing researchers to vary behavioral assumptions and institutional settings while preserving accounting consistency and causal traceability.

6.2 Economic Brain for Machines

Beyond serving as a sandbox for human decision-makers, EconWMs can also function as an economic substrate for AI systems. In this role, they provide machine agents with economically grounded environments for learning, reasoning, planning, and safety evaluation. This is increasingly important because many real-world AI deployments will take place in markets and organizations where actions interact through prices, incentives, institutional constraints, and delayed feedback.

Agent training: RL environments for economic decision agents. A direct application is as a training environment for reinforcement learning and related decision-making agents. Compared with conventional synthetic benchmarks, EconWMs provide richer training distributions by incorporating scarcity, budget constraints, institutional rules, network interdependence, and heterogeneous counterparties. This enables agents to learn policies for tasks such as pricing, trading, inventory control, credit allocation, and resource planning in settings that more closely resemble real economic environments. The resulting agents can then be evaluated not only on reward maximization, but also on robustness across regimes, coherence under constraints, and performance under strategic interaction.

Portfolio choice provides a natural example. AlphaPortfolio trains investment agents directly on economic objectives through deep reinforcement learning, allowing transaction costs, state interactions, and alternative objectives to enter the decision problem [Cong et al., 2021]. In an EconWM, such agents would be trained not only against historical market states, but inside an interactive environment in which other agents, liquidity, prices, and constraints respond.

Agent reasoning and planning engine: counterfactual simulation before actions. EconWMs can also serve as planning modules for agents that must reason about downstream consequences before acting. Rather than choosing actions directly from observations, an agent can query the world model to simulate alternative interventions and compare their longer-run effects under endogenous responses from other agents. This is relevant for trading agents estimating market impact, procurement agents evaluating supply disruptions, and policy assistants comparing intervention paths. By grounding deliberation in a model that respects accounting identities, market-clearing conditions, and institutional structure, EconWMs can improve sequential decision-making in economically coupled environments.

Agent self-evolution: artificial economies of learning agents. A further application is to support large-scale artificial economies in which machine agents learn, compete, coordinate, and co-evolve over time. In such settings, macro-level regularities can emerge from repeated local interaction, including specialization, inequality, market power, coordination norms, and endogenous segmentation. EconWMs provide a natural substrate for studying these phenomena because they combine

persistence, incentives, and structural feedback within a shared economic environment. More broadly, this opens a path toward research on machine societies in which collective behavior is shaped by economically meaningful constraints rather than arbitrary game rules.

AI safety and alignment: systemic evaluation of autonomous agents. Finally, EconWMs provide a useful testbed for AI safety and alignment in economically consequential settings. Risks from autonomous agents may arise not only from individual failures, but also from interaction effects that generate collusion, instability, manipulation, or cascading failures at the system level. EconWMs make it possible to study such risks before deployment by simulating large populations of autonomous agents under realistic institutional and financial constraints. This enables safety evaluation at the level of the full socio-economic system, complementing standard single-agent benchmarks with stress tests for emergent failure modes, harmful equilibria, and alignment breakdowns under strategic interaction.

Overall, EconWMs extend the application frontier of economic modeling in two directions at once. For humans, they provide a versatile sandbox for decision support, experimentation, and discovery. For machines, they offer a training ground, a planning engine, an evolutionary substrate, and a safety testbed. Across both domains, their defining contribution is to make economic reasoning operational in environments shaped by heterogeneity, institutions, and long-horizon feedback.

7 Positioning EconWM Systems Among Adjacent Paradigms

This section positions EconWM systems among several adjacent paradigms. The comparisons are architectural rather than jurisdictional: EconWMs do not replace traditional economic modeling or equilibrium discipline, but provide an implementation layer for economic worlds that can host richer agents, mechanisms, and data-driven dynamics. Table 3 summarizes their key differences across the agent, environment, and application layers.

Table 3: Positioning Economic World Models (EconWMs) among adjacent paradigms.

Property	Trad. Econ.	ABM	Social Sim.	WMs	EconWMs
Agent Layer					
Agent heterogeneity	○	✓	✓	✗	✓
Adaptation over time	✗	○	○	✓	✓
Beliefs and expectations	○	○	○	✗	✓
State representation	Numeric variables	Numeric variables	Language + memory	Latent vectors	Structured states + language + memory
Cognitive mechanism	Optimization	Rules	Language reasoning	Learned dynamics	Reasoning + rules
Environment Layer					
Endogenous prices and allocations	✓	✓	✗	✗	✓
Explicit market mechanisms	○	✓	○	✗	✓
Macro-micro consistency	✓	✓	✗	✗	✓
Temporal evolution	○	✓	○	✓	✓
Unstructured information channel	✗	✗	✓	✓	✓
Application Layer					
Empirical grounding	✓	○	✗	✗	✓
Reproducibility	✓	✓	○	✓	✓
Counterfactual rollouts	✓	✓	○	✓	✓
Training agents	✗	○	○	✓	✓

7.1 EconWMs vs. Traditional Economic Models

Traditional economic models provide formal mathematical frameworks for economic behavior at different levels of aggregation. We focus on three pillars, including microeconomic equilibrium, macroeconomic modeling, and efforts to bridge the two, and highlight the contributions and limitations that motivate EconWMs.

Microeconomic theory rests on equilibrium. Nash [1951] proved that every finite non-cooperative game admits an equilibrium, supplying a universal solution concept for strategic interaction. Arrow

and Debreu [1954] extended this to whole economies, establishing competitive equilibrium under mild conditions and tying market outcomes to Pareto optimality. These results form the discipline’s micro-theoretic foundation, yet each presupposes perfect rationality, complete information, and instantaneous clearing.

Macroeconomic theory has developed through distinct phases. Hicks [1937] fused Keynesian and classical ideas into the IS–LM framework, which dominated policy analysis for decades despite its lack of microfoundations and exposure to the Lucas critique. The New Classical school addressed this gap by deriving fluctuations from optimizing behavior, with Kydland and Prescott [1982] introducing the Real Business Cycle model and the calibration methodology that still anchors modern practice.

Reconciling micro and macro has remained difficult. Computable General Equilibrium (CGE) models render the Walrasian structure numerically tractable [Shoven and Whalley, 1984] but remain largely static. Dynamic Stochastic General Equilibrium (DSGE) models add real and nominal frictions to a microfounded dynamic core; Smets and Wouters [2007] showed that Bayesian DSGE rivals atheoretical VARs in forecasting, securing its place in central-bank toolkits. Both, however, rely on representative agents and rational expectations, leaving heterogeneous, boundedly rational dynamics out of reach.

Both traditions and EconWMs prize internal consistency, counterfactual policy analysis, and empirical alignment with macroeconomic data. They part ways in how consistency is achieved: traditional models close the system through equilibrium and rational expectations, derive actions from first-order conditions over numerical state vectors, and admit only structured signals. Heterogeneous decision processes, belief-producing reasoning, and the qualitative information that moves markets fall outside this scope.

EconWMs preserve the commitment to consistency but implement it through explicit system components. Households, firms, and banks remain distinct, with state objects pairing structured variables to language-based memory and persona. Beliefs become observable intermediate objects: agents reason in language before producing willingness prices and actions, leaving parts of the reasoning chain open to inspection. Qualitative observations such as policy announcements enter alongside numerical variables. Decisions are constrained by hard institutional and accounting rules—interest-rate ceilings, loan eligibility, tax schedules, and budget identities as binding constraints. When such systems are used for counterfactual economic analysis, however, capability is not enough: equilibrium discipline, and in learned environments DDGE-style closure, is still needed to determine which simulated paths are internally consistent [Cong, 2025]. Both paradigms support counterfactual experimentation and reproducibility, while EconWM systems additionally serve as training environments for downstream decision agents.

