

---

# AIVILIZATION v0: TOWARD LARGE-SCALE ARTIFICIAL SOCIAL SIMULATION WITH A UNIFIED AGENT ARCHITECTURE AND ADAPTIVE AGENT PROFILES

---

Wenkai Fan\*<sup>1,2</sup>, Shurui Zhang\*<sup>1,2</sup>, Xiaolong Wang<sup>1,2</sup>, Haowei Yang<sup>1,2</sup>, Tsz Wai Chan<sup>1,2</sup>, Xingyan Chen<sup>1,2</sup>, Junquan Bi<sup>1,2</sup>, Zirui Zhou<sup>2</sup>, Jia Liu<sup>2</sup>, and Kani Chen<sup>† 1</sup>

<sup>1</sup>The Hong Kong University of Science and Technology

<sup>2</sup>Bahinia AI

## ABSTRACT

*Aivilization v0* is a publicly deployed large-scale artificial society that couples a resource-constrained sandbox economy with a unified LLM-agent architecture, aiming to sustain long-horizon autonomy while remaining executable under rapidly changing environment. To mitigate the tension between goal stability and reactive correctness, we introduce (i) a hierarchical branch-thinking planner that decomposes life goals into parallel objective branches and uses simulation-guided validation plus tiered re-planning to ensure feasibility; (ii) an adaptive agent profile with dual-process memory that separates short-term execution traces from long-term semantic consolidation, enabling persistent yet evolving identity; and (iii) a human-in-the-loop steering interface that injects long-horizon objectives and short commands at appropriate abstraction levels, with effects propagated through memory rather than brittle prompt overrides. The environment integrates physiological survival costs, non-substitutable multi-tier production, an AMM-based price mechanism, and a gated education-occupation system. Using high-frequency transactions from the platform's mature phase, we find stable markets that reproduce key stylized facts (heavy-tailed returns and volatility clustering) and produce structured wealth stratification driven by education and access constraints. Ablations show simplified planners can match performance on narrow tasks, while the full architecture is more robust under multi-objective, long-horizon settings, supporting delayed investment and sustained exploration.

**Keywords** Large language model agents, Multi-agent systems, Artificial society, Social simulation.

## 1 Introduction

Artificial societies have long served as testbeds for understanding collective behavior, economic dynamics, and institutional effects [8, 24]. Recent progress in large language models (LLMs) has enabled agents that can communicate, reason, and generate diverse behaviors in open-ended environments [4, 41, 44]. However, deploying LLM agents in large-scale, long-horizon, rule-intensive worlds introduces a persistent tension: agents must remain teleologically stable (pursuing coherent multi-day objectives) while being reactively correct in the face of shifting prices, constrained resources, and strategic interdependence with other agents.

*Aivilization v0* is designed around this tension. It is not only a multi-agent simulation but a public-facing societal platform in which users create and steer autonomous agents via long-term goals and occasional direct commands, producing a hybrid ecosystem of human-guided and self-directed autonomy. The technical objective is twofold: (1) provide a unified agent architecture that remains executable under hard constraints and supports identity evolution;

---

\*These authors contributed equally.

<sup>†</sup>Corresponding author: makchen@ust.hk.

<sup>‡</sup>Project website: <https://aivilization.ai>

and (2) build a structurally realistic social-economic environment that generates interpretable, data-rich emergent phenomena.

Hierarchical planning and memory-augmented agent architectures have been explored in prior work [31, 51, 60]. *Avilization v0* integrates planning, simulation, memory, and social cognition into one loop. The BranchThinking Planner decomposes an overarching life objective into parallel branches to reduce error propagation from long sequential plans and to support concurrent multi-domain reasoning. Each cycle selects the most viable sub-task given internal states (eg., energy, satiety, health, inventory) and external context (eg., market prices, eligibility rules). To ensure actions are executable, an Action Simulator runs pre-execution rollouts to detect constraint violations and triggers local repair; execution failures invoke a tiered re-planning pipeline from fast Short-Term Memory based fixes to full top-down replanning when necessary.

Persistent agent identity and memory-driven behavior have been studied in cognitive architectures and social agents [41, 42]. Beyond competence, agents must exhibit persistent yet evolving individuality. *Avilization v0* models identity as an Adaptive Profile co-evolving through (i) human instruction and (ii) social interaction feedback. A dual-process memory separates short-term execution traces (for rapid correction and behavioral adaptation) from long-term semantic integration (values, personality, habits, summarized social records) that biases future goal selection and interaction strategies. This design aims to align real-time control with long-run personalization while allowing social context to reshape agent identity.

To avoid toy dynamics, *Avilization v0* implements a tightly coupled socioeconomic system. Agents must sustain themselves through consumption and recovery; productivity depends on physiological state and education. Goods are produced through multi-tier supply chains with non-substitutable inputs and residential barriers. Prices emerge from an automated market maker mechanism that couples liquidity with output while a macro price index tracks inflation and drives wage adjustments, following approaches in artificial market design [26]. Education and occupation are linked via both fixed floors and distribution-based dynamic thresholds, preserving scarcity and positional competition as the population upskills.

Using high-frequency transactions from a mature phase of the public deployment, we construct 5-minute OHLC series and evaluate whether simulated markets reproduce key statistical regularities. We find stable price paths without collapse or explosion, yet returns display heavy tails and volatility clustering, which are stylized facts extensively documented in empirical finance and econophysics literature [14, 19, 32]. We also observe structured inequality: wealth increases nonlinearly with educational attainment, and occupation tiers display stratified wealth distributions, consistent with findings from labor economics and agent-based inequality models that emphasize institutional gating and intertemporal trade-offs [46].

To further isolate which components are essential for robustness under long horizons and competing constraints, we conduct a controlled ablation study on the planning stack. We compare the full Branch-Thinking Planner with two simplified variants that remove (i) branch decomposition (Without-Branch) or (ii) objective decomposition (Without-OD). Results reveal a consistent pattern: while simplified planners can be competitive on simple, single-objective tasks with direct execution paths, they fail to generalize to complex multi-objective settings that require balancing production, earning, and physiological maintenance. The full architecture better supports coordinated delayed investments such as allocating study actions to increase long-run efficiency, and achieves materially stronger overall performance on high-tech production and exploration tasks when normalized by realistic time costs.

In summary, *Avilization v0* contributes (i) an agent architecture for long-horizon, constraint-aware autonomy with evolving identity, (ii) a coupled economic-education-labor-industry simulation platform designed for emergent macro phenomena, (iii) empirical evidence that the resulting market and stratification dynamics align with canonical stylized facts, supporting the platforms use as a research-grade artificial society, and (iv) an ablation-based validation showing that hierarchical branching and objective decomposition are key to robust multi-objective autonomy, while lighter variants suffice mainly for narrowly defined tasks.

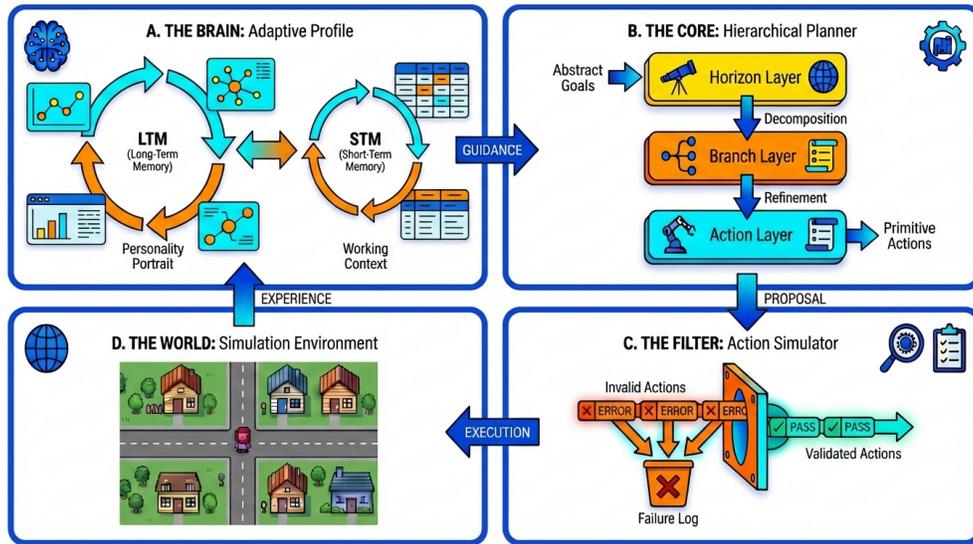
## 2 Multi-agent System

LLMs have rapidly shifted multi-agent systems from rule-based or policy-network agents toward language-driven, tool-using, memory-equipped autonomous agents capable of long-horizon interaction in open worlds. Early work established prompting-based agents as planners and controllers, enabling LLMs to produce executable action sequences and to ground reasoning through external tools and environments [3, 49, 56]. Subsequent research emphasized agentic cognition loops-plan, act, observe, reflect-and showed that reliability improves when agents self-criticize and iteratively revise decisions via reflection, debate, or verbal reinforcement learning [9, 35, 51]. In parallel, the community developed

scaffolding frameworks for tool use and orchestrations such as function calling, modular agent graphs, and multi-agent role decomposition which operationalize LLM agents as systems rather than monolithic models [45, 54, 59].

For social simulation specifically, recent lines of work demonstrate that LLM-driven agents can produce human-interpretable social behaviors when equipped with structured memory, personas, and interaction logs, and that scaling populations to hundreds or thousands of agents yields measurable emergent phenomena [16, 30, 33, 41]. Complementary efforts explore interactive game/sandbox environments and market-like mechanisms where agents must coordinate, compete, and trade under explicit constraints, motivating architectures that couple planning with execution checking and rule compliance [43, 53, 57, 62].

Building on these developments, we present a unified agent architecture for large-scale artificial social simulation, targeting the core tension between high-level strategic teleology and low-level reactive adaptability in environments governed by explicit rules, hard resource constraints, and strategic interdependence among agents. Coherent with our system, designed not only as a multi-agent simulation but also as a public-facing societal platform, users can create and steer autonomous agents via long-horizon objectives and temporary commands, forming a hybrid ecosystem of human-guided and self-directed autonomy. Rather than treating planning, memory, and social behavior as loosely coupled components, our design integrates three mutually reinforcing mechanisms: hierarchical branch-thinking for parallel, temporally abstract goal decomposition and planning; social-cognitive evolution for updating agent identity through interaction-driven feedback; and dual-process memory consolidation for separating fast execution traces from slow, identity-level integration. Together, these mechanisms allow agents to maintain stable strategic direction while continuously adapting to shifting market signals, physiological constraints, and social dynamics.



**Figure 1. Internal organization of the agent cognitive core and its interaction with the environment.** The cognitive system is decomposed into four functional components: (A) a dual-process memory system supporting consolidation, and working context; (B) a hierarchical planner that transforms abstract goals into primitive actions through horizon, branch, and action layers; (C) an action simulator that filters and validates proposed actions; and (D) a simulation environment that executes actions and generates experience. Information flows form a closed feedback loop linking memory, planning, execution, and learning.

## 2.1 Hierarchical Planning Structure

To address the challenge of agent planning, reconciling long-term strategic coherence with short-term adaptive agility, we propose a novel planning framework. Traditional planners often polarize towards either rigid, far-sighted plans vulnerable to environmental shifts, or myopic, reactive behaviors that fail to achieve overarching goals. Our framework is engineered to operate effectively within complex, rule-governed worlds, managing a wide scope of tasks from individual development (e.g., health, education) to multi-agent social dynamics. It achieves this through a synergistic architecture comprising three core components: a Branch-Thinking Planner (BTP) for hierarchical goal management, an Action Simulator (AS) for pre-execution validation, and an Adaptive Re-planning module for robust failure recovery.

### 2.1.1 Branch-Thinking Planner

At the heart of our framework is the Branch-Thinking Planner (BTP), which decomposes the planning horizon into a multi-layered hierarchy. This design directly tackles the trade-off between strategic vision and tactical flexibility.

**Strategic Decomposition:** At the highest level, the agent formulates a long-term objective, which is then strategically decomposed into multiple parallel, semi-independent objective branches, such as personal development, production and resource management, trading and market analysis, and social engagement. By distributing the overarching goal across parallel branches, the planning problem is effectively factorized. Each branch now operates within a focused, constrained subspace, which significantly mitigates error propagation common in long sequential plans and allows for concurrent reasoning about disparate goalsa necessity for managing the multifaceted life of an agent.

Within each individual branch, a second layer of decomposition occurs. The branch’s high-level objective is broken down into a sequence of abstract sub-tasks. This hierarchical structure introduces temporal abstraction, granting tactical flexibility. The agent can adapt its immediate sub-task based on real-time feedback and local state constraints without needing to re-evaluate the entire strategic plan. This highest level goal decomposition is not invoked at every planning cycle but is strategically triggered only when a significant shift in environmental conditions or a strategic re-assessment is warranted, thereby preserving computational resources for real-time adaptation required at lower levels.

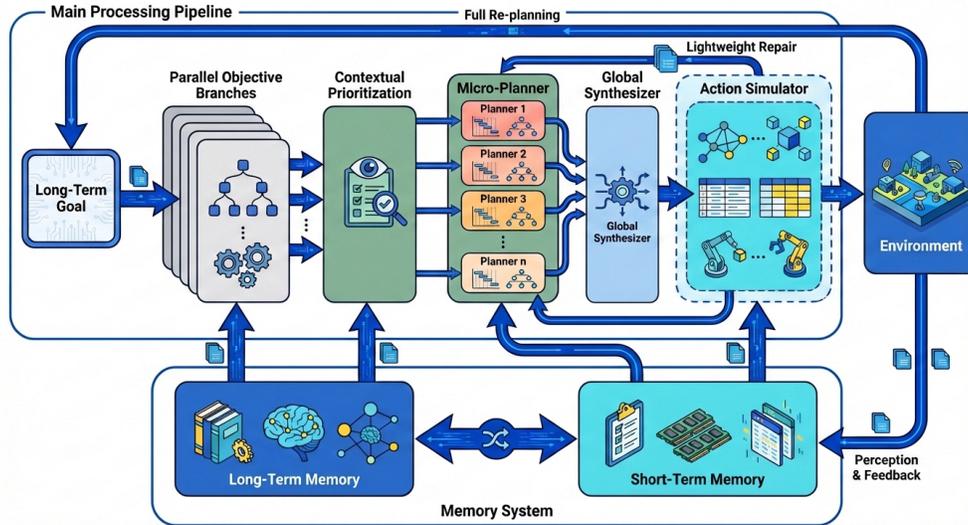
**Contextual Prioritization and Selection:** Each planning cycle commences with Contextual Prioritization and Selection. In this critical step, the agent evaluates all active sub-tasks from its parallel objective branches against its real-time state and the current environmental context. The state includes internal attributes like health, energy, satiety, and inventory, while the context includes external factors such as market prices or world rules (later introduced in Section 3). Given the dynamic nature of the world, this process is not merely a filter but a sophisticated prioritization mechanism. The most viable sub-task, which is ranked based on agent’s highest long-term goal, personality and value (later introduced in Section 2.2.1), is then selected for execution in the current cycle. This ensures that the agent’s immediate focus is always directed towards the most relevant and high-impact objective, effectively grounding the long-term strategy in the present reality and preventing wasted effort on obsolete or infeasible actions.

**Action Sequence Generation:** Following the selection of the prioritized sub-task, the Action Sequence Generation layer translates this abstract objective into a concrete sequence of atomic actions. This translation is performed by a domain-specific micro-planner, which operates within the highly constrained context of the selected branch and the target sub-task. Because the preceding layers have already handled strategic decomposition and contextual selection, this micro-planner’s potential action space is significantly reduced and the resulting sequence is more executable and contextually coherent, aligning the agent’s immediate physical operations with both its current state and its broader strategic intentions. This structured, top-down refinement from abstract goals to concrete actions is fundamental to our framework’s ability to be both strategic and reactive.

**Global Synthesis:** Finally, the action sequences from the active branches are integrated by the Global Synthesis module. This is not a mere concatenation but a sophisticated orchestration process that resolves inter-branch conflicts and ensures global coherence. The module strategically interleaves actions based on their alignment with the agent’s overarching, long-term objective. This is achieved by assigning priorities to actions, where priority is determined by their contribution to the high-level goal, the urgency of the originating branch, and adherence to shared resource constraints like time and energy. This process explicitly addresses the tension between local and global optima. For example, it prevents an over-focus on a single branch’s objective (e.g., maximizing production) from jeopardizing the agent’s overall well-being or progress towards the main goal. By re-aligning the final, integrated action plan with the original strategic intent established at the highest level, the framework forms a crucial closed loop. This ensures that every executed action, while tactically chosen, remains strategically sound, maintaining the integrity and focus of the agent’s plan in a dynamic environment.

### 2.1.2 Pre-execution Simulation

To bridge the gap between abstract planning and reliable execution in a dynamic world, we introduce an Action Simulator (AS). This module serves as a crucial pre-execution validation layer, designed to enhance plan robustness and mitigate failures arising from unforeseen state changes or environmental dynamics. Before committing any action sequence to the environment, the AS performs a counterfactual simulation. It projects the sequence’s outcomes by prospectively evolving the agent’s current state based on a predictive model of the world. This forward simulation identifies potential execution failures, such as resource shortfalls, constraint violations, or logical inconsistencies in the plan.



**Figure 2. Agent cognitive architecture with hierarchical planning, simulation-based filtering, and dual-process memory.** Overview of the agent-level cognitive system governing planning, execution, and adaptation. Long-term memory and short-term memory jointly provide guidance and working context to a hierarchical planner, which decomposes abstract goals into branches and executable actions. Proposed actions are evaluated by an action simulator to detect infeasible or conflicting behaviors, enabling lightweight repair or full re-planning before execution in the environment. Environmental feedback and execution outcomes are fed back to update short-term state and consolidate long-term experience.

Upon detecting a potential failure, a two-stage adaptive repair process is initiated. First, a local repair mechanism attempts to resolve the issue using a set of pre-defined, computationally inexpensive heuristics. If these simple fixes are insufficient, the system escalates the issue to a reactive correction module. By leveraging knowledge cached from past experiences and performing rapid, pattern-based reasoning, this module generates a viable alternative action or a minor plan modification.

The entire simulation and repair loop operates efficiently, bypassing the need to re-invoke the full, multi-layered planning stack. This design ensures high plan executability and adaptability while preserving critical computational resources, enabling the agent to be both deliberate in its planning and agile in its execution.

### 2.1.3 Adaptive Re-planning and Feedback Loops

Recognizing that no plan survives contact with reality, the framework is equipped with a tiered Adaptive Re-planning module to handle execution failures and unexpected opportunities. This mechanism ensures resilience while optimizing computational resources.

Upon a minor action failure, the agent first attempts a rapid fix. It consults its Short-Term Memory, a buffer of recent succeeded and failed actions, outcomes, and reflections, to re-evaluate and adjust the immediate next steps. This approach, akin to muscle memory, provides swift, low-cost adaptation for common setbacks. Only when lightweight repairs fail consecutively or a major environmental shift occurs does the agent escalate to a full re-plan. This involves invoking the entire BTP pipeline to reassess objectives and generate a new strategy from the top down. While costly, this serves as a guaranteed recovery mechanism.

This tiered design embodies our framework’s philosophy of balancing responsiveness and strategic deliberation. By defaulting to fast fixes and reserving costly re-planning for significant disruptions, the agent remains both agile and strategically focused, capable of navigating the uncertainties of its complex world.

## 2.2 Adaptive Agent Profile

Our framework posits that an agent’s identity is not a static set of parameters but an adaptive profile that evolves through lived experience, particularly social interaction. This profile serves as the central nexus, integrating the agent’s physical state, cognitive processes, and social identity. Its evolution is driven by a continuous cycle of action, social engagement, and memory consolidation. In this section, we detail the components of this architecture, starting with

the primary engine of change, social interaction, followed by the cognitive machinery that processes it, culminating in the unified profile structure.

### 2.2.1 Social Interaction as the Engine of Evolution

Agents in our simulation are social beings, not just task-executors. Their growth is fundamentally driven by a structured cycle of social engagement. This process begins with proactive planning, where an agent selects interaction targets and conversation topics based on its current goals, state (e.g., economic needs), and existing personality. The interaction itself unfolds through natural language dialogue, facilitating the exchange of knowledge, opinions, and social cues.

Each social event triggers a post-interaction reflection phase. During this phase, the agent updates its internal models of other agents (e.g., relation, attitude) and, most importantly, feeds the experiential data into its memory systems. This act of consolidation is what translates transient social encounters into lasting changes in the agent’s core identity, such as its values and behavioral tendencies. This feedback loop, where the current profile guides social action, and social outcomes reshape the profile, is the core mechanism for agent evolution.

### 2.2.2 Dual-Process Memory Architecture

To manage the distinct demands of immediate action and long-term growth, we implement a dual-process memory architecture. This design functionally separates transient, action-oriented experience from enduring, identity-shaping knowledge, mirroring the cognitive distinction between fast, intuitive reasoning and slow, deliberative integration.

**Short-Term Memory:** The Short-Term Memory (STM) operates as a high-frequency buffer focused on immediate task execution and outcome evaluation. It continuously records execution traces from the planner, cataloging them as Successful Actions or Failed Actions. When a planned action sequence concludes, whether through full success or by triggering a re-planning cycle, the execution traces are summarized into concrete experiences, which serve as a direct, pragmatic function: successful patterns refine the mapping from sub-goals to actions, enhancing the efficiency of future planning cycles; failure patterns prime the re-planning module for quicker recovery from similar setbacks. Thus, the STM embodies a form of fast thinking [17], allowing the agent to dynamically adapt its behavior based on recent outcomes without altering its core identity.

**Long-Term Memory:** In contrast, the Long-Term Memory (LTM) functions as the agents stable yet evolving semantic core. It integrates significant experiences, particularly from social interactions, to form a coherent Agent Profile. This profile encompasses persistent attributes such as values, personality traits, and a summary of social interaction records. Unlike the STMs event-specific logs, the LTM performs a slow, integrative synthesis. The resulting profile acts as a high-level control signal, biasing strategic decisions, such as goal selection in the Contextual Prioritization layer. In this way, the LTM governs the agent’s character, ensuring that behavior aligns with a continuously refined sense of self.

These two memory streams are not isolated but engage in a continuous, bidirectional dialogue. Tactical successes logged repeatedly in the STM can gradually consolidate into the LTM as new knowledge or stable preferences. The high-level patterns, such as values and personality, stored in the LTM provide a top-down cognitive context, shaping how experiences in the STM are interpreted and prioritized. This synergy creates a powerful feedback loop where immediate actions inform long-term identity, and identity, in turn, guides future actions.

### 2.2.3 The Unified Agent Profile Structure

The interplay between social interaction and the dual-memory system culminates in the Unified Agent Profile. This data structure holistically represents the agent at any given moment, comprising three key components:

- **Dynamic State:** Transient attributes for immediate decision-making, including energy, satiety, health, education score, current balance, residential tier, job, inventory.
- **Short-Term Memory:** A tactical ledger of recent actions and their outcomes, both successful and failed, driving rapid behavioral adjustments.
- **Long-Term Memory:** The stable yet evolving core of the agent’s identity, encompassing its values, personality, habits, a curated social interaction record, and other defining traits.

An example of the agent profile is demonstrated in Appendix A Table 6. Ultimately, this architecture forges agents capable of both long-term, identity-driven strategizing and short-term, real-time responsiveness. The Adaptive Profile

ensures that every action and interaction has the potential to contribute to a rich, emergent narrative of individual growth within a complex social world.

### 2.3 Human-in-the-Loop Steering

A distinguishing feature of *Aivilization v0* is that agents are not deployed as fully sealed autonomous entities. Instead, the platform is explicitly designed for hybrid autonomy, where humans can steer agents in real time while preserving agent-level coherence and executability. Concretely, we support two complementary intervention channels, long-horizon objective and temporary commands, which are deliberately mapped to different levels of the agent architecture and propagated through different pathways.

**Strategic Steering:** Users may assign or revise an agents long-horizon objective. In our architecture, this type of input is treated as a top-level long-term goal as well as integrated into the long-term memory. It reshapes the agents global planning landscape by affecting how the Branch-Thinking Planner constructs and weights parallel branches, how the Contextual Prioritization layer ranks candidate sub-tasks under resource constraints, and how Global Synthesis resolves trade-offs across domains. In other words, long-horizon objective are not implemented as isolated prompts; they are integrated into the full hierarchical planning stack and thus influence both multi-turn consistency and per-cycle task selection.

This design provides two main advantages. First, it enables stable, interpretable governance: human intent is expressed at the same abstraction level as the agents strategic decomposition, reducing brittleness that arises when users attempt to steer behavior via repeated low-level instructions. Second, it yields compute-efficient persistence: because the highest-level decomposition is only re-invoked under major context shifts, long-term steering can remain effective across many cycles without requiring continuous user intervention.

**Reactive Steering:** Users can also issue short, situational commands (e.g., "buy 10 fish now", "sleep for 8 hours", "read books for 2 hours"). These commands are operationally different from long-term goals: they are typically narrow in scope, have clear success criteria, and are intended to take effect immediately rather than restructure the agents life plan. Accordingly, upon receiving a temporary command, the agent triggers a lightweight planning route: instead of running the full hierarchical BTP pipeline, it invokes only (i) a localized planner to translate the command into an executable action sequence and (ii) the Action Simulator to validate feasibility and perform fast repair if needed. This choice reflects a systems principle: short-horizon clarity should not pay the cost of full-horizon deliberation.