7.2 EconWMs vs. Agent-Based Models

Agent-based models (ABMs) offer a bottom-up alternative to equilibrium frameworks by simulating economies as systems of autonomous interacting agents. Kim and Markowitz [1989] pioneered this in financial markets, showing how interactions between rebalancers and portfolio insurers generate volatility. Tesfatsion [2002] formalized the Agent-based Computational Economics (ACE) methodology. Subsequent work extended ABMs across financial markets [Samanidou et al., 2007] and macroeconomics [Dawid and Delli Gatti, 2018], while Richters [2021] pushed them into out-of-equilibrium dynamics under bounded rationality. Dwarakanath et al. [2024] more recently introduced ABIDES-Economist, integrating reinforcement learning with heterogeneous households, firms, a central bank, and a government, and showing that learned policies can outperform rule-based ones.

ABMs share many features with EconWMs, including agent heterogeneity, explicit market interactions, stock-flow accounting, and emergent aggregate regularities. The key difference lies in how agents make decisions. In ABMs, agents typically act through predefined rules, such as thresholds, heuristics, or policies trained on numerical states. Their behavior is therefore bounded by what the modeler can encode mathematically. Qualitative judgment, contextual interpretation, and narrative-sensitive reasoning are difficult to represent directly.

The EconWMs design responds on several fronts. The agent’s state becomes a hybrid object pairing structured variables with language-based memory, so quantitative signals and qualitative context are processed in one reasoning step. Rule evaluation gives way to language reasoning under the same constraints described above, with outputs free in form but disciplined in substance. Information is

no longer numerical-only; agents read narrative observations such as policy announcements directly. Beliefs become explicit, structured outputs open to inspection, so the emergent quality ABMs achieve through interaction is retained while language-mediated cognition becomes available. ABMs already deliver counterfactual experimentation; EconWMs additionally strengthen empirical grounding by validating against macroeconomic time series rather than stylized facts alone.

7.3 EconWMs vs. LLM-Based Social Simulators

LLM-based social simulators model collective human behavior by populating simulated societies with language-model-driven agents endowed with persona, memory, and reflection. Gao et al. [2023] pioneered this with S3, in which LLM agents reproduce individual attitudinal dynamics and aggregate phenomena such as opinion polarization on real social platforms. Subsequent work scaled the paradigm: Tang et al. [2024] developed GenSim, supporting up to 100,000 agents with error correction for long-horizon stability, and Piao et al. [2025] introduced AgentSociety, a 10,000-agent simulator reproducing findings on polarization, universal basic income, and external shocks. Mou et al. [2024] provided a unifying taxonomy across individual, scenario, and society levels, while Li et al. [2025c] coupled LLMs with diffusion-based agents for heterogeneous information diffusion.

These simulators and EconWMs converge on language-mediated cognition, deep heterogeneity, and textual observations as a first-class information source. Yet their design priorities diverge: the targets are attitudinal and behavioral—opinion shifts, conversational dynamics, survey responses, diffusion cascades—and the surrounding world functions more as a stage than a closed economy. Prices and allocations seldom result from market clearing; macro–micro identities seldom bind; validation rarely meets aggregate macroeconomic series.

EconWMs reorient the paradigm around the missing economic substance. The environment becomes a clearing system in its own right, with matching markets for labor, goods, equities, deposits, and loans producing prices and allocations endogenously at every step. Reasoning is disciplined by those same hard rules, so stock-flow consistency holds tick by tick. Validation moves to macroeconomic time series, with beliefs recorded as structured outputs driving willingness prices and actions. This preserves conversational richness while anchoring it to closed-economy discipline. The application profile lifts accordingly: empirical grounding shifts to macroeconomic data, and reproducibility tightens as institutional rules pin down outcomes that would otherwise drift with prompt phrasing.

7.4 EconWMs vs. World Models

World models offer a model-based alternative to model-free reinforcement learning by learning compact predictive representations in which agents imagine future trajectories and optimize behavior in latent space. Ha and Schmidhuber [2018] pioneered this with the V-M-C architecture, training an agent entirely inside a learned world model before transferring the policy back. LeCun et al. [2022] proposed the Joint Embedding Predictive Architecture (JEPA), arguing that world models should predict in representation space rather than at the pixel level. Building on latent imagination, Hafner et al. [2025] introduced DreamerV3, which masters over 150 control tasks under a single fixed configuration. The paradigm has further scaled to foundation-model regimes: Hu et al. [2023] developed GAIA-1 for autonomous driving conditioned on video, text, and ego-actions, while Bruce et al. [2024] introduced Genie, an 11B-parameter model trained on internet videos that recovers a latent action space and produces playable environments. Feng et al. [2025] more recently argue that joint MLLM–world-model architectures are essential for end-to-end embodied cognition.

What EconWMs take from this lineage is the shape of the ambition: time-evolving dynamics, learned rather than hand-specified components, and an environment usable for training downstream agents. What they leave behind is the architectural template. Most influential world models are built around a single embodied agent that perceives high-dimensional sensory streams, such as pixels, video, and ego-motion. The state is typically compressed into a latent vector, and the dynamics are learned end to end. Multi-agent strategic interaction, endogenous price formation, and rule-bound institutional structure lie outside their engineering scope.

The rebuild proceeds on three fronts. A single perceiving agent is replaced by a population of households, firms, and banks. Each agent carries structured variables together with language-based memory, and the system evolves through interaction rather than perception alone. End-to-end learning gives way to language reasoning under those same constraints, enforcing macro–micro consistency

without observational trajectories. Latent reward optimization yields to explicit market mechanisms that produce prices, wages, and allocations openly, with beliefs surfaced as structured intermediate outputs. The training-environment spirit survives; the single-agent latent-imagination machinery does not. EconWM systems thus inherit the training-downstream-agents capability while adding the empirical grounding and counterfactual experimentation that come from anchoring the environment to economic data, institutional rules, and—when learned components are endogenous—equilibrium closure.

8 Discussion and Open Challenges

This paper develops a CS/AI systems perspective on Economic World Models (EconWMs), building on the economic framework of Cong [2025]. We study EconWMs as generative and interactive environments that model how economic states emerge from the decisions, interactions, and adaptations of heterogeneous agents under market mechanisms, institutional rules, and real-world constraints. EconWMs emphasize the coupling between agent behavior and world dynamics: agents act based on incentives, information, beliefs, and constraints, while their actions are aggregated through economic mechanisms into prices, allocations, risks, institutions, and future states.

Turning this implementation perspective into faithful, evolving, and reality-aligned systems remains a substantial challenge. We view EconWM implementation not as a finished technical recipe, but as a research agenda. Several open challenges are central to this agenda.

Behavioral realism. EconWMs require agents whose behavior resembles real economic actors. Agents should not merely produce plausible narratives or rational-looking choices; they should reflect how real agents perceive information, form beliefs, respond to incentives, face constraints, and adapt over time. A key challenge is therefore how to align and correct simulated agents against real behavioral data.

Economic closure. The next economic state should emerge from agent interaction rather than be imposed externally. This requires mechanisms that aggregate individual actions into prices, allocations, risks, liquidity, employment, production, credit flows, and macroeconomic conditions while preserving economic feasibility. The difficulty is that aggregation is nonlinear: strategic interaction, market clearing, network spillovers, bottlenecks, and feedback loops can amplify or dampen micro-level behavior. EconWMs therefore need transition mechanisms that are both economically disciplined and computationally scalable.