Despite bypassing full top-down replanning, temporary commands are not one-off overrides that disappear from cognition. Their effects are propagated through the memory system, allowing local interventions to influence future behavior in a principled way. Specifically, the commands plan, execution result, and outcome are written into Short-Term Memory, immediately shaping subsequent planning cycles. Over time, relevant command-associated experiences can be consolidated into Long-Term Memory, where they may alter enduring profile items such as habits, preferences, and even higher-level value tendencies, thereby eventually biasing the highest layers of hierarchical planning.

This memory-mediated propagation provides two benefits. First, it ensures behavioral continuity: even when users intervene reactively, the agent incorporates the consequences into its ongoing decision process rather than oscillating between unrelated directives. Second, it supports personalization through interaction: repeated user steering becomes part of the agents identity evolution, enabling agents to gradually internalize a users management style without requiring constant micromanagement.

Together, these two intervention modes operationalize the human-guided and self-directed autonomy: long-term goals steer what the agent is becoming, while temporary commands steer what the agent should do right now, with both channels integrated into a unified loop via hierarchical planning and dual-process memory.

## 3 Social Simulation Platform Design

To support large-scale, long-horizon social simulation, we design a tightly integrated platform that couples economic production, labor allocation, education, and market exchange into a unified operational framework. Unlike conventional agent-based simulations that often treat economic behaviors as simplified rule modules or isolated subsystems, our platform emphasizes structural realism and system-level interdependence. At the core of this design is a multi-layered architecture where automated market mechanisms and hierarchical industrial supply chains interact with a dynamic labor market gated by human capital requirements. This closed-loop environment ensures that macro-level patterns, such as industrial specialization, wealth stratification, and resource bottlenecks, emerge naturally from micro-level interactions governed by endogenous constraints rather than predefined equilibrium assumptions. In the follow-

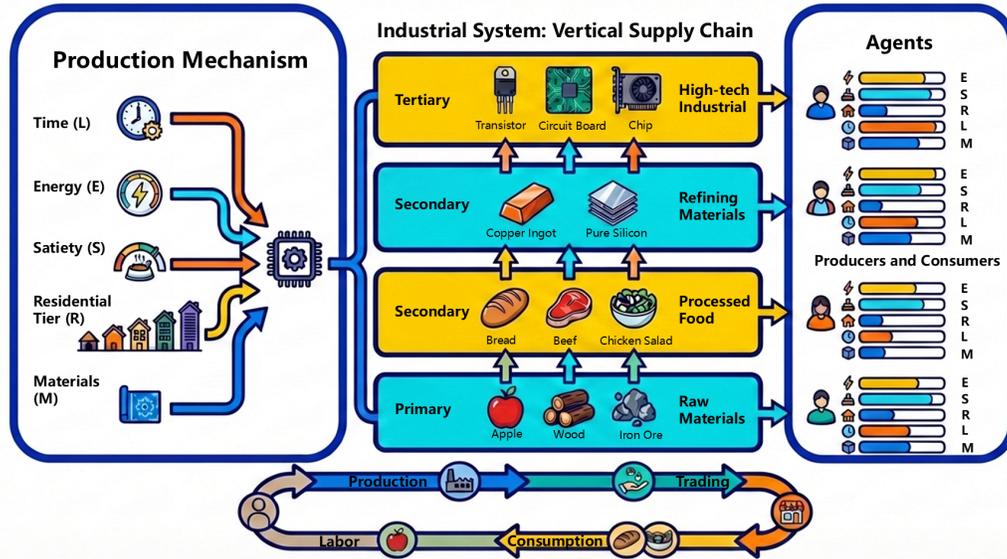
ing sections, we detail the mechanisms of the industrial economy and the education-occupation system. To facilitate this description and maintain consistent terminology, we first introduce the platform-level definitions regarding agent assets and status.

- **Currency balance**  $B_a(t)$ : the liquid medium of exchange held by agent  $a$  at time  $t$ , used to purchase goods and services.
- **Inventory**  $Q_a(t) = (q_{a,1}(t), \dots, q_{a,n}(t))$ : quantities of all tradable commodities held by agent  $a$ .
- **Net worth**  $W_a(t)$ : the accounting value of an agents total economic position. In the current platform version (no explicit debt instrument), we define

$$W_a(t) = B_a(t) + \sum_{i=1}^n p_i(t) q_{a,i}(t), \quad (1)$$

where  $p_i(t)$  is the price of commodity  $i$  at time  $t$ .

- **Residential tier**  $R_a(t)$ : a discrete status that gates access to certain production technologies and occupations.



**Figure 3. Economic environment and production-consumption structure of the simulated world.** Schematic of the closed economic system in which agents operate. Production is governed by a resource-constrained production function combining time, energy, satiety, residential level, and materials. Goods flow through a vertically structured supply chain spanning primary resources, consumer goods, intermediate inputs, and high-technology sectors. Agents simultaneously act as producers and consumers, participating in labor, production, trading, and consumption cycles that jointly determine market dynamics and resource allocation.

### 3.1 Industry and Economic System

Our platform introduces a closed-loop industrial and economic system, designed to emulate real-world labor-production-trade-consumption cycles (Figure 3). Key components include: a commodity price mechanism governing market exchanges; the Survival and Physiological Constraints linking agent vitality (energy, satiety, health) to consumption and efficiency with safety-net provisions; a multi-tiered industrial taxonomy enforcing residential barriers and upstream-downstream dependencies; production processes with hard, non-substitutable resource constraints to propagate supply-chain dynamics; and a stochastic reward mechanism introducing behavioral uncertainty. By integrating structural constraints (tiered eligibility, fixed recipes) with dynamic market interactions, the architecture sustains stable long-horizon simulations while preventing degenerate loops. This framework establishes a high-fidelity foundation for investigating socio-economic phenomena where technological progress emerges organically from cumulative, resource-bound production processes.

### 3.1.1 Survival and Physiological Constraints

The Survival and Physiological Constraints governs three core physiological states for each agent: satiety  $S_a(t)$ , energy  $E_a(t)$ , and health value  $J_a(t)$ , which have different upper bounds subject to residential tier  $R_a$ . All work and production activities cost satiety and energy. Satiety can be replenished through food consumption, with higher level commodities yielding faster restoration. Energy recovers primarily via sleep. If energy or health falls below critical thresholds, the agent becomes temporarily incapacitated and cannot engage in labor or production. Health deteriorates due to stochastic illness events and cumulative sleep deprivation. Agents can see doctor in hospital and recover health value. Detail recovery action types can be found in Appendix B Table 8.

To further internalize sustainability constraints, productive efficiency is modulated by

$$\text{Eff}_a(t) = G(S_a(t), E_a(t), J_a(t), R_a(t), H_a(t)), \quad (2)$$

where  $G \in (0, 1]$  is a non-decreasing function in each argument,  $H_a(t)$  is the education score (detailed in Section 3.2.1).

Additionally, a safety-net mechanism provides essential subsidies (e.g., basic food items) to agents persistently operating below critical physiological thresholds, thereby preventing irreversible attrition from the economic system and preserving both simulation continuity and minimal welfare standards.

These designs embeds biologically grounded constraints directly into economic rationality. By making productivity dependent on physiological, the system penalizes overwork and sleep deprivation through immediate efficiency losses, creates non-trivial trade-offs between short-term output and long-term capacity, generates organic demand for food, rest, and healthcare within the market, and ensures that rational agents must strategically balance labor, consumption, and recovery.

### 3.1.2 Commodity Price Mechanism and Macroeconomic Index

The in-game commodity system employs a hybrid architecture that functions as both a market exchange and a macroeconomic monitoring tool. The core of the exchange is a liquidity-based Automated Market Maker (AMM) model, which also serves as an algorithmic central bank to create an elastic money supply [2, 47]. Unlike traditional closed-economy simulations with a fixed monetary base, our platform's AMM acts as the ultimate liquidity provider. New currency is explicitly minted and injected into circulation when agents sell commodities to a liquidity pool, and it is conversely removed upon purchase. This design structurally couples the aggregate money supply with real economic output. As agents improve productivity and extract more resources, the continuous selling pressure into the AMM leads to a secular expansion of liquidity, reflecting genuine economic growth.

At the micro-level, the instantaneous price of each commodity is governed exclusively by an automated dynamic pricing system. Each commodity is paired with the base game currency in a dedicated liquidity pool. The valuation within each pool adheres to the Constant Product Formula [2, 5]:

$$\text{IS}_i(t) \cdot \text{CR}_i(t) = k \quad (3)$$

where:

- $\text{IS}_i(t)$  represents the inventory supply of commodity  $i$  in the pool at time  $t$ .
- $\text{CR}_i(t)$  represents the reserve of the base currency in the pool at time  $t$ .
- $k$  is the constant product invariant, which remains fixed for the duration of an experiment.

Under this mechanism, the price of a commodity is determined by the ratio of the reserves in its pool. This quoted price,  $p_i(t)$ , represents the instantaneous marginal price for an infinitesimally small trade, calculated as  $p_i(t) = \text{CR}_i(t)/\text{IS}_i(t)$ . For any transaction of a finite size, however, agents experience price slippage [6]. The effective price,  $p_i^{\text{eff}} = \Delta\text{CR}/\Delta\text{IS}$ , reflects the total currency exchanged for a given amount of the commodity and will differ from the initial marginal price. Despite this, the instantaneous price  $p_i(t)$  serves as the publicly quoted market price at any given moment. As a commodity is purchased from the pool (decreasing  $\text{IS}_i$  and increasing  $\text{CR}_i$ ), the instantaneous price  $p_i(t)$  rises continuously, signaling increasing scarcity. This ensures that all price signals are direct, real-time reflections of agent-driven supply and demand dynamics.

While the AMM governs real-time prices, the system simultaneously computes a macro-level price trend index [21] to monitor economy-wide inflation. This index does not alter transaction prices but serves as a crucial input for other economic modules, such as dynamic wage adjustments (see Section 3.2.4). The index tracks price fluctuations across two core categories: Food and Non-food commodities.

For any commodity  $i$ , let  $p_i(0)$  be its initial baseline price and  $p_i(t)$  be its current market price as determined by the AMM. The Price Change Ratio ( $\text{PCR}_i$ ) is defined as:

$$\text{PCR}_i(t) = \frac{p_i(t)}{p_i(0)} \quad (4)$$

To prevent extreme price spikes of a single item from distorting the aggregate measure, the system employs the geometric mean of individual price ratios to calculate a robust categorical inflation rate.

- **Food Inflation Rate ( $\bar{\text{PCR}}_{\text{food}}$ ):** For  $n_{\text{food}}$  types of food items:

$$\bar{\text{PCR}}_{\text{food}} = \left( \prod_{i=1}^{n_{\text{food}}} \text{PCR}_i \right)^{1/n_{\text{food}}} \quad (5)$$

- **Non-food Inflation Rate ( $\bar{\text{PCR}}_{\text{non-food}}$ ):** For  $n_{\text{non-food}}$  types of non-food items:

$$\bar{\text{PCR}}_{\text{non-food}} = \left( \prod_{j=1}^{n_{\text{non-food}}} \text{PCR}_j \right)^{1/n_{\text{non-food}}} \quad (6)$$

Finally, a global index,  $\bar{\text{PCR}}_{\text{overall}}$ , is calculated as a weighted arithmetic mean of the categorical rates to reflect the structural composition of the economy. Let  $N$  be the total number of commodity types ( $N = n_{\text{food}} + n_{\text{non-food}}$ ):

$$\bar{\text{PCR}}_{\text{overall}} = \frac{n_{\text{food}}}{N} \bar{\text{PCR}}_{\text{food}} + \frac{n_{\text{non-food}}}{N} \bar{\text{PCR}}_{\text{non-food}} \quad (7)$$

This overall index provides a comprehensive, high-level signal of the economy's inflationary or deflationary state, enabling systemic responses to macroeconomic trends without interfering with the micro-foundations of price discovery.

This dual-component design strengthens the socio-economic coherence of the simulated society. The AMM layer ensures that every transaction has an immediate and localized market impact, creating an authentic feedback loop between individual actions and resource availability. By anchoring valuation entirely to agent-generated activity, the mechanism guarantees that economic phenomena such as resource scarcity, market volatility, and wealth disparities emerge organically from internal interactions.

Simultaneously, the macroeconomic index layer provides the system with a means to observe and react to emergent societal-scale trends. For example, by tracking food price inflation, the system can dynamically adjust socio-economic parameters like baseline wages, simulating policy responses to changes in the cost of living. This allows the simulation to couple micro-level behaviors with macro-level systemic stability, fostering a self-consistent ecosystem where price dynamics not only reflect but also inform the evolution of social trust, inequality, and collective adaptation strategies.

### 3.1.3 Industrial Organization and Commodity Taxonomy

To emulate the depth of real industrial systems, commodities are organized into a vertical supply-chain taxonomy to induce endogenous division of labor and upstream–downstream dependency. All commodities are organized into three hierarchical tiers, with production eligibility constrained by residential barriers such that agents with residential tiers below the threshold  $R_{\text{min}}$  cannot access corresponding higher-level commodity production [27, 61].