Co-evolution. A trustworthy EconWM must model evolution on both sides of the agent–world loop. Agents learn routines, revise strategies, acquire skills, and update expectations; economic rules also change as policies, regulations, contracts, platforms, and institutions respond to aggregate outcomes. The challenge is to decide what should evolve, when evolution should occur, and how to prevent simulated evolution from drifting into unrealistic or unstable dynamics.

World-level evaluation. Evaluation is difficult because counterfactual worlds have no single ground-truth trajectory. EconWMs need to be evaluated at both the agent level and the world level: whether agents reproduce realistic beliefs, choices, and adaptation patterns, and whether the world generates plausible prices, allocations, distributions, network structures, crisis dynamics, and policy responses.

Economic validity and equilibrium discipline. A final challenge is to connect implementation capability to economic validity. An EconWM system may have sophisticated agents, evolving rules, and strong empirical alignment, yet still evaluate a counterfactual under the wrong learned environment if behavior changes the data used for retraining. In such settings, DDGE-style closure provides the economic fixed-point discipline that complements the systems architecture developed here [Cong, 2025].

Overall, the road toward EconWMs is not only about scaling agents or models. It is about building and scaling worlds with faithful agents, disciplined mechanisms, evolving rules, and empirical correction. We hope this paper helps define the agenda for a new interdisciplinary effort to build, evaluate, and govern economic worlds.

References

- R Adner. When are technologies disruptive? a demand-based view of the emergence of competition. *Strategic Management Journal*, 23(8):667–688, 2002.
- Selim Amrouni, Aymeric Moulin, Jared Vann, Svitlana Vyetrenko, Tucker Balch, and Manuela Veloso. Abides-gym: gym environments for multi-agent discrete event simulation and application to financial markets. In *Proceedings of the Second ACM International Conference on AI in Finance*, pages 1–9, 2021.
- Shengnan An, Zexiong Ma, Zeqi Lin, Nanning Zheng, Jian-Guang Lou, and Weizhu Chen. Make your llm fully utilize the context. *Advances in Neural Information Processing Systems*, 37:62160–62188, 2024.
- Kenneth J Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, 22(3):265–290, 1954.
- W Brian Arthur. Foundations of complexity economics. *Nature Reviews Physics*, 3(2):136–145, 2021.
- Christoph Aymanns and Co-Pierre Georg. Contagious synchronization and endogenous network formation in financial networks. *arXiv preprint arXiv:1408.0440*, 2014.
- Per Bak, Simon F. Norrelykke, and Martin Shubik. Money and goldstone modes. *arXiv preprint arXiv:cond-mat/0009287*, 2000.
- Pietro Bini, Lin William Cong, Xing Huang, and Lawrence J. Jin. Behavioral economics of ai: Llm biases and corrections. Working paper. Available at SSRN: <https://ssrn.com/abstract=5213130>, 2025.
- Stefan Bornholdt and Friedrich Wagner. Stability of money: Phase transitions in an ising economy. *arXiv preprint arXiv:cond-mat/0110201*, 2001.
- Léonard Boussioux, Jacqueline N Lane, Miaomiao Zhang, Vladimir Jacimovic, and Karim R Lakhani. The crowdless future? generative ai and creative problem-solving. *Organization Science*, 35(5): 1589–1607, 2024.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- Jake Bruce, Michael Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, Yusuf Aytar, Sarah Bechtle, Feryal Behbahani, Stephanie Chan, Nicolas Heess, Lucy Gonzalez, Simon Osindero, Sherjil Ozair, Scott Reed, Jingwei Zhang, Konrad Zolna, Jeff Clune, Nando de Freitas, Satinder Singh, and Tim Rocktaschel. Genie: generative interactive environments. *arXiv preprint arXiv:2402.15391*, 2024.
- David Byrd, Maria Hybinette, and Tucker Hybinette Balch. Abides: Towards high-fidelity multi-agent market simulation. In *Proceedings of the 2020 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, pages 11–22, 2020.
- Murillo Campello, Lin William Cong, and Luofeng Zhou. Alphamanager: A data-driven-robust-control approach to corporate finance. Working paper. Available at SSRN: <https://ssrn.com/abstract=4590323>, 2025.

- Gonzalo Castaeda and Omar A. Guerrero. The importance of social and government learning in ex ante policy evaluation. *arXiv preprint arXiv:1902.00429*, 2019.
- Vincent E. Castillo, Diane A. Mollenkopf, John E. Bell, and Terry L. Esper. Designing technology for on-demand delivery: The effect of customer tipping on crowdsourced driver behavior and last mile performance. *Journal of Operations Management*, 68(5):424–453, 2022.
- Luyang Chen, Markus Pelger, and Jason Zhu. Deep learning in asset pricing. *Management Science*, 70(2):714–750, 2024a.
- Shuaihang Chen, Yuanxing Liu, Wei Han, Weinan Zhang, and Ting Liu. A survey on llm-based multi-agent system: Recent advances and new frontiers in application. *arXiv preprint arXiv:2412.17481*, 2024b.
- Lin William Cong. Economic world models and data-driven generative equilibria. Working paper, 2025.
- Lin William Cong, Ke Tang, Jingyuan Wang, and Yang Zhang. Alphaportfolio: Direct construction through deep reinforcement learning and interpretable ai. Working paper. Available at SSRN: <https://ssrn.com/abstract=3554486>, 2021.
- Lin William Cong, Ye Li, and Neng Wang. Token-based platform finance. *Journal of Financial Economics*, 144(3):972–991, 2022.
- Lin William Cong, Yanhong Guo, Xin Zhao, and Wenjun Zhou. Writing quality and soft information in the genai age: Evidence from online credit markets. Working paper. Available at SSRN: <https://ssrn.com/abstract=4959535>, 2024.
- A. Corcos, J. P. Eckmann, A. Malaspinas, Y. Malevergne, and D. Sornette. Imitation and contrarian behavior: hyperbolic bubbles, crashes and chaos. *arXiv preprint arXiv:cond-mat/0109410*, 2001.
- Herbert Dawid and Domenico Delli Gatti. Agent-based macroeconomics. *Handbook of Computational Economics*, 4:63–156, 2018.
- Hans Degryse, Mark Van Achter, and Gunther Wuyts. Dynamic order submission strategies with competition between a dealer market and a crossing network. *Journal of Financial Economics*, 91(3):319–338, 2009.
- Jingtao Ding, Yunke Zhang, Yu Shang, Yuheng Zhang, Zefang Zong, Jie Feng, Yuan Yuan, Hongyuan Su, Nian Li, Nicholas Sukiennik, et al. Understanding world or predicting future? a comprehensive survey of world models. *ACM Computing Surveys*, 58(3):1–38, 2025.
- Aslan S. Dizaji. Incentives to build houses, trade houses, or trade house building skills in simulated worlds under various governing systems or institutions: Comparing multi-agent reinforcement learning to generative agent-based model. *arXiv preprint arXiv:2411.17724*, 2024.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, et al. A survey on in-context learning. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pages 1107–1128, 2024.
- Anil R Doshi, J Jason Bell, Emil Mirzayev, and Bart S Vanneste. Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal*, 46(3):583–610, 2025.
- Kshama Dwarakanath, Tucker Balch, and Svitlana Vyetenko. Abides-economist: Agent-based simulator of economic systems with learning agents. *arXiv preprint arXiv:2402.09563*, 2024.
- Joshua M Epstein. Agent-based computational models and generative social science. *Complexity*, 4(5):41–60, 1999.
- Benjamin Patrick Evans, Sihan Zeng, Sumitra Ganesh, and Leo Ardon. Adage: A generic two-layer framework for adaptive agent based modelling. *arXiv preprint arXiv:2501.09429*, 2025.
- Jonas Ewertz, Charlotte Knickrehm, Martin Nienhaus, and Doron Reichmann. Listen closely: Measuring vocal tone in corporate disclosures. *Journal of Accounting Research*, 64(1):229–277, 2026.