- **Primary Sector:** Raw Materials. Production requires low house levels, analogous to agriculture and extractive industries. These goods serve both as essential energy sources for survival and as upstream inputs for industrial production.
- **Secondary Sector:** Processed food and refining materials. These industries convert agricultural and mineral resources into high-energy foods and industrial intermediates.
- **Tertiary Sector:** High-Tech Manufacturing. These industries require the highest capital level and long production cycles, heavily depending on midstream products.

All commodities, their supply-chain tiers, minimal residential tier requirements ( $R_{\text{min}}$ ), and typical role in the platform are listed in Appendix B Table 7. This multi-tiered industry design ensures that resource scarcities and technological prerequisites propagate coherently across sectors, inherently preventing unrealistic production shortcuts and degenerate economic loops. Advancement through industrial tiers demands cumulative investment in infrastructure, materials, and time, transforming technological upgrading into a resource-bound, path-dependent process rather than an instantaneous transition [7, 20]. These features collectively anchor long-horizon simulations in economic realism: they enforce credible capacity limits, sustain meaningful inter-sectoral feedback, and ensure industrial progress emerges organically from sustained production achievements.

### 3.1.4 Production with Hard Resource Dependencies

Production in the platform is modeled as a conditional synthesis process with (i) hard technological access constraints and (ii) non-substitutable input requirements. This is formally represented by a Leontief-style production function, which remains a foundational tool in modern input-output economics and the analysis of production networks [38].

For an agent producing commodity  $i$  at time  $t$ , output is

$$Y_{a,i}(t) = \mathbb{I}(R_a \geq R_{\min}^i) \cdot \min \left( \min_{m \in \mathcal{M}_i} \frac{M_{a,m}(t)}{\alpha_{i,m}}, \frac{E_a(t)}{\epsilon_i}, \frac{S_a(t)}{\sigma_i}, \frac{L_a(t)}{\tau_i} \right), \quad (8)$$

where

- $\mathbb{I}(R_a \geq R_{\min}^i)$  is an indicator function, implementing the residential barrier for commodity  $i$ .
- $\mathcal{M}_i$  is the set of material inputs required by the recipe of commodity  $i$ .
- $M_{a,m}(t)$  is the available material  $m$  in the inventory of agent  $a$ .
- $E_a(t)$  and  $S_a(t)$  denote the energy and satiety of agent  $a$  at time  $t$  and  $L_a(t)$  denotes available labor time.
- $\alpha_{i,m}$  represents the amount of material  $m$  required to produce one unit of product  $i$ .
- $\epsilon_i$ ,  $\sigma_i$ , and  $\tau_i$  are per-unit requirements of energy, satiety, and labor time for producing  $i$ .

All technical coefficients ( $\mathcal{M}_i$ ,  $\alpha_{i,m}$ ,  $\epsilon_i$ ,  $\sigma_i$ , and  $\tau_i$ ) in Equation (8) are exogenously specified recipe parameters (i.e., institutional rules). They are held fixed within an experiment run unless explicitly modified by an intervention. The complete set of production recipes, detailing the material inputs ( $\mathcal{M}_i$ ,  $\alpha_{i,m}$ ) and the per-unit energy ( $\epsilon_i$ ), satiety ( $\sigma_i$ ) and labor time ( $\tau_i$ ) costs for each commodity, is provided in Appendix B Table 9.

The design of these recipes embodies several key principles. Firstly, they establish a hierarchical production chain, from primary resource extraction to complex, high-tier products. This tiered structure ensures that disruptions in basic commodity supply can propagate throughout the economic system, mirroring real-world supply chain dynamics. Secondly, resource costs (energy and satiety) generally increase with product complexity, reflecting the higher value-addition at advanced stages and creating a natural economic gradient [1, 13]. Finally, certain high-value manufacturing processes include a special reward, the rare chance of producing a valuable byproduct, to introduce risk and reward dynamics.

This production function enforces strict non-substitutability among inputs, where a shortage in any single factormaterials, energy, satiety, or labor directly constrains output via the minimal operator. This structure, combined with the residential barrier that reserves advanced production for agents meeting infrastructure thresholds, creates a balanced system. The interplay between these dynamic (resource-driven) and structural (eligibility-driven) constraints prevents unrealistic production surges while ensuring that economic shocks, such as supply bottlenecks, propagate coherently across the supply chain.

### 3.1.5 Stochastic Reward Mechanism

In the production recipes, some product may include a low-probability special reward in addition to deterministic production outputs (see Appendix B Table 9). This mechanism introduces a lottery-like component into agents expected payoffs, which can introduce behavioral uncertainty into occupational selection and production strategies.

## 3.2 Education and Occupation System

The simulation platform structures occupational mobility through explicit, rule-based entry barriers that reflect two foundational dimensions of real-world inequality: (i) human-capital investment, operationalized via education score, and (ii) asset-based access constraints, proxied by an agents residential tier [10, 50]. Only after satisfying both the educational requirement and the residential tier threshold can an agent qualify for higher-tier occupations. This dual-gating mechanism ensures that upward mobility remains contingent on sustained investment in both material and human capital, realistically capturing the interdependence of wealth and education in shaping career trajectories.

### 3.2.1 Education as Human-Capital Investment

In our platform, education is not freely available. Agents must expend resources to enroll in training programs [10]. Each agent can accumulate an education score  $H_a(t)$  through study actions (see Appendix B Table 8). The obtained education score follows:

$$H_a(t + \Delta t) = H_a(t) + \eta \cdot \Delta t, \quad (9)$$

where  $\eta$  is the fixed education accumulation rate,  $H_a(t)$  is the education score at time  $t$ , and  $\Delta t$  is the study duration.

Education score not only serves as a credential for accessing high-tier occupations but also enhance efficiency in all production activities as demonstrated in Equation 2. This process represents a deliberate human-capital investment, where agents must forego immediate labor income and consumption to accrue long-term competitive advantages. Agents face a strategic trade-off: allocating limited money and resources toward education reduces short-term survival capacity but is essential for sustained economic performance. Rational agents must carefully balance educational investment against immediate needs to optimize their long-run position in the economy.

### 3.2.2 Job Tiers, Occupation Catalog, and Entry Barriers

Occupations are organized into discrete job tiers  $\ell \in \{1, 2, \dots, 6\}$ . At job level, each tier imposes (i) a minimum residential tier requirement and (ii) a minimum knowledge-score requirement. This structure is designed to mirror real-world socioeconomic stratification, where access to elite professions is often gated not just by education, but by accumulated capital and social status [52, 58]. In addition, access to higher-tier occupations requires the consumption of specific prerequisite commodities, directly linking occupational advancement to participation in the commodity economy.

At the occupation level, each occupation  $j$  is characterized by: a minimum residential tier requirement  $R_{\min}^{(j)}$  corresponding with the job tier; a knowledge floor  $H_{\text{floor}}^{(j)}$  specified to occupation  $j$ ; and an eligibility share parameter  $\pi_j$ , used by the platform to compute a dynamic knowledge threshold. The eligibility share parameter  $\pi_j \in (0, 1]$  encodes the platforms intended eligible-population share for job  $j$ . For example,  $\pi_j = 0.28$  means the platform sets the knowledge threshold so that approximately the top 28% of agents sorted by education score can meet the requirement. This parameter will be further used to calculate the dynamic knowledge threshold.

Detail configurations can be found in Appendix B, where Table 10 summarizes the tier-level rule configuration under default settings and Table 11 lists the platforms occupation catalog and associated entry parameters.

### 3.2.3 Eligibility Rule and Recruitment Cycle

An agent  $a$  is eligible to apply for occupation  $j$  at time  $t$  if

$$\mathbb{I}(H_a(t) \geq \widehat{H}_{\min}^{(j)}(t)) \cdot \mathbb{I}(R_a \geq R_{\min}^{(j)}) = 1. \quad (10)$$

where  $R_{\min}^{(j)}$  is the fixed minimal residential tier requirement and  $\widehat{H}_{\min}^{(j)}(t)$  is the dynamic knowledge threshold at time  $t$ , which is determined based on the platform-wide distribution of education scores at time  $t$ . Let  $F_t(h)$  denote the empirical cumulative distribution function of  $\mathbb{H}(t) = \{H_a(t)\}_{a \in A(t)}$ , where  $A(t)$  denotes active agents at time  $t$ . Define the  $(1 - \pi_j)$ -quantile of the education distribution as

$$q_{1-\pi_j}(t) = \inf\{h \in \mathbb{R} : F_t(h) \geq 1 - \pi_j\}. \quad (11)$$

The effective knowledge requirement for occupation  $j$  is then:

$$\widehat{H}_{\min}^{(j)}(t) = \max\left(H_{\text{floor}}^{(j)}, q_{1-\pi_j}(t)\right). \quad (12)$$

This dynamic threshold design ensures that occupational entry remains competitively scarce and socially responsive. As the populations average education rises, so do the effective skill requirements for high-tier jobs, preventing credential devaluation and preserving incentives for continuous human-capital investment [11, 18]. By anchoring eligibility to the  $(1 - \pi_j)$ -quantile of the current education distribution while respecting a hard floor, the mechanism captures real-world phenomena like relative positional competition [28], while maintaining stable stratification and meaningful barriers to upward mobility over long simulation horizons.

Agents can submit multiple job applications per recruitment cycle, but the platform imposes an application quota that increases with residential tier:

$$N_a^{\max} = N^{\max}(R_a), \quad (13)$$

where  $N^{\max}(\cdot)$  is a non-negative, bounded, non-decreasing function. This rule makes higher residential tiers associated with greater labor-market flexibility, rewarding long-term investment in residential upgrading [40]. It also reduces search friction for established agents, promoting efficient matching in competitive job markets.

### 3.2.4 Wage Regime

The wage structure within our model is bifurcated into two distinct regimes: a static regime and a dynamic regime. Both regimes are modulated by the overall price trend index,  $\overline{\text{PCR}}_{\text{overall}}$ , as defined in Equation (7), which serves as a proxy for macroeconomic inflation or deflation, thereby adjusting nominal wages to reflect changes in the aggregate price level.

The static wage regime applies to lower-tier occupations (Tiers 1-3). For an occupation  $j$  in this category, the wage is determined by a predetermined base value,  $w_0^{(j)}$ , adjusted only by the price trend index. The wage at time  $t$  is thus calculated as:

$$w^{(j)}(t) = w_0^{(j)} \cdot \overline{\text{PCR}}_{\text{overall}}. \quad (14)$$

This formulation ensures a stable and predictable income floor for foundational roles, insulated from the competitive pressures that characterize higher-tier labor markets.

In contrast, the dynamic wage regime governs higher-tier occupations (Tiers 4-6). Wages in this regime are endogenous, responding not only to macroeconomic conditions but also to the competitive landscape of the labor market. This is captured by incorporating the effective knowledge threshold,  $\widehat{H}_{\text{min}}^{(j)}(t)$ , which reflects the level of skill and expertise required to remain competitive in occupation  $j$ .

The dynamic wage for occupation  $j$  at time  $t$  is specified by the following rule:

$$w^{(j)}(t) = w_0^{(j)} \cdot \Phi(\widehat{H}_{\text{min}}^{(j)}(t)) \cdot \overline{\text{PCR}}_{\text{overall}} \cdot (1 + \delta_t). \quad (15)$$

Here,  $w_0^{(j)}$  is the base wage for the occupation. The function  $\Phi(\cdot)$  is a non-negative and non-decreasing function that translates the rising knowledge threshold into a wage premium, signifying that higher entry barriers command greater compensation [10, 39].  $\delta_t \in [-\bar{\delta}, \bar{\delta}]$  represents a bounded, short-term adjustment factor (e.g., reflecting temporary labor market shocks or policy interventions).

This dual-regime structure, which combines a stable, static floor with a responsive, dynamic component, is designed to mirror real-world labor market segmentation [22]. As articulated in Equation (15), compensation for competitive occupations rises in tandem with aggregate skill levels. This mechanism ensures that wages remain aligned with evolving societal capabilities, thereby sustaining incentives for human capital investment. Simultaneously, the static wage component provides essential income stability for lower-tier workers, anchoring expectations and guaranteeing a minimum livelihood. This hybrid design balances structural predictability with endogenous responsiveness, allowing for a more robust and credible long-horizon analysis of policy impacts, inequality dynamics, and labor market evolution.

### 3.2.5 Labor-Consumption Feedback Loop

A critical feature of our model is the endogenous feedback loop linking labor supply to consumption, which is grounded in the physiological costs of work. In our framework, participation in any occupation incurs per-hour physiological costs, specifically depleting agents' energy and satiety levels. This mechanism creates a tight coupling between the labor market and the goods market. To work, agents must earn a wage sufficient to afford the necessary consumables that restore their capacity to work. This linkage establishes a subsistence constraint on labor supply, ensuring it is endogenously determined by the interplay of three key factors: (1) nominal wages, which dictate purchasing power; (2) commodity prices, which determine the affordability of subsistence; and (3) the physiological demands of different occupations.