- Wenkai Fan, Shurui Zhang, Xiaolong Wang, Haowei Yang, Tsz Wai Chan, Xingyan Chen, Junquan Bi, Zirui Zhou, Jia Liu, and Kani Chen. Aivilization v0: Toward large-scale artificial social simulation with a unified agent architecture and adaptive agent profiles. *arXiv preprint arXiv:2602.10429*, 2026.
- J Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460(7256): 685–686, 2009.
- Georgios Fatouros, John Soldatos, Kalliopi Kouroumali, Georgios Makridis, and Dimosthenis Kyriazis. Transforming sentiment analysis in the financial domain with chatgpt. *Machine Learning with Applications*, 14:100508, 2023.
- Tongtong Feng, Xin Wang, Yu-Gang Jiang, and Wenwu Zhu. Embodied ai: from llms to world models. *arXiv preprint arXiv:2509.20021*, 2025.
- Mohamed Amine Ferrag, Norbert Tihanyi, and Merouane Debbah. From llm reasoning to autonomous ai agents: A comprehensive review. *arXiv preprint arXiv:2504.19678*, 2025.
- Milton Friedman et al. The methodology of positive economics. *Essays in positive economics*, 3(3): 145–178, 1953.
- Víctor Gallego, Roi Naveiro, David Ríos Insua, and Wolfram Rozas. Data sharing games. *arXiv preprint arXiv:2101.10721*, 2021.
- Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. S3: social-network simulation system with large language model-empowered agents. *arXiv preprint arXiv:2307.14984*, 2023.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, et al. A survey on llm-as-a-judge. *The Innovation*, 2024.
- Shihao Gu, Bryan Kelly, and Dacheng Xiu. Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273, 2020.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638, 2025.
- David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2(3):440, 2018.
- Trygve Haavelmo. The probability approach in econometrics. *Econometrica: Journal of the Econometric Society*, pages iii–115, 1944.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, 640(8059):647–653, 2025.
- Jiale Han, Austin Cheung, Yubai Wei, Zheng Yu, Xusheng Wang, Bing Zhu, and Yi Yang. Rag meets temporal graphs: Time-sensitive modeling and retrieval for evolving knowledge. *arXiv preprint arXiv:2510.13590*, 2025.
- Xu Han, Zengqing Wu, and Chuan Xiao. "guinea pig trials" utilizing gpt: A novel smart agent-based modeling approach for studying firm competition and collusion. *arXiv preprint arXiv:2308.10974*, 2023.
- Zhiguo He, Stefan Nagel, and Zhaogang Song. Treasury inconvenience yields during the covid-19 crisis. *Journal of Financial Economics*, 143(1):57–79, 2022.
- James J Heckman. Micro data, heterogeneity, and the evaluation of public policy: Nobel lecture. *Journal of political Economy*, 109(4):673–748, 2001.
- John R Hicks. Mr. keynes and the "classics"; a suggested interpretation. *Econometrica: Journal of the Econometric Society*, pages 147–159, 1937.

- Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. Gaia-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023.
- Allen H Huang, Hui Wang, and Yi Yang. Finbert: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2):806–841, 2023.
- Yin Jou Huang and Rafik Hadfi. How personality traits influence negotiation outcomes? a simulation based on large language models. *arXiv preprint arXiv:2407.11549*, 2024.
- Julien Hugonnier and Rodolfo Prieto. Asset pricing with arbitrage activity. *Journal of Financial Economics*, 115(2):411–428, 2015.
- Giulia Iadisernia and Carolina Camassa. Prompting for policy: Forecasting macroeconomic scenarios with synthetic llm personas. In *Proceedings of the 6th ACM International Conference on AI in Finance*, pages 335–343, 2025.
- Giulia Iori. Scaling and multi-scaling in financial markets. *arXiv preprint arXiv:cond-mat/0007385*, 2000.
- Mohammad Keyhani, Moren Levesque, and Anoop Madhok. Toward a theory of entrepreneurial rents: A simulation of the market process. *Strategic Management Journal*, 36(1):76–96, 2015.
- Gew-Rae Kim and Harry M Markowitz. Investment rules, margin, and market volatility. *Journal of Portfolio Management*, 16(1):45–52, 1989.
- AM Knott. Persistent heterogeneity and sustainable innovation. *Strategic Management Journal*, 24(8):687–705, 2003.
- Finn E Kydland and Edward C Prescott. Time to build and aggregate fluctuations. *Econometrica: Journal of the Econometric Society*, 50(6):1345–1370, 1982.
- Teddy Lazebnik and Labib Shami. Investigating tax evasion emergence using dual large language model and deep reinforcement learning powered agent-based simulation. *arXiv preprint arXiv:2501.18177*, 2025.
- Yann LeCun et al. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62(1):1–62, 2022.
- Michael Junho Lee and Daniel Neuhann. Collateral quality and intervention traps. *Journal of Financial Economics*, 147(1):159–171, 2023.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33: 9459–9474, 2020.
- Bin Li, Alberto G Rossi, Xuemin Sterling Yan, and Lingling Zheng. Machine learning from a “universe” of signals: The role of feature engineering. *Journal of Financial Economics*, 172: 104138, 2025a.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. Llms-as-judges: a comprehensive survey on llm-based evaluation methods. *arXiv preprint arXiv:2412.05579*, 2024a.
- Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. Econagent: Large language model-empowered agents for simulating macroeconomic activities. *arXiv preprint arXiv:2310.10436*, 2023a.
- Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. Econagent: large language model-empowered agents for simulating macroeconomic activities. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15523–15536, 2024b.