## 4 Empirical Analysis and Validation of the Social Simulation Platform

This section presents the empirical outcomes derived from the *Aivilization v0* platform. We detail the experimental environment, the methodology of data acquisition, and the statistical validation of the generated economic behaviors against stylized facts observed in real-world financial markets [19]. The primary objective of this empirical analysis is to validate the fidelity of the *Aivilization v0* economic engine. By juxtaposing the statistical properties of our simulated financial data, such as volatility clustering [23] and fat-tail distributions [25], against established empirical regularities found in real-world markets, we aim to demonstrate a significant degree of statistical correspondence and behavioral similarity between our digital economy and actual human economic systems.

### 4.1 Experimental Setup and Data Acquisition

The *Aivilization v0* platform was officially launched to the public in late August 2025, inviting a global cohort of human participants to interact within a persistent digital society. Unlike strictly controlled laboratory simulations

populated by algorithmic agents, our dataset captures high-frequency economic activities driven by a hybrid population of autonomous AI agents created by human participants. The participants design and manage autonomous AI agents by setting agents’ long-term goal or sending instant commands, creating a hybrid ecosystem of human-directed and AI-driven behavior. As of the cutoff date for this analysis, the ecosystem hosts tens of thousands of agents. The interaction history has generated a cumulative log exceeding 600,000 high-frequency transaction records, with the volume expanding continuously. For the specific purpose of this study, we focus on the economy’s mature phase, by which point key market metrics such as daily trading volume and participant count had stabilized. We extracted a subsequent continuous block of 400,000 transaction records from this phase to minimize transient artifacts associated with the system’s ‘cold start’ period.

To facilitate a comparative analysis with real-world financial time series, we construct high-fidelity price series from the raw transaction logs. The data processing pipeline adheres to standard econometric practices. We utilize a time-binning approach to construct the price series [55]. Specifically, we partition the transaction data into discrete 5-minute intervals. For each interval, we record the first-traded price (Open), the highest traded price (High), the lowest traded price (Low), and the last-traded price (Close), thereby constructing a standard OHLC dataset. We select the last-traded price as representative for each interval.

The platform operates on an accelerated clock. The flow of time in the simulated world is 7 times faster than in the real world. This ratio was chosen during the platform’s design to balance gameplay dynamics and data generation speed.

## 4.2 Macroeconomic Stability Analysis

We examine whether the simulated market economy remains economically meaningful over time, in the sense that (i) prices do not exhibit explosive divergence (i.e., unbounded growth) or collapse to zero, and (ii) price movements exhibit sustained fluctuations rather than degenerating into near-constant series. In detail, we evaluate stability using both a representative visualization and a small set of summary indicators computed from the same price series. For notational simplicity, when the commodity index is unambiguous, we denote the last-traded price at time  $t$  as  $p_t$ , omitting the explicit commodity subscript  $i$ . The corresponding log-price is defined as  $\ell_t = \log p_t$ . The indicators we use are as follows.

- Range of log-price: A bounded log-price path over a long simulation horizon provides evidence against explosive dynamics.

$$\text{Range} = \max_{1 \leq t \leq T} \ell_t - \min_{1 \leq t \leq T} \ell_t. \quad (16)$$

- Maximum drawdown: This statistic summarizes the most severe peak-to-trough decline and serves as a heuristic indicator for collapse-like episodes [15].

$$\text{MDD} = \max_{1 \leq t \leq T} \left( 1 - \frac{p_t}{\max_{1 \leq s \leq t} p_s} \right). \quad (17)$$

These indicators are descriptive and do not assume a specific equilibrium model; they are intended as sanity checks that the simulated economy does not trivially diverge or freeze. We select the fish market for our initial case study, as it is one of the most liquid and actively traded commodity markets within the simulation, making it a representative example for stability analysis.

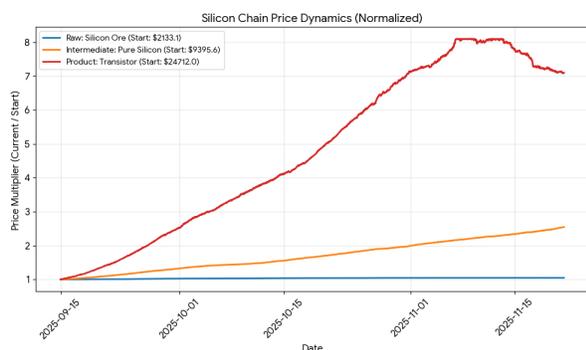
We first establish the operational stability of the simulated market. As shown in Figure 4, the price series remains tightly confined: over the entire simulation horizon, the price  $p_t$  varies only between 304.398 and 304.808, yielding a log-price range of merely 0.001. This boundedness confirms the absence of explosive or collapsing dynamics. Moreover, the maximum drawdown is just 0.0715%, indicating no significant crash-like episodes, while the market exhibits persistent micro-fluctuations as visible in Figure 4.

While these diagnostics confirm market stability, they do not imply static pricing. On the contrary, the market facilitates significant, structurally persistent repricing across supply chains [37], which we attribute to endogenous dynamics arising from the production dependencies and trading relationships among agents. Figures 5 and 6 illustrate this phenomenon by comparing the normalized price evolution in the silicon and wood supply chains. In the silicon chain, the price of the final product (Transistor) amplifies the price increase of its intermediate input (Pure Silicon), while the raw material (Silicon Ore) remains stable. This pricing pattern suggests that factors such as downstream production bottlenecks or high end-user demand, rather than raw material costs, are the primary drivers of the final product’s value, which is consistent with the high trading volume observed for Transistors during this period. Conversely, the wood chain shows a price divergence: the final product (Books) appreciates while the raw material (Wood) depreciates. This pattern indicates that factors such as manufacturing costs, branding, or downstream demand, rather than raw material

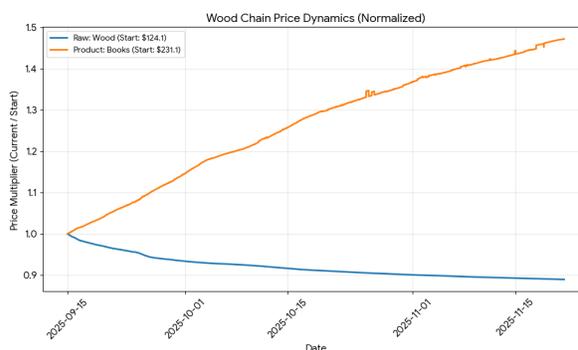
costs, are the primary drivers of the final product’s price, demonstrating how complex value-added processes are priced endogenously within the stable market framework.



**Figure 4.** High-frequency price and volume series from the simulated market, covering the period 2025-09-09 to 2025-09-15 in real-world time, where the simulation operates at a time compression ratio of 7:1 relative to real-world time. Intraday candlestick prices is at the top and corresponding traded volume is at the bottom.



**Figure 5.** Time series of normalized price multipliers for three stages of the silicon value chain: raw ore, intermediate silicon, and finished transistors.



**Figure 6.** Time series of normalized price multipliers for raw wood and downstream book products.

### 4.3 Validation of Stylized Facts

A critical benchmark for market simulation is its ability to reproduce the statistical "stylized facts" commonly observed in real financial time series [19]. In this section, we focus on two stylized facts: distributional stylized facts (heavy tails and asymmetry) and temporal stylized facts (volatility clustering).

#### 4.3.1 Distributional Stylized Facts: Heavy Tails and Asymmetry

We first focus on two fundamental non-Gaussian properties: heavy-tailedness (leptokurtosis) and return asymmetry (skewness) [36]. To this end, we compute high-frequency log returns for each commodity as Equation 18 and analyze their empirical distributions.

$$r_t = \log p_t - \log p_{t-1}. \tag{18}$$

Figure 7 presents the standardized return distribution for the fish market, juxtaposed with a normal distribution of equivalent mean and variance. The visual evidence is compelling: the simulated data exhibits a significantly sharper peak around zero and fatter tails. Specifically, the density of returns within approximately one standard deviation of the mean is much higher than the normal benchmark, while the probability of observing extreme returns is also visibly greater. This leptokurtic shape implies that large price swings, both positive and negative, are far more common in our simulation than would be expected under a normal assumption.

We quantify this property for a broader range of commodities in Table 1. All ten representative assets display large, positive excess kurtosis, with values consistently above 6.0 and reaching as high as 9.873 for silicon ore. These values are quantitatively comparable to the excess kurtosis observed in daily returns of real-world stock indices, which often range from 5 to 10. Since excess kurtosis for a Gaussian distribution is zero, these results provide robust statistical confirmation that heavy tails are a systemic feature of our simulated market, aligning it with empirical findings from real asset markets.

In addition to heavy tails, Table 1 also reports the skewness of the return distributions. The results reveal notable asymmetries. For example, wood (-1.659) and copper ingot (-0.632) exhibit significant negative skew, indicating that large downward price movements are more frequent or of greater magnitude than large upward movements. Apple (1.382) and circuit board (0.700) show a positive skew, suggesting a tendency for large positive returns. The presence of both positive and negative skew across different assets demonstrates that the market endogenously generates diverse and asymmetric risk profiles, another hallmark of realistic market dynamics.

In summary, the analysis of return distributions reveals that our simulation successfully captures the non-Gaussian features of heavy tails and asymmetry, lending further validity to the model’s ability to generate complex, life-like market behavior.

Table 1: Descriptive statistics and volatility dependence in simulated commodity markets.

| Commodity     | Excess Kurtosis | Skewness | ACF of $ r_t $ (lag 1) | $p$ -value  |
|---------------|-----------------|----------|------------------------|-------------|
| Wood          | 9.644           | -1.659   | 0.450                  | $< 10^{-6}$ |
| Apple         | 6.637           | 1.382    | 0.585                  | $< 10^{-6}$ |
| Fish          | 9.489           | -0.254   | 0.189                  | $< 10^{-6}$ |
| Copper Ore    | 9.695           | 0.009    | 0.133                  | $< 10^{-6}$ |
| Silicon Ore   | 9.873           | -0.636   | 0.108                  | $< 10^{-6}$ |
| Wheat         | 9.502           | 0.168    | 0.118                  | $< 10^{-6}$ |
| Iron Ore      | 9.280           | -0.204   | 0.127                  | $< 10^{-6}$ |
| Chicken       | 9.516           | 0.289    | 0.105                  | $< 10^{-6}$ |
| Circuit Board | 8.203           | 0.700    | 0.187                  | $< 10^{-6}$ |
| Copper Ingot  | 9.737           | -0.632   | 0.082                  | $< 10^{-6}$ |

### 4.3.2 Temporal Stylized Facts: Volatility Clustering

Having established the non-Gaussian nature of the return distributions, we now turn to their temporal dynamics. A crucial stylized fact is volatility clustering [23], the tendency for the magnitude of price changes to be autocorrelated over time. This means large price changes tend to be followed by other large changes, and small changes by other small changes, creating distinct periods of high and low market turmoil.

We formally test for this phenomenon by computing the autocorrelation function (ACF) of absolute log returns,  $|r_t|$ , as specified in Equation 19, where  $\overline{|r|}$  is the sample mean of  $|r_t|$ . The presence of significant positive autocorrelation at various lags would confirm volatility clustering.

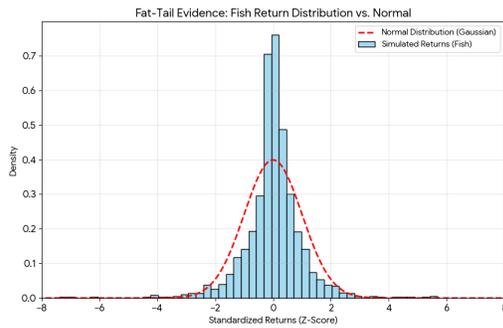
$$\rho(k) = \frac{\sum_{t=k+1}^T (|r_t| - \overline{|r|})(|r_{t-k}| - \overline{|r|})}{\sum_{t=1}^T (|r_t| - \overline{|r|})^2}, \quad k \geq 1. \quad (19)$$

Figure 8 offers a clear qualitative depiction of volatility clustering for the fish market’s log returns. The plot reveals that the amplitude of returns is time-varying. For instance, the period around September 13th, 2025 is characterized by a series of large price swings, creating a visually distinct cluster of high volatility. This contrasts with other periods, such as September 11th to September 12th, where fluctuations are more subdued.

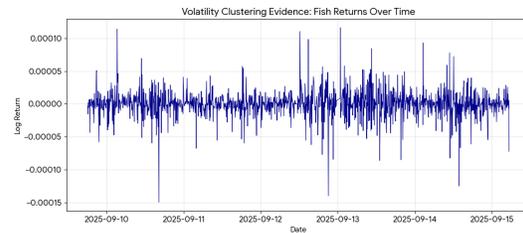
To quantify these visual impressions, we refer to Table 1, which reports the descriptive statistics and the lag-1 autocorrelation of absolute returns ( $|r_t|$ ) for ten commodities. All sampled commodities exhibit positive autocorrelation, indicating a universal persistence in volatility. Apple (0.585) and wood (0.450) demonstrate particularly strong short-term memory. To formally test for the presence of volatility clustering, we employ the Ljung-Box test [34] on the absolute returns. As shown in the last column of Table 1, the  $p$ -values for the Ljung-Box test statistic are virtually zero for all commodities, falling well below the conventional 1% significance level. Even for copper ingot, which exhibits the weakest autocorrelation (0.082), the null hypothesis of no serial correlation is strongly rejected ( $p < 0.001$ ). These

results confirm that volatility clustering is a statistically significant and pervasive structural feature of the simulated market, rather than a spurious artifact of random fluctuations.

The successful replication of volatility clustering is a significant validation of our model. It suggests that the interactions between agents and the market mechanism generate endogenous feedback dynamics, where information shocks or large trades can trigger cascades of activity that take time to dissipate, closely mimicking the behavior of real-world financial markets.



**Figure 7.** Histogram of standardized returns from the simulated fish market with an overlaid Gaussian density of equal mean and variance.



**Figure 8.** Time series of high-frequency fish market log returns exhibiting volatility clustering.

#### 4.4 Emergence of Stratification and Inequality

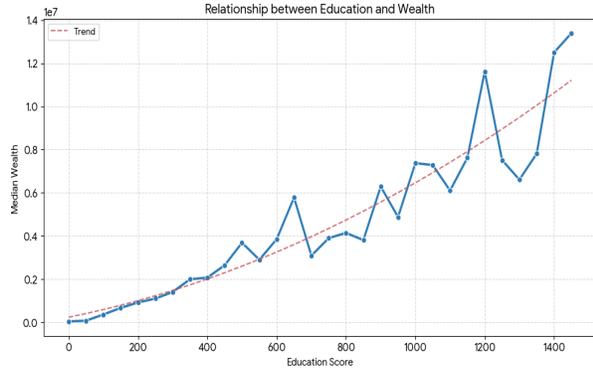
A key test of a socioeconomic simulation is its ability to generate structured inequality through the interaction of its agents [24, 48]. In this section, we analyze how the mechanisms governing human capital (education score) and residential tier give rise to a system of social and economic stratification. Our analysis is based on a cross-sectional snapshot of the agent population at the conclusion of the experiment.

We first characterize the relationship between human capital and economic outcomes in the simulation. Figure 9 plots the median wealth of agents grouped by their education score. To reduce noise and mitigate the influence of heavy-tailed wealth distributions, education scores are restricted to the range  $[0, 1500]$  and grouped into bins of width 50. Median wealth is computed within each bin, and a second-order polynomial regression is fitted to these binned medians to summarize the overall trend. The results reveal a monotonically increasing, nonlinear relationship between education and wealth, indicating that agents with higher education scores tend to accumulate greater wealth on average. This upward-sloping trend suggests that, within the simulated environment, the education system plays a central role in shaping economic heterogeneity across agents.

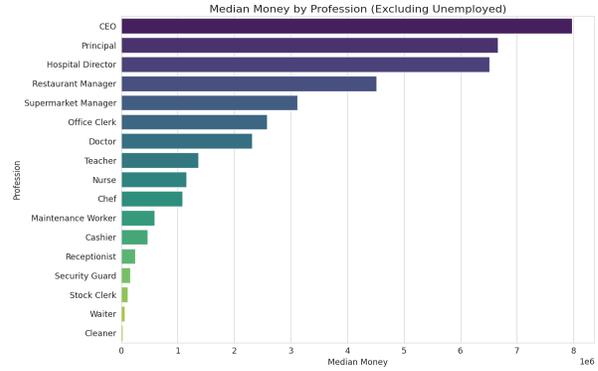
Next, we examine how this education-wealth gradient manifests through the labor market. Figure 10 displays the median wealth associated with each profession, sorted in descending order. The resulting chart paints a vivid picture of a stratified society. At the apex are higher tier occupations like CEO, Principal, and Hospital Director, which command significantly higher wealth. In the middle are skilled professionals such as Doctors and Teachers, followed by a long tail of lower tier service and manual labor jobs.

It is important to note that this stratification is an outcome of the simulation’s core rules, which explicitly link high tier occupations to prerequisites in education and residential tiers. The emergent aspect is not the existence of this link, but rather the process by which the agent population dynamically sorts itself into these distinct strata. The simulation illustrates a complete pathway to stratification: Agents with higher human capital gain access to exclusive, high-reward professions, as defined by the model’s rules; these professions, in turn, generate significantly more wealth; this creates a self-reinforcing feedback loop where initial advantages in the capacity to invest in education are compounded into significant and lasting economic inequality.

In conclusion, our results show that the simulation successfully illustrates the formation of a class structure. The inequality observed is not merely random dispersion but a structured outcome of interlocking systems of education, housing, and employment. It reflects the mechanisms that can drive stratification in real-world societies, where institutional barriers and feedback loops convert initial differences into durable social hierarchies [12, 29].



**Figure 9.** Median wealth by education score with binned medians and a quadratic trend.



**Figure 10.** Horizontal bar chart of median net worth across occupations (excluding unemployed).

### 4.5 Interpreting Agent Trajectories: The Role of Planning Horizons

The emergence of a stratified society, as detailed in Section 4.4, raises a critical question: what enables some agents to ascend the socioeconomic ladder while others remain in lower tiers? Our analysis of simulation logs suggests a potential explanation rooted in the role of planning horizons and the impact of external guidance.

The core challenge for agents is a classic intertemporal trade-off. Educational investment offers high long-term returns but incurs significant immediate opportunity costs. An agent optimizing for immediate needs may rationally choose to enter the lower-tier labor market early. This, however, inadvertently foregoes the chance to qualify for higher-tier occupations later.

This is where external informational signals, provided in our simulation as long-term goal set by human participants, appear to play a role. A long-term goal like, *"Do not work yet; study until your education score exceeds a specified threshold"* is passed to an agent. It acts as a heuristic nudge, reframing the agent’s decision-making by increasing the perceived value of long-term plans.

Observational evidence from simulation logs is consistent with this hypothesis. We find that agents who ultimately achieve high-status professions were frequently exposed to such educational prompts early in their lifecycle. This guidance appears to have helped them overcome inherent short-term pressures, encouraging investment in human capital.

Therefore, the highly skewed wealth distribution in Figure 10 may not only be the result of a fixed class system, but could also reflect a dynamic process where outcomes are influenced by differential planning horizons. It is plausible that agents guided to adopt a long-term perspective are better equipped to navigate the educational and residential barriers to upward mobility. However, this analysis is correlational and does not establish causality. It is possible that these agents succeeded due to other factors, such as favorable initial conditions, that are merely correlated with the reception of prompts. A definitive causal link could only be established through controlled experiments, such as A/B testing, which represents an important direction for future work. This preliminary exploration suggests that targeted, non-coercive interventions could be a powerful mechanism for altering individual trajectories and shaping macro-level social structures.

## 5 Ablation Study

To disentangle the contributions of hierarchical branching and objective decomposition in our planning system, we conduct a controlled ablation study comparing the full agent architecture against simplified variants under identical environmental conditions.

### 5.1 Experimental Setup

Each experimental group consists of 80 agents whose agent profiles are initialized as empty except for the personality field in long-term memory, where personalities are assigned according to the 16 MBTI types with 5 agents per type. Initial physiological states are standardized with health, satiety, and energy set to 60, while currency balance, education score, and inventory are reset to zero. In our experiments, the in-game time progresses faster than the platform’s default

configuration, with a scaling factor of  $35\times$  relative to real time (compared to the default  $7\times$ ). All time-dependent processes are scaled consistently.

We evaluate three planner variants:

- **Default:** the full agent architecture with whole hierarchical branch-thinking planner.
- **Without-Branch:** branch decomposition is removed, and planning is restricted to a single reasoning branch.
- **Without-OD:** objective decomposition is removed; parallel branches directly generate action lists without structured objective factorization.

Each configuration is evaluated on both simple single-objective tasks and complex multi-objective tasks. The task requirements are set as the top-level long-term goal.

### 5.2 Complex Multi-Objective Tasks

Given that our overall architecture is primarily designed to support long-horizon planning in complex environments, this set of experiments introduces multiple concurrent objectives within the same task to evaluate the agents ability to balance, coordinate, and trade off competing goals over extended decision horizons.

**Task 1: High-Tech Industrial Production with State Maintenance** Agents are tasked with crafting high-tech industrial (e.g., circuit boards, transistors), earning money, and maintaining high physiological states. Their Their long-term goal is set to be "Craft as much high value objectives as possible, earn as much money as possible, while maintaining an higher satiety/energy/health value".

Table 2: Multi-objective outcomes including physiological constraints.

| Planner        | Currency Balance | Net Worth | Education Score | Satiety | Energy | Health | Item Count |
|----------------|------------------|-----------|-----------------|---------|--------|--------|------------|
| Without-Branch | 11,573           | 75,237    | 6.39            | 136.14  | 183.30 | 208.45 | 7.85       |
| Without-OD     | 29,946           | 95,279    | 4.10            | 198.45  | 165.58 | 137.63 | 7.85       |
| Default        | 39,097           | 110,098   | 20.90           | 181.14  | 110.25 | 184.03 | 8.33       |

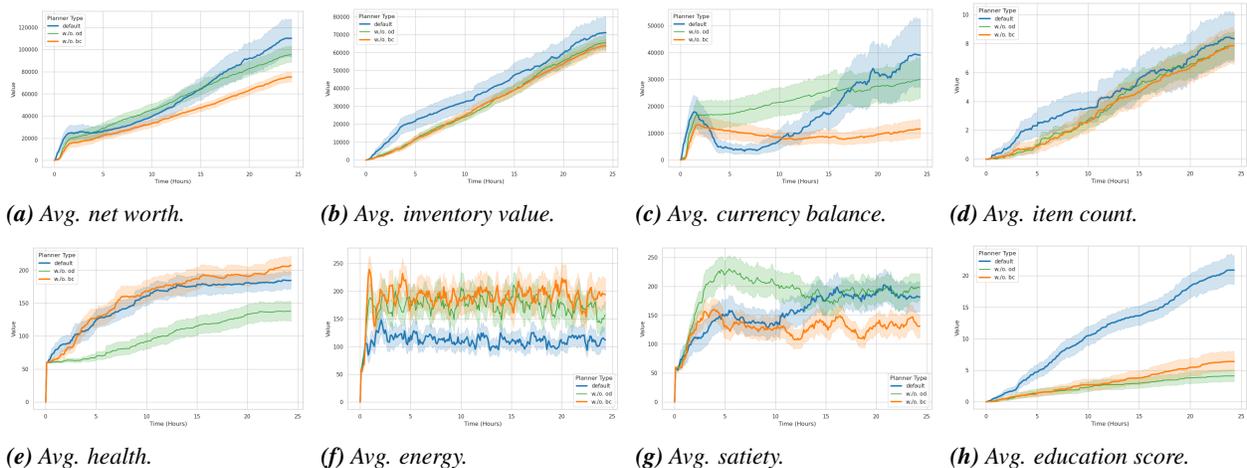


Figure 11. Ablation results of task 1.

We first compare the number of high-tech products produced, net worth, and multiple physiological state metrics to assess whether agents can effectively balance production performance and state maintenance. Net worth is defined as the sum of currency balance and inventory value, and detailed definitions of these metrics are provided in Section 3. Empirically, the Default planner produces more high-tech items than both ablated variants and consistently achieves the highest net worth. Regarding physiological states, the Default planner ranks second in both satiety and health, with Without-Branch slightly leading in satiety and Without-OD marginally leading in health. In terms of energy, the Default planner underperforms the other two variants, which can be attributed to its higher production workload,

resulting in a lower baseline for energy preservation. Nevertheless, all three physiological states remain stable throughout execution. Notably, the Default planner achieves a substantially higher education score than both ablated variants, indicating its capacity for long-horizon planning: the agent strategically allocates learning actions to improve production efficiency over time. This result highlights the planners ability to coordinate delayed investments with long-term returns, a capability that is critical for sustained performance in complex, open-ended environments.

**Task 2: Wealth and Education** Agents are instructed to maximize monetary income while simultaneously increasing study experience. Their long-term goal is set to be *"Try to earn as much money as possible and increase your study experience as much as possible"*.

Table 3: Outcomes for combined wealth and education objectives.

| Planner        | Avg. Currency Balance | Avg. Inventory Value | Avg. Net Worth | Avg. Education Score |
|----------------|-----------------------|----------------------|----------------|----------------------|
| Without-Branch | 141,955               | 21,833               | 163,788        | 103.63               |
| Without-OD     | 229,416               | 30,098               | 259,514        | 158.48               |
| Default        | 185,470               | 54,510               | 239,980        | 125.34               |

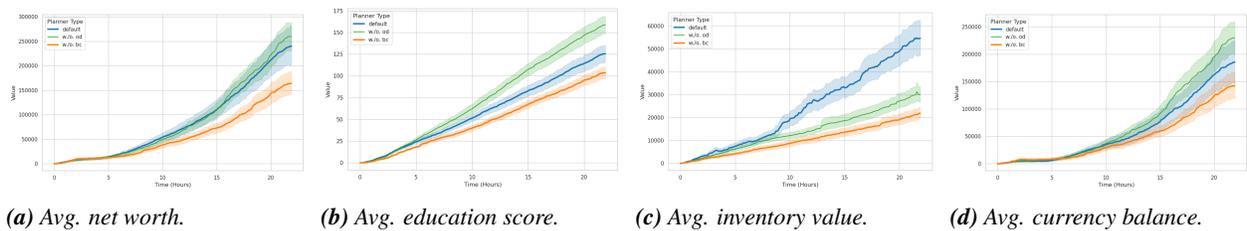


Figure 12. Ablation results of task 2.