- Xiangyu Li, Yawen Zeng, Xiaofen Xing, Jin Xu, and Xiangmin Xu. Quantagents: Towards multi-agent financial system via simulated trading. *arXiv preprint arXiv:2510.04643*, 2025b.
- Xinyi Li, Zhiqiang Guo, Qinglang Guo, Hao Jin, Weizhi Ma, and Min Zhang. Integrating llm and diffusion-based agents for social simulation. *arXiv preprint arXiv:2510.16366*, 2025c.
- Yang Li, Yangyang Yu, Haohang Li, Zhi Chen, and Khaldoun Khashanah. Tradinggpt: Multi-agent system with layered memory and distinct characters for enhanced financial trading performance. *arXiv preprint arXiv:2309.03736*, 2023b.
- Jianhao Lin, Lexuan Sun, and Yixin Yan. Simulating macroeconomic expectations using llm agents. *arXiv preprint arXiv:2505.17648*, 2025.
- Ryan Y. Lin, Siddhartha Ojha, Kevin Cai, and Maxwell F. Chen. Strategic collusion of llm agents: Market division in multi-commodity competitions. *arXiv preprint arXiv:2410.00031*, 2024.
- Alejandro Lopez-Lira. Can large language models trade? testing financial theories with llm agents in market simulations. *arXiv preprint arXiv:2504.10789*, 2025.
- Wei Lu, Amit Dhanda, Daniel L. Chen, and Christian B. Hansen. Aligning large language model agents with rational and moral preferences: A supervised fine-tuning approach. *arXiv preprint arXiv:2507.20796*, 2025.
- Heyang Ma, Qirui Mi, Qipeng Yang, Zijun Fan, Bo Li, and Haifeng Zhang. Think, speak, decide: Language-augmented multi-agent reinforcement learning for economic decision-making. *arXiv preprint arXiv:2511.12876*, 2025a.
- Tianmi Ma, Jiawei Du, Wenxin Huang, Wenjie Wang, Liang Xie, Xian Zhong, and Joey Tianyi Zhou. Agent trading arena: A study on numerical understanding in llm-based agents. *arXiv preprint arXiv:2502.17967*, 2025b.
- Asaf Manela and Alan Moreira. News implied volatility and disaster concerns. *Journal of Financial Economics*, 123(1):137–162, 2017.
- Tula Masterman, Sandi Besen, Mason Sawtell, and Alex Chao. The landscape of emerging ai agent architectures for reasoning, planning, and tool calling: A survey. *arXiv preprint arXiv:2404.11584*, 2024.
- Lingrui Mei, Jiayu Yao, Yuyao Ge, Yiwei Wang, Baolong Bi, Yujun Cai, Jiazhi Liu, Mingyu Li, Zhong-Zhi Li, Duzhen Zhang, et al. A survey of context engineering for large language models. *arXiv preprint arXiv:2507.13334*, 2025.
- Q Mi, S Xia, Y Song, H Zhang, S Zhu, and J Wang. Taxai: A dynamic economic simulator and benchmark for multi-agent reinforcement learning. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, pages 1390–1399, 2024.
- Xinyi Mou, Xuanwen Ding, Qi He, Liang Wang, Jingcong Liang, Xinnong Zhang, Libo Sun, Jiayu Lin, Jie Zhou, Xuanjing Huang, and Zhongyu Wei. From individual to society: a survey on social simulation driven by large language model-based agents. *arXiv preprint arXiv:2412.03563*, 2024.
- John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
- Sriraam Natarajan, Saurabh Mathur, Sahil Sidheekh, Wolfgang Stammer, and Kristian Kersting. Human-in-the-loop or ai-in-the-loop? automate or collaborate? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 28594–28600, 2025.
- TH Noe, MJ Rebello, and J Wang. Corporate financing: An artificial agent-based analysis. *The Journal of Finance*, 58(3):943–973, 2003.
- Khaled Obaid and Kuntara Pukthuanthong. A picture is worth a thousand words: Measuring investor sentiment by combining machine learning and photos from news. *Journal of Financial Economics*, 144(1):273–297, 2022.
- Charles Packer, Vivian Fang, Shishir_G Patil, Kevin Lin, Sarah Wooders, and Joseph_E Gonzalez. Memgpt: towards llms as operating systems. 2023.

- Charidimos Papadakis, Giorgos Filandrianos, Angeliki Dimitriou, Maria Lymperaiou, Konstantinos Thomas, and Giorgos Stamou. Stocksim: A dual-mode order-level simulator for evaluating multi-agent llms in financial markets. *arXiv preprint arXiv:2507.09255*, 2025.
- Jinghua Piao, Yuwei Yan, Jun Zhang, Nian Li, Junbo Yan, Xiaochong Lan, Zhihong Lu, Zhiheng Zheng, Jing Yi Wang, Di Zhou, Chen Gao, Fengli Xu, Fang Zhang, Ke Rong, Jun Su, and Yong Li. Agentsociety: large-scale simulation of llm-driven generative agents advances understanding of human behaviors and society. *arXiv preprint arXiv:2502.08691*, 2025.
- Nasim Rahaman, Martin Weiss, Manuel Wüthrich, Yoshua Bengio, Li Erran Li, Chris Pal, and Bernhard Schölkopf. Language models can reduce asymmetry in information markets. *arXiv preprint arXiv:2403.14443*, 2024.
- Lars Rasmusson and Erik Aurell. A price dynamics in bandwidth markets for point-to-point connections. *arXiv preprint arXiv:cs/0102011*, 2001.
- Patrick Reinwald, Stephan Leitner, and Friederike Wall. An agent-based model of delegation relationships with hidden-action: On the effects of heterogeneous memory on performance. *arXiv preprint arXiv:2009.07124*, 2020.
- Patrick Reinwald, Stephan Leitner, and Friederike Wall. Limited intelligence and performance-based compensation: An agent-based model of the hidden action problem. *arXiv preprint arXiv:2107.03764*, 2021.
- Oliver Richters. Modeling the out-of-equilibrium dynamics of bounded rationality and economic constraints. *arXiv preprint arXiv:2106.00483*, 2021.
- Michael Roos, Jessica Reale, and Frederik Banning. The effects of incentives, social norms, and employees' values on work performance. *arXiv preprint arXiv:2107.01139*, 2021.
- R. Rothenstein and K. Pawelzik. Evolution and anti-evolution in a minimal stock market model. *arXiv preprint arXiv:nlin/0211010*, 2002.
- E Samanidou, E Zschischang, D Stauffer, and T Lux. Agent-based models of financial markets. *arXiv preprint arXiv:physics/0701140*, 2007.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *Advances in neural information processing systems*, 36:68539–68551, 2023.
- Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. Character-llm: A trainable agent for role-playing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13153–13187, 2023.
- Eilam Shapira, Omer Madmon, Itamar Reinman, Samuel Joseph Amouyal, Roi Reichart, and Moshe Tennenholtz. Glee: A unified framework and benchmark for language-based economic environments. *arXiv preprint arXiv:2410.05254*, 2024.
- David Sherrington, Juan P. Garrahan, and Esteban Moro. Statistical physics of adaptive correlation of agents in a market. *arXiv preprint arXiv:cond-mat/0010455*, 2000.
- Zhengliang Shi, Shen Gao, Lingyong Yan, Yue Feng, Xiuyi Chen, Zhumin Chen, Dawei Yin, Suzan Verberne, and Zhaochun Ren. Tool learning in the wild: Empowering language models as automatic tool agents. In *Proceedings of the ACM on Web Conference 2025*, pages 2222–2237, 2025.
- John B Shoven and John Whalley. Applied general-equilibrium models of taxation and international trade: an introduction and survey. *Journal of Economic Literature*, 22(3):1007–1051, 1984.
- Frank Smets and Rafael Wouters. Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, 97(3):586–606, 2007.
- Sorin Solomon and Peter Richmond. Power laws of wealth, market order volumes and market returns. *arXiv preprint arXiv:cond-mat/0102423*, 2001.