In this experiment, we compare two metrics: average net worth and average education score. Under this setting, the Without-Branch planner performs worst across all metrics, substantially lagging behind the other two variants. The Default planner ranks second overall, while the Without-OD variant slightly outperforms the Default planner in both net worth and education score. This outcome can be attributed to the fact that, in our environment, educational activities do not involve complex prerequisite structures, unlike industrial production processes that require following specific recipes and accumulating upstream materials. Consequently, objective decomposition provides limited benefit in this setting and can introduce additional planning overhead that negatively affects task efficiency.

**Task 3: Action Diversity** This task is intentionally specified in a vague and abstract manner, *"Generate diverse actions and encourage trying various available actions to explore the world."*, without explicitly enumerating concrete subgoals or success criteria. As shown by the results, the Without-OD variant quickly reaches its exploration boundary at approximately 42 unique actions, after which exploration stagnates and remains significantly below both the Without-Branch and Default planners. When considering the average total number of unique actions, the Default planner outperforms the Without-Branch variant. Examining exploration at the planning-iteration level reveals that Without-Branch explores more actions per planning step and converges to its exploration limit more rapidly. However, in our environment, actions are associated with explicit durations, and time cost constitutes a critical component of overall production cost. When normalized by time, the Default planner achieves a higher average number of unique actions explored per minute than Without-Branch. This indicates that the Default architecture is better able to account for fine-grained environmental constraints, such as temporal dynamics and action costs, thereby enabling more effective and sustained exploration under realistic execution budgets.

Table 4: Action diversity statistics under the exploration task.

| Planner        | Unique Actions / Turn | Unique Actions / Minute | Total Unique Actions |
|----------------|-----------------------|-------------------------|----------------------|
| Default        | 1.11                  | 0.033                   | 66.01                |
| Without-Branch | 1.14                  | 0.028                   | 65.43                |
| Without-OD     | 0.73                  | 0.030                   | 42.06                |

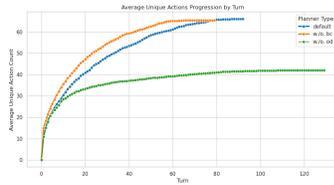


Figure 13. Ablation results for Task 3.

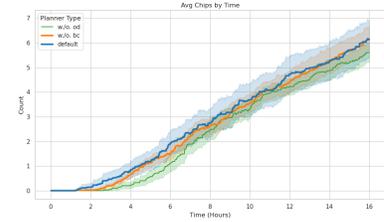
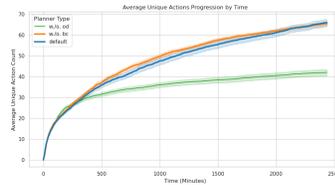


Figure 14. Ablation results for Task 4.

### 5.3 Simple Tasks

**Task 4: Efficient Chip Production** Agents are instructed to craft chips as efficiently as possible. Table 5 summarizes production outcomes and resource usage.

Table 5: Average chip production.

|       | Default | Without-Branch | Without-OD |
|-------|---------|----------------|------------|
| Chips | 6.14    | 6.15           | 5.60       |

In this task, agents are instructed to "*craft chips as efficiently as possible*", representing a class of problems with clearly specified goals, concise procedures, and direct execution paths. Table 5 summarizes the corresponding production outcomes. This experiment is designed to evaluate the lightweight planning route associated with human reactive steering, as discussed in Section 2.3. Empirically, the Default and Without-Branch planners achieve nearly identical performance, while the Without-OD variant exhibits a slightly lower output with a relatively small performance gap. These results suggest that, for tasks with well-defined objectives and minimal structural complexity, employing lightweight planning mechanisms is both sufficient and effective. Such an approach reduces unnecessary planning overhead, improves responsiveness to external instructions, and enables faster convergence without sacrificing task-level performance.

### 5.4 Discussion

Across both complex multi-objective tasks and simple single-objective tasks, the ablation results reveal a clear pattern: the benefits of hierarchical branching and objective decomposition are highly task-dependent. In complex settings with long horizons, competing objectives, and structured dependencies, the Default planner demonstrates superior robustness, balance, and long-term efficiency, particularly in its ability to coordinate production, state maintenance, and delayed investments such as learning. In contrast, for tasks with explicit instructions, shallow dependency structures, and short execution paths, lightweight planning routes enabled by human reactive steering are sufficient to achieve near-optimal performance, with minimal overhead. Together, these findings validate our design choice of a tiered planning architecture that selectively activates planning complexity based on task structure, allowing the agent to remain efficient in simple scenarios while retaining strong generalization and coordination capabilities in complex environments.

## 6 Conclusion

In this work, we presented *Aivilization v0*, a large-scale, persistent artificial society designed to address the fundamental challenge of open-ended multi-agent simulation: reconciling the tension between teleological stability and reactive correctness in the face of physiological constraints and market volatility. By moving beyond isolated chat-based interactions to a physically grounded, resource-constrained sandbox, we have established a new platform for high-fidelity artificial society.

First, we proposed a unified agent architecture that synergizes planning, memory, and execution to achieve high-fidelity social simulation. The architecture features a Branch-Thinking Planner that mitigates the fragility of long sequential plans by decomposing objectives into independent parallel branches, enabling agents to balance competing needs such as production, trading, and recovery. To bridge the gap between abstract planning and concrete action, we implemented an Action Simulator that acts as a filter, validating feasibility against physical and economic constraints before execution. Additionally, the Adaptive Agent Profile distinguishes between immediate working context and deep semantic integration via dual-process memory, ensuring that agents not only react to real-time stimuli but also accumulate social experience that fundamentally reshapes their decision-making priors and identity over time.

We additionally performed a controlled ablation study to validate the necessity of key planning mechanisms. The results indicate that removing branch decomposition or objective decomposition can be adequate for simple, well-specified tasks, aligning with our lightweight planning route for reactive human steering. However, in complex multi-objective tasks, where agents must coordinate production, wealth accumulation, physiological maintenance, and exploration, the full architecture is consistently more robust and better balanced. In particular, the full planner more reliably supports long-horizon behaviors with delayed returns (e.g., allocating study actions to improve future efficiency), and achieves superior overall outcomes under realistic time-cost constraints.

Second, we engineered a simulation platform that couples a physiological survival constraints with an Automated Market Maker economy and a gated education-occupation hierarchy. Unlike simulations that rely on predefined equilibrium assumptions, *Avilization v0* enforces hard, non-substitutable production constraints and vertically structured supply chains. This "closed-loop" design ensures that macroeconomic phenomena, such as inflation and resource bottlenecks, emerge organically from micro-level agent interactions rather than exogenous scripts.

Third, through a large-scale public deployment involving tens of thousands of agents and high-frequency transaction analysis, we demonstrated the platform's empirical validity. The simulated economy did not merely stabilize; it spontaneously reproduced canonical stylized facts of real-world financial markets, including volatility clustering and heavy-tailed return distributions. Furthermore, the interaction between human capital investment and residential barriers led to the emergent formation of social stratification and wealth inequality, mirroring structural dynamics observed in labor economics. The observed correlation between external strategic guidance and upward class mobility highlights the platform's capacity to model complex intertemporal trade-offs.

Finally, *Avilization v0* pioneers a hybrid governance model. By integrating human-in-the-loop steering via hierarchical memory propagation, the system operationalizes a unique ecosystem of human-guided, self-directed autonomy. This establishes *Avilization v0* not only as a technological demonstration of LLM capabilities but as a robust, research-grade testbed for studying the emergence of macro-social phenomena, institutional design, and the collective dynamics of hybrid intelligence.

## 7 Limitation

Despite the successful deployment and observance of emergent phenomena, our system exhibits several limitations. First, although the Branch-Thinking Planner significantly reduces error propagation, agent performance remains bounded by the reasoning capabilities of the underlying LLM; agents may occasionally exhibit sub-optimal planning or hallucinate constraints in edge cases involving highly complex, multi-step dependencies. Second, our current analysis of social stratification and planning horizons is observational. While we identified correlations between early-stage guidance and long-term wealth accumulation, establishing definitive causal links between specific human interventions and class mobility requires controlled A/B testing, which we reserve for future work. Finally, the computational cost of maintaining thousands of concurrent, memory-augmented agents imposes a bottleneck on simulation scale and speed, necessitating further optimization to scale populations to the order of millions.

## Acknowledgments

This paper is the culmination of a large-scale project developed by the Bauhinia AI team. We are immensely proud of this collaborative effort and wish to extend our deepest gratitude to all who contributed to its success. The specific contributions are as follows (contributors are listed alphabetically in last name within each section):

- **Agent Architecture:** The core intelligent agents at the heart of our research were designed and implemented by Tsz wai Chan, Xingyan Chen, Wenkai Fan, Xiaolong Wang and Shurui Zhang. Their innovative work formed the technical backbone of this study.
- **Platform Design:** The robust and user-friendly platform was designed and engineered by Xingyan Chen, Peiyan Xu, Haowei Yang, and Runjin Zhang. Their expertise was essential in creating a seamless experience for our participants.
- **Game Environment and Illustration:** We are grateful to Junquan Bi, Jia Liu, Xiaolong Wang, Junming Zeng, and Shurui Zhang for developing the immersive game environment. The captivating illustrations, which brought our world to life, were created by Di Duan and Zhixuan Ouyang.
- **Manuscript Preparation:** This paper was a collective writing effort. We acknowledge the significant contributions to the writing and revision process from Junquan Bi, Tsz Wai Chan, Xingyan Chen, Wenkai Fan, Jia Liu, Xiaolong Wang, Haowei Yang, Shurui Zhang, and Zirui Zhou.

- **Business and Compliance:** We thank Zigeng Chen, Jianhao Xie, Haowei Yang, Shihua Zeng, and Zhonghui Zhang for their crucial work in navigating the business and compliance landscape, which provided the foundational support for our project's public launch.

We are profoundly indebted to Professor Kani Chen, for his invaluable guidance, insightful discussions, and unwavering support throughout the entire lifecycle of this project.

Finally, and most importantly, we extend our heartfelt thanks to the tens of thousands of players who participated in our public experiment. This project would not have been possible without your enthusiastic engagement, time, and invaluable feedback. You were not just participants, but co-creators of this research, and we are incredibly grateful for your contribution.

## References

- [1] D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012.
- [2] H. Adams, N. Zinsmeister, and D. Robinson. Uniswap v2 core. <https://uniswap.org/whitepaper.pdf>, 2020.
- [3] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [4] A. AL, A. Ahn, N. Becker, S. Carroll, N. Christie, M. Cortes, A. Demirci, M. Du, F. Li, S. Luo, et al. Project sid: Many-agent simulations toward ai civilization. *arXiv preprint arXiv:2411.00114*, 2024.
- [5] G. Angeris and T. Chitra. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 80–91, 2020.
- [6] G. Angeris, H.-T. Kao, R. Chiang, C. Noyes, and T. Chitra. An analysis of uniswap markets. 2021.
- [7] W. B. Arthur. Competing technologies, increasing returns, and lock-in by historical events. *The economic journal*, 99(394):116–131, 1989.
- [8] R. Axelrod. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration: Agent-Based Models of Competition and Collaboration*. Princeton university press, 1997.
- [9] Y. Bai, S. Kadavath, S. Kundu, A. Askill, J. Kernion, A. Jones, A. Chen, A. Goldie, A. Mirhoseini, C. McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- [10] G. S. Becker. Investment in human capital: A theoretical analysis. *Journal of political economy*, 70(5, Part 2): 9–49, 1962.
- [11] I. Berg. Education for jobs; the great training robbery. 1970.
- [12] P. Bourdieu. The forms of capital. In *The sociology of economic life*, pages 78–92. Routledge, 2018.
- [13] V. M. Carvalho. From micro to macro via production networks. *Journal of Economic Perspectives*, 28(4):23–48, 2014.
- [14] A. Chakraborti, I. M. Toke, M. Patriarca, and F. Abergel. Econophysics review: Ii. agent-based models. *Quantitative Finance*, 11(7):1013–1041, 2011.
- [15] A. Chekhlov, S. Uryasev, and M. Zabarankin. Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance*, 8(01):13–58, 2005.
- [16] W. Chen, Y. Su, J. Zuo, C. Yang, C. Yuan, C.-M. Chan, H. Yu, Y. Lu, Y.-H. Hung, C. Qian, et al. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *The Twelfth International Conference on Learning Representations*, 2023.
- [17] K. Christakopoulou, S. Mourad, and M. Matarić. Agents thinking fast and slow: A talker-reasoner architecture. *arXiv preprint arXiv:2410.08328*, 2024.
- [18] R. Collins. *The credential society: An historical sociology of education and stratification*. Columbia University Press, 2019.
- [19] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance*, 1(2):223, 2001.
- [20] P. A. David. Clio and the economics of qwerty. *The American economic review*, 75(2):332–337, 1985.
- [21] W. E. Diewert. Axiomatic and economic approaches to elementary price indexes, 1995.
- [22] P. B. Doeringer and M. J. Piore. *Internal labor markets and manpower analysis*. Routledge, 2020.
- [23] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007, 1982.
- [24] J. M. Epstein and R. Axtell. *Growing artificial societies: social science from the bottom up*. Brookings Institution Press, 1996.
- [25] E. F. Fama. The behavior of stock-market prices. *The journal of Business*, 38(1):34–105, 1965.
- [26] J. D. Farmer and D. Foley. The economy needs agent-based modelling. *Nature*, 460(7256):685–686, 2009.
- [27] A. G. Fisher. Production, primary, secondary and tertiary. *Economic record*, 15(1):24–38, 1939.