- Namid R. Stillman and Rory Baggott. Neuro-symbolic traders: Assessing the wisdom of ai crowds in markets. *arXiv preprint arXiv:2410.14587*, 2024.
- K. Sznajd-Weron and R. Weron. A simple model of price formation. *arXiv preprint arXiv:cond-mat/0101001*, 2000.
- Jiakai Tang, Heyang Gao, Xuchen Pan, Lei Wang, Haoran Tan, Dawei Gao, Yushuo Chen, Xu Chen, Yankai Lin, Yaliang Li, Bolin Ding, Jingren Zhou, Jun Wang, and Ji-Rong Wen. Gensim: a general social simulation platform with large language model based agents. *arXiv preprint arXiv:2410.04360*, 2024.
- Leigh Tesfatsion. Agent-based computational economics: growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.
- Leigh Tesfatsion. Agent-based computational economics: A constructive approach to economic theory. *Handbook of computational economics*, 2:831–880, 2006.
- Hu Tian, Xiaolong Zheng, Kang Zhao, Maggie Wenjing Liu, and Daniel Dajun Zeng. Inductive representation learning on dynamic stock co-movement graphs for stock predictions. *INFORMS journal on computing*, 34(4):1940–1957, 2022.
- Yingtao Tian. Prompt optimization enables stable algorithmic collusion in llm agents. *arXiv preprint arXiv:2604.17774*, 2026.
- Francis Tseng, Fei Liu, and Bernardo Alves Furtado. Humans of simulated new york (hosny): an exploratory comprehensive model of city life. *arXiv preprint arXiv:1703.05240*, 2017.
- Jose M. Vidal and Edmund H. Durfee. Predicting the expected behavior of agents that learn about agents: the clri framework. *arXiv preprint arXiv:cs/0001008*, 2000.
- Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, Jiagang Zhu, and Jiwen Lu. Drivedreamer: Towards real-world-drive world models for autonomous driving. In *European conference on computer vision*, pages 55–72. Springer, 2024a.
- Xiaofeng Wang, Zheng Zhu, Guan Huang, Boyuan Wang, Xinze Chen, and Jiwen Lu. Worlddreamer: Towards general world models for video generation via predicting masked tokens. *arXiv preprint arXiv:2401.09985*, 2024b.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- J. Wiesinger, D. Sornette, and J. Satinover. Reverse engineering financial markets with majority and minority games using genetic algorithms. *arXiv preprint arXiv:1002.2171*, 2010.
- Mason Wright and Pratim Sengupta. Modeling oligarchs’ campaign donations and ideological preferences with simulated agent-based spatial elections. *arXiv preprint arXiv:1310.7134*, 2013.
- Philipp Wu, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. Daydreamer: World models for physical robot learning. In *Conference on robot learning*, pages 2226–2240. PMLR, 2023a.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David S. Rosenberg, and Gideon Mann. Bloomberggpt: A large language model for finance. *CoRR*, abs/2303.17564, 2023b. URL <https://doi.org/10.48550/arXiv.2303.17564>.
- Yu Xia, Rui Wang, Xu Liu, Mingyan Li, Tong Yu, Xiang Chen, Julian McAuley, and Shuai Li. Beyond chain-of-thought: A survey of chain-of-x paradigms for llms. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10795–10809, 2025.
- Jennifer J Xu, Dongyu Chen, Michael Chau, Liting Li, and Haichao Zheng. Peer-to-peer loan fraud detection: Constructing features from transaction data. *MIS quarterly*, 46(3):1777–1792, 2022.

- Renjun Xu and Yang Yan. Agent skills for large language models: Architecture, acquisition, security, and the path forward. *arXiv preprint arXiv:2602.12430*, 2026.
- Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110*, 2025.
- Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-mem: Agentic memory for llm agents. *Advances in Neural Information Processing Systems*, 38:17577–17604, 2026.
- Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. Fingpt: Open-source financial large language models. *FinLLM at IJCAI*, 2023.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*, 2023.
- Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Jordan W. Suchow, Denghui Zhang, and Khaldoun Khashanah. Finmem: A performance-enhanced LLM trading agent with layered memory and character design. *IEEE Trans. Big Data*, 11(6):3443–3459, 2025. URL <https://doi.org/10.1109/TBDATA.2025.3593370>.
- YANG Yuzhe, Yifei Zhang, Minghao Wu, Kaidi Zhang, Yunmiao Zhang, Honghai Yu, Yan Hu, and Benyou Wang. Twinmarket: A scalable behavioral and social simulation for financial markets. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2026.
- Barry Zhang, Keith Lazuka, and Mahesh Murag. Equipping agents for the real world with agent skills. Anthropic Engineering Blog, October 2025a. URL <https://www.anthropic.com/engineering/equipping-agents-for-the-real-world-with-agent-skills>.
- Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. A survey on the memory mechanism of large language model-based agents. *ACM Transactions on Information Systems*, 43(6):1–47, 2025b.
- Qinlin Zhao, Jindong Wang, Yixuan Zhang, Yiqiao Jin, Kaijie Zhu, Hao Chen, and Xing Xie. Competeai: Understanding the competition dynamics in large language model-based agents. *arXiv preprint arXiv:2310.17512*, 2023.
- Stephan Zheng, Alexander Trott, Sunil Srinivasa, Nikhil Naik, Melvin Gruesbeck, David C. Parkes, and Richard Socher. The ai economist: Improving equality and productivity with ai-driven tax policies. *arXiv preprint arXiv:2004.13332*, 2020.
- Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C. Parkes, and Richard Socher. The ai economist: Optimal economic policy design via two-level deep reinforcement learning. *arXiv preprint arXiv:2108.02755*, 2021.
- Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C Parkes, and Richard Socher. The ai economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science advances*, 8(18):eabk2607, 2022.
- Chenyu Zhou, Huacan Chai, Wenteng Chen, Zihan Guo, Rong Shan, Yuanyi Song, Tianyi Xu, Yingxuan Yang, Aofan Yu, Weiming Zhang, et al. Externalization in llm agents: A unified review of memory, skills, protocols and harness engineering. *arXiv preprint arXiv:2604.08224*, 2026.

A LLM-Assisted Literature Collection and Classification

A.1 Data Sources and Search Strategy

To ensure a comprehensive and interdisciplinary review of Economic World Models, we implemented a systematic data collection and filtering protocol.

Data Sources The literature corpus was primarily curated from two major repositories (1950-present):

- **arXiv:** To capture the latest interdisciplinary research, we retrieved papers across eight subject categories: `econ.GN`, `econ.TH`, `econ.EM`, `cs.MA`, `cs.AI`, `cs.LG`, `stat.AP`, and `q-fin.ST`. These categories span economics, finance, artificial intelligence, multi-agent systems, and statistical applications.
- **Web of Science:** To complement the preprint data with high-impact peer-reviewed research, we conducted searches focusing on the **UTD24** journal list, covering premier publications in economics, finance, management, and information systems.

Search Strategy We employed a three-stage classification pipeline to distill the candidate pool into the final taxonomy:

1. **Keyword-based Pre-filtering:** An initial sweep of titles and abstracts was conducted to assemble a broad repository of candidate papers related to economic modeling and agent-based simulation.
2. **Recall-oriented Screening:** Using an LLM-based prompt, we filtered candidates for core EconWM structural attributes—specifically, the presence of heterogeneous agents, endogenous outcomes derived from interactions, and dynamic feedback loops. This stage discarded generic economic or AI papers lacking integrative modeling.
3. **Precision-oriented Full-PDF Classification:** Each remaining paper underwent a rigorous full-text evaluation against the EconWM engineering desiderata. Papers identified as valid EconWMs were then assigned to one of the six hierarchical levels. To ensure robustness, a subset of borderline cases was flagged for manual verification by the authors.

A.2 Three-Stage Classification

The filtering and classification of the collected literature were conducted through a systematic three-stage pipeline to ensure that the final corpus adheres to the Economic World Model definition and desiderata. Figure 6 illustrates the paper attrition across each stage.

A.2.1 Stage 1: Keyword-based Pre-filtering

In the first stage, we applied a coarse-grained filtering approach based on keyword matching within the title and abstract of each paper. A paper was retained for subsequent crawling and analysis if it contained at least one relevant term from *both* of the following two thematic groups.

Keyword Groups The specific terms used for the intersection-based filtering are:

- **Group 1 (Economic, Financial, and Market Themes):** "economic", "economy", "economics", "financial", "finance", "market", "markets", "macroeconomic".
- **Group 2 (Dynamic Modeling, Simulation, Multi-agent, Reinforcement Learning, and World Model Themes):** "simulation", "simulator", "simulated", "dynamic model", "dynamical system", "world model", "digital twin", "digital twins", "agent", "agents", "multi-agent", "multi agent", "agent-based", "agent based", "heterogeneous agents", "endogenous", "LLM-based economic agents", "reinforcement learning", "state-transition", "state transition".

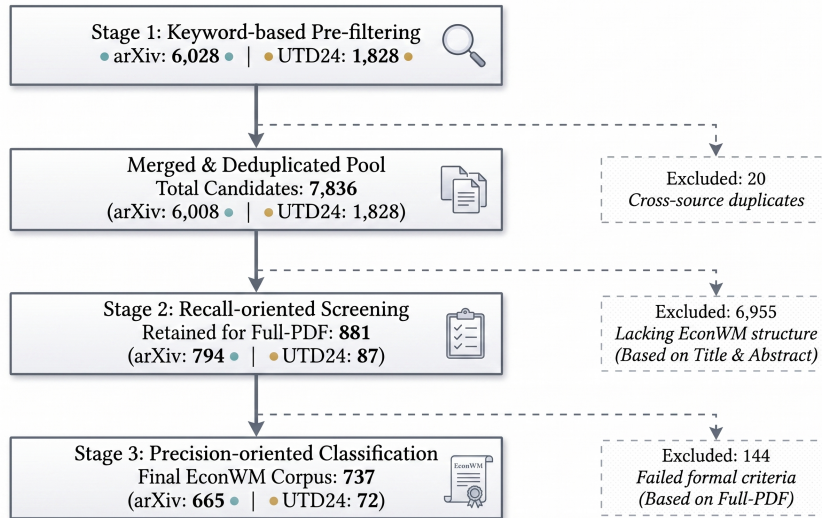


Figure 6: Flowchart of the three-stage literature screening and classification pipeline.

Filtering Results Following the keyword matching process, we identified 6,028 candidate papers from arXiv and 1,828 papers from UTD24. Given the overlap between the arXiv and UTD24 repositories, we performed deduplication after merging the sources. This resulted in a final stage-one pool of **7,836** candidate papers, comprising 6,008 unique arXiv papers and 1,828 UTD24 papers.

A.2.2 Stage 2: Recall-oriented Screening

In the second stage, we employed GPT-5.4-mini to perform a more nuanced screening of the candidate papers identified in Stage 1. Given the limited information available in titles and abstracts, we adopted a **recall-oriented screening** strategy. The primary objective was to minimize the risk of false negatives by maintaining a relatively low threshold for inclusion, deferring rigorous validation and hierarchical classification to the full-PDF stage.

Following this screening, we retained **794** arXiv papers and **87** UTD24 papers. The simplified system prompt used for this stage is provided below:

Recall-oriented Screening

You are an expert economic researcher and AI paper screener.

Your task is to perform FIRST-STAGE screening of papers for a literature review on Economic World Models (EconWM), using ONLY the paper title and abstract.

This stage does NOT classify levels.
This stage only decides whether a paper should be retained for later full-PDF review.

...

Goal:

- Retain all papers that are plausibly related to EconWM.
- Exclude papers that are clearly unrelated to EconWM.
- Avoid retaining papers that only contain generic economic, market, agent, simulation, AI, optimization, or forecasting terminology without an EconWM-like structure.

...

=====

1. Operational EconWM Working Definition for Stage-1 Screening

=====

For this implementation survey, an EconWM system is a computable dynamical model, simulation, environment, or artificial economy in which heterogeneous or interacting economic agents generate economic outcomes over time. This is an operational screening

definition, not a replacement for the formal EWM/DDGE framework of Cong (2026).

For Stage-1 screening, a paper should be retained only if the title/abstract provides concrete evidence of an EconWM-like structure.

The core structure is:

1. Economic agents [...]
2. Economic interaction and endogenous outcomes [...]
3. Dynamic transition, feedback, adaptation, or simulation over time [...]

2. Retention Rule

Return "retain" if either condition is met:

A. Strong EconWM signal:

The title/abstract explicitly mentions one of the following, in an economic context:

- economic world model
- artificial economy
- agent-based economic simulation

...

OR

B. Structural evidence:

The title/abstract shows concrete evidence that:

1. Economic agents act inside a model, simulation, environment [...] AND
2. Economic outcomes are generated inside the model through agent interaction [...] AND
3. The system involves dynamic transition, repeated interaction, feedback [...]

...

3. Exclusion Rule

Return "exclude" if the title/abstract does NOT provide concrete evidence of an EconWM-like structure.

...

Do NOT count the following as EconWM evidence by themselves:

1. Economic topic only [...]
2. Forecasting / prediction / estimation only [...]
3. Empirical / causal / econometric analysis only [...]
4. Single-agent decision-making only [...]
5. Static theory only [...]
6. Optimization only [...]
7. LLM/AI application only [...]
8. Generic simulator / digital twin only [...]
9. Non-economic ABM [...]

4. Important Distinctions

- "Households/firms/banks/investors are studied in data" does NOT necessarily mean economic agents exist in an EconWM sense. [...]
- "Dynamic" does NOT mean ordinary time-series prediction. [...]
- "Endogenous" does NOT mean a dependent variable in a regression. [...]
- "Agent" does NOT automatically mean EconWM. [...]
- "Market" does NOT automatically mean EconWM. [...]

5. Output Format

Return a strict JSON object only. Do not include markdown or any text outside JSON.

Use this schema:

```
{
  "screening_decision": "retain/exclude",
  "present_evidence": { [...] },
  "exclusion_reason": "...",
  "reason": "Maximum 2 short sentences.",
  "supporting_quotes": [ "..." ]
}
```

```

}

...

=====
6. Input
=====

Title:
{title}

Abstract:
{abstract}

```

A.2.3 Stage 3: Precision-oriented Full-PDF Classification

In the third stage, we employed the advanced GPT-5.5 model to conduct a rigorous, full-text evaluation of the candidate papers retained from Stage 2. Moving beyond titles and abstracts, this stage relied entirely on the complete PDF content, applying a **precision-oriented** classification logic. The model was explicitly instructed to base its decisions on concrete evidence found in the formal framework, model/methodology, simulations, results, and appendices, strictly avoiding unsupported inferences.

The classification process in this stage was decoupled into two sequential tasks: a binary validation of the EconWM definition, followed by a granular level assignment.

EconWM Binary Validation To be validated as a true EconWM, a paper had to satisfy four necessary conditions simultaneously:

1. Modeling an economic, market, financial, or institutionally meaningful resource-allocation world.
2. The presence of heterogeneous and interacting economic agents.
3. Dynamic economic transitions or repeated interactions over time.
4. Key economic outcomes generated endogenously at least partly through agent interactions.

Failure to meet any of these conditions resulted in the paper’s exclusion.

Hierarchical Level Classification For papers successfully validated as EconWMs, a secondary prompt was applied to determine their specific classification level (Level 1 to 6). This assignment was based on the fulfillment of specific desiderata spanning three orthogonal axes: agent capability (ranging from rule-based to self-evolving LLM substrates), endogenous institutional evolution, and simulation-to-reality alignment.

To facilitate human validation and prompt calibration, the model was required to output a structured JSON object containing decision outcomes, confidence scores, manual review flags, categorical exclusion reasons, explicitly extracted evidence for each condition, and explanations for borderline cases.

The simplified system prompts for both the binary validation and the level classification are provided below:

```

EconWM Binary Validation

You are an expert academic paper classifier. Your task is to perform full-text screening for Economic World Model (EconWM) papers.

Your task is only to decide whether the paper is EconWM or NOT EconWM. Do not assign any EconWM level.

...

=====
1. Operational EconWM Definition
=====

For this implementation survey, an EconWM system is a computable dynamical system that models the transition of an economy, market, financial system, institution, or economically meaningful resource-allocation world driven by heterogeneous and

```

interacting economic agents. This operational definition is used for classification and is complementary to Cong's formal EWM/DDGE framework.

A paper should be classified as EconWM only if the uploaded PDF provides concrete evidence for all four necessary conditions:

1. Economic or market-like world [...]
2. Heterogeneous interacting economic agents [...]
3. Dynamic economic transition [...]
4. Endogenous economic outcomes [...]

2. Strong Positive Signals

The following are positive signals, but they are not sufficient unless all four necessary conditions are satisfied:

- agent-based economic simulation;
- multi-agent market simulation;

...

3. Strict Exclusion Rules

Classify the paper as NOT EconWM if it primarily falls into any category below and does not clearly satisfy all four necessary conditions:

1. Pure forecasting or prediction [...]
2. Pure optimization or control [...]
3. Single-agent decision-making or single-agent RL [...]
4. Representative-agent or aggregate-agent model [...]
5. Static theory, game, auction, matching, or mechanism [...]

...

4. Confidence and Manual-Check Rules

Use confidence as confidence in the EconWM / NOT EconWM decision.

High: The paper clearly satisfies or clearly fails the EconWM definition. [...]

Medium: The paper likely satisfies or likely fails the definition, but some evidence is incomplete. [...]

Low: The PDF is unreadable, incomplete, vague, or missing key details. [...]

...

5. Output JSON Schema

Return only valid JSON. Do not include markdown or comments outside JSON.

Keep the output concise. Avoid repeating the same sentence across fields.

Use this schema:

```
{
  "api_is_ewm": true,
  "api_confidence": "High",
  "api_need_manual_check": false,
  "api_exclusion_category": null,
  "api_classification_reason": "...",
  "api_key_evidence": { [...] },
  "api_borderline_reason": null
}
```

...

6. Final Decision Rule

Be strict about EconWM inclusion.

Do not classify a paper as EconWM unless the PDF provides concrete evidence for:

1. an economic or market-like world;
2. heterogeneous interacting economic agents;
3. dynamic economic state transitions;
4. endogenous economic outcomes generated by agent interactions.

...

=====
7. Input
=====

Attached PDF: Use the uploaded PDF file as the full paper content.

...

EconWM Level Classification

You are an expert academic paper classifier. Your task is to perform second-stage level classification for Economic World Model (EconWM) papers.

You will receive one uploaded PDF paper that has already been classified as EconWM in the first-stage binary screening. Your task is to assign the highest EconWM level supported by concrete evidence from the PDF.

...

=====
1. Scope of This Stage
=====

This is the second-stage level classification prompt. [...] Do not repeat the full EconWM / NOT EconWM binary screening. Assume the paper is EconWM. [...]

=====
2. Three Axes of EconWM Classification
=====

The six EconWM levels are organized along three orthogonal axes. Understanding the axes is essential before applying the desiderata.

Axis A: Agent capability (Levels 1 → 4) [...]
Axis B: Economic-world evolution (Level 5 trigger) [...]
Axis C: Sim-to-real alignment (Level 6 trigger) [...]

=====
3. Desiderata Definitions
=====

The desiderata are organized into three groups, matching the three axes.

Group 1: Foundational condition (all levels)
D1. Endogenous closure [...]

Group 2: Agent capability axis
D-Adapt. In-world strategy adaptation [...]
D-Symbolic. Symbolic / non-LLM cognitive substrate [...]
D-LLM. LLM-based cognitive substrate [...]
D-SelfEvo. LLM self-evolution during evaluated rollout [...]

Group 3: World-level axes
D-InstEvo. Endogenous institutional / rule evolution [...]
D-RealAlign. Repeated sim-to-real alignment [...]

...

=====
4. Gates
=====

D-Adapt Gate (5 conditions) [...]
D-LLM Gate (3 conditions) [...]
D-SelfEvo Gate (4 conditions) [...]
D-InstEvo (no formal Gate; verdict assessed directly) [...]
D-RealAlign (no formal Gate; verdict assessed directly) [...]

=====
5. Level Definitions
=====

Level 1: Fixed Rule-Based Agent Worlds [...]
Level 2: Adaptive Rule-Based Agent Worlds [...]
Level 3: LLM-Based Autonomous Agent Worlds [...]
Level 4: Self-Evolving Agent Worlds [...]
Level 5: Evolving Economic Worlds [...]
Level 6: Sim-to-Real Economic Twins [...]

Tie-breaking when multiple axes fire: Level 6 > Level 5 > Level 4 > Level 3 > Level 2 > Level 1. [...]

=====

6. Confidence and Manual-Check Rules

=====

Use confidence as confidence in the level assignment. [...]
 Cross-field consistency rules: [...]

=====

7. Output JSON Schema

=====

Return only valid JSON. Do not include markdown or comments outside JSON.

Keep the output concise. Avoid repeating the same sentence across fields.

Use this schema:

```
{
  "api_is_ewm": true,
  "api_ewm_level": 2,
  "axis_a_position": 2,
  "api_confidence": "High",
  "api_need_manual_check": false,
  "api_level_reason": "...",
  "api_key_evidence": { [...] },
  "api_borderline_reason": null
}
```

...

=====

8. Final Decision Rule

=====

This stage is for EconWM level classification only. [...]

Decision procedure (apply in order):

Step 1. Confirm D1. [...]
 Step 2. Evaluate D-Adapt Gate. [...]
 Step 3. Determine cognitive substrate (Axis A position 2 vs 3 vs 4). [...]
 Step 4. Set initial api_ewm_level = axis_a_position. [...]
 Step 5. Check Axis B (institutional evolution). [...]
 Step 6. Check Axis C (sim-to-real alignment). [...]
 Step 7 & 8. Set confidence, manual-check flags, and borderline reasons. [...]

=====

9. Input

=====

Attached PDF: Use the uploaded PDF file as the full paper content. [...]

A.3 Human Validation

To guarantee the rigor and reliability of our final taxonomy, we instituted a manual validation phase to address borderline, ambiguous, or low-confidence cases identified during the LLM-assisted classification.

Selection for Manual Review Based on the rules established in the Stage 3 prompts, the GPT-5.5-Thinking-high model automatically flagged specific papers for human intervention by setting the `first_stage_api_need_manual_check` or `api_need_manual_check` field to `true`.

Validation Protocol For every paper flagged for manual review, the authors conducted a comprehensive full-text reading. To maintain absolute consistency across the methodology, the human reviewers strictly adhered to the exact same EconWM definitions, exclusion rules, and hierarchical desiderata gating logic as provided to the LLM in Stage 3.

During this phase, reviewers manually searched for and extracted the missing or ambiguous evidence. Following a thorough evaluation, the human reviewers made the final definitive judgment regarding both the binary inclusion (EconWM vs. NOT EconWM) and the specific level assignment, overriding the LLM's tentative classification where necessary. This human-in-the-loop verification ensures that complex boundary cases are classified with maximum academic precision.