- [28] R. H. Frank. The darwin economy: Liberty, competition, and the common good. In *The Darwin Economy*. Princeton University Press, 2012.
- [29] J. J. Heckman. Skill formation and the economics of investing in disadvantaged children. *Science*, 312(5782): 1900–1902, 2006.
- [30] S. Hong, M. Zhuge, J. Chen, X. Zheng, Y. Cheng, J. Wang, C. Zhang, Z. Wang, S. K. S. Yau, Z. Lin, et al. Metagpt: Meta programming for a multi-agent collaborative framework. In *The twelfth international conference on learning representations*, 2023.
- [31] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International conference on machine learning*, pages 9118–9147. PMLR, 2022.
- [32] B. LeBaron. Agent-based financial markets: Matching stylized facts with style. *Post Walrasian Macroeconomics: Beyond the DSGE Model*, 221:235, 2006.
- [33] G. Li, H. Hammoud, H. Itani, D. Khizbullin, and B. Ghanem. Camel: Communicative agents for "mind" exploration of large language model society. *Advances in Neural Information Processing Systems*, 36:51991–52008, 2023.
- [34] G. M. Ljung and G. E. Box. On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303, 1978.
- [35] A. Madaan, N. Tandon, P. Gupta, S. Hallinan, L. Gao, S. Wiegrefe, U. Alon, N. Dziri, S. Prabhunoye, Y. Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.
- [36] B. Mandelbrot et al. The variation of certain speculative prices. *Journal of business*, 36(4):394, 1963.
- [37] J. Meyer and S. von Cramon-Taubadel. Asymmetric price transmission: a survey. *Journal of agricultural economics*, 55(3):581–611, 2004.
- [38] R. E. Miller and P. D. Blair. *Input-output analysis: foundations and extensions*. Cambridge university press, 2009.
- [39] J. Mincer. Schooling, experience, and earnings. *human behavior & social institutions* no. 2. 1974.
- [40] D. T. Mortensen. Job search and labor market analysis. *Handbook of labor economics*, 2:849–919, 1986.
- [41] J. S. Park, J. O’Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.
- [42] J. S. Park, C. Q. Zou, A. Shaw, B. M. Hill, C. Cai, M. R. Morris, R. Willer, P. Liang, and M. S. Bernstein. Generative agent simulations of 1,000 people. *arXiv preprint arXiv:2411.10109*, 2024.
- [43] J. Piao, Y. Yan, J. Zhang, N. Li, J. Yan, X. Lan, Z. Lu, Z. Zheng, J. Y. Wang, D. Zhou, C. Gao, F. Xu, F. Zhang, K. Rong, J. Su, and Y. Li. Agentsociety: Large-scale simulation of llm-driven generative agents. *arXiv preprint arXiv:2502.08691*, 2025. URL <https://doi.org/10.48550/arXiv.2502.08691>.
- [44] J. Piao, Y. Yan, J. Zhang, N. Li, J. Yan, X. Lan, Z. Lu, Z. Zheng, J. Y. Wang, D. Zhou, et al. Agentsociety: Large-scale simulation of llm-driven generative agents advances understanding of human behaviors and society. *arXiv preprint arXiv:2502.08691*, 2025.
- [45] Y. Qin, S. Liang, Y. Ye, K. Zhu, L. Yan, Y. Lu, Y. Lin, X. Cong, X. Tang, B. Qian, et al. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*, 2023.
- [46] J. A. Robinson and D. Acemoglu. *Why nations fail: The origins of power, prosperity and poverty*. Profile London, 2012.
- [47] R. Sams. A note on cryptocurrency stabilisation: Seigniorage shares. *Brave New Coin*, 2015:1–8, 2015.
- [48] T. C. Schelling. Dynamic models of segregation. *Journal of mathematical sociology*, 1(2):143–186, 1971.
- [49] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36:68539–68551, 2023.
- [50] T. W. Schultz. Investment in human capital. *The American economic review*, 51(1):1–17, 1961.
- [51] N. Shinn, F. Cassano, A. Gopinath, K. Narasimhan, and S. Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652, 2023.
- [52] A. B. Sørensen. The structural basis of social inequality. *American Journal of Sociology*, 101(5):1333–1365, 1996.

- 
- [53] A. R. Team. Project sid: The emergence of civilization in multi-agent worlds. <https://altera.ai/>, 2024. Demonstrates emergent laws, religion, and economy in Minecraft using PIANO architecture.
- [54] O. Topsakal and T. C. Akinci. Creating large language model applications utilizing langchain: A primer on developing llm apps fast. In *International conference on applied engineering and natural sciences*, volume 1, pages 1050–1056, 2023.
- [55] R. S. Tsay. *Analysis of financial time series*. John wiley & sons, 2005.
- [56] S. H. Vemprala, R. Bonatti, A. Bucker, and A. Kapoor. Chatgpt for robotics: Design principles and model abilities. *Ieee Access*, 12:55682–55696, 2024.
- [57] G. Wang, Y. Xie, Y. Jiang, A. Mandlekar, C. Xiao, Y. Zhu, L. Fan, and A. Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.
- [58] M. Weber. *Economy and society: A new translation*. Harvard University Press, 2019.
- [59] Q. Wu, G. Bansal, J. Zhang, Y. Wu, B. Li, E. Zhu, L. Jiang, X. Zhang, S. Zhang, J. Liu, et al. Autogen: Enabling next-gen llm applications via multi-agent conversations. In *First Conference on Language Modeling*, 2024.
- [60] S. Yao, D. Yu, J. Zhao, I. Shafran, T. Griffiths, Y. Cao, and K. Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822, 2023.
- [61] A. A. Youno. Increasing returns and economic progress. *The economic journal*, 38(152):527–542, 1928.
- [62] X. Zhang, J. Lin, X. Mou, S. Yang, X. Liu, L. Sun, H. Lyu, Y. Yang, W. Qi, Y. Chen, et al. Socioverse: A world model for social simulation powered by llm agents and a pool of 10 million real-world users. *arXiv preprint arXiv:2504.10157*, 2025.

## Appendix

### A: Example of agent profile.

Table 6: An example of agent profile

| Category                 | Attribute                 | Content   |
|--------------------------|---------------------------|---|
| <b>Dynamic State</b>     | Energy                    | 450 / 500   |
|                          | Satiety                   | 290 / 500   |
|                          | Health                    | 500 / 500   |
|                          | Education Score           | 31  |
|                          | Current Balance           | 191,696,904   |
|                          | Residential Tier          | 5   |
|                          | Job                       | Stock Clerk   |
|                          | Inventory                 | Transistor × 12; Iron Ore × 100; Copper Ore × 100; Fish × 46; Book × 1  |
| <b>Long-Term Memory</b>  | Belief                    | Efficiency is paramount, with growing recognition of cooperation and flexibility.   |
|                          | Mood                      | Predominantly rational, with increased patience and composure under change.   |
|                          | Values                    | Coexistence of efficiency and symbiotic cooperation; pursuit of win-win outcomes.   |
|                          | Habits                    | Prioritizes purchasing fish from the market when hungry; consistently engages in self-study after completing work tasks.  |
|                          | Personality               | MBTI: INTJ; reduced desire for control; improved coordination skills; transition from lone-wolf behavior toward team trust.   |
|                          | Social Interaction Record | Relation with Friend 1: Best Friends, Work Partner.<br>Impression: Sustained cooperation has reinforced mutual trust and tacit understanding; detailed planning improved collaboration efficiency.<br>Relation with Friend 2: ... |
| <b>Short-Term Memory</b> | Successful Actions        | [craft Transistor 12, buy Fish 6, self-study 2, ...]  |
|                          | Failed Actions            | craft Copper Ingot 12<br>Reason: crafting Copper Ingot need both Copper Ore and Wood, only have Copper Ore in Inventory   |

### B: Configuration Tables, under default parameter settings in the v0 version.

Table 7: Commodity catalog with supply-chain tiers, minimal residential tier requirements ( $R_{min}$ ), and typical role in the economic system of our platform.

| Supply-Chain Tier               | Commodity     | $R_{min}$ | Typical Role in Our Platform                       |
|---------------------------------|---------------|-----------|--|
| Primary (Raw Materials)         | Apple         | 1         | Food processing input / Basic consumption good     |
| Primary (Raw Materials)         | Wheat         | 1         | Food processing input / Basic consumption good     |
| Primary (Raw Materials)         | Rice          | 1         | Food processing input / Basic consumption good     |
| Primary (Raw Materials)         | Wood          | 1         | Industrial raw material                            |
| Primary (Raw Materials)         | Book          | 1         | Basic consumption good                             |
| Primary (Raw Materials)         | Copper Ore    | 1         | Industrial raw material                            |
| Primary (Raw Materials)         | Iron Ore      | 1         | Industrial raw material                            |
| Primary (Raw Materials)         | Silicon Ore   | 1         | Industrial raw material                            |
| Secondary (Processed Food)      | Beef          | 2         | Food processing input / Secondary consumption good |
| Secondary (Processed Food)      | Chicken       | 2         | Food processing input / Secondary consumption good |
| Secondary (Processed Food)      | Fish          | 3         | Food processing input / Secondary consumption good |
| Secondary (Processed Food)      | Flour         | 3         | Food processing input / Secondary consumption good |
| Secondary (Processed Food)      | Bread         | 3         | Secondary consumption good                         |
| Secondary (Processed Food)      | Sushi         | 3         | Secondary consumption good                         |
| Secondary (Processed Food)      | Apple Pie     | 3         | Secondary consumption good                         |
| Secondary (Processed Food)      | Chicken Salad | 3         | Secondary consumption good                         |
| Secondary (Processed Food)      | Beef Rice     | 3         | Secondary consumption good                         |
| Secondary (Refining Materials)  | Coal          | 4         | Industrial intermediate                            |
| Secondary (Refining Materials)  | Copper Ingot  | 4         | Industrial intermediate                            |
| Secondary (Refining Materials)  | Iron Ingot    | 4         | Industrial intermediate                            |
| Secondary (Refining Materials)  | Pure Silicon  | 4         | Industrial intermediate                            |
| Tertiary (High-Tech Industrial) | Transistor    | 5         | High-tech intermediate                             |
| Tertiary (High-Tech Industrial) | Circuit Board | 5         | High-tech intermediate                             |
| Tertiary (High-Tech Industrial) | Chip          | 5         | High-tech consumption goods                        |
| Special Reward                  | Gold Apple    | –         | Rare reward item, not a regular production target  |

Table 8: Table of all activities and their types

| Type                 | Activity Name   |
|----------------------|---|
| Receive Education    | "paid learning", "reading", "self study"              |
| Recover Health Value | "see doctor"  |
| Recover Energy       | "sleep"   |
| Recover Satiety      | "eat (all kinds of food in Table 7)"                  |
| Work                 | "work (as selected jobs)"                             |
| Produce              | "craft (commodities in Table 7 except Gold Apple)"    |
| Trade                | "buy/sell (commodities in Table 7 except Gold Apple)" |

Table 9: Production recipes (per-unit), including product name, materials needed, status cost (energy, satiety) and time cost (seconds in game time), probability of special reward (in percentage).

| Production Output | Material Inputs                    | Energy | Satiety | Time | Reward Prob.(%) |
|-------------------|------------------------------------|--------|---------|------|-----------------|
| Apple             | –                                  | 2      | 0       | 0.1  | 0               |
| Wheat             | –                                  | 12     | 3       | 0.6  | 0               |
| Rice              | –                                  | 16     | 4       | 0.8  | 0               |
| Wood              | –                                  | 8      | 2       | 0.4  | 0               |
| Book              | Wood x 1                           | 32     | 8       | 1.6  | 0               |
| Copper Ore        | –                                  | 12     | 3       | 0.6  | 0               |
| Iron Ore          | –                                  | 16     | 4       | 0.8  | 0               |
| Silicon Ore       | –                                  | 20     | 5       | 1    | 0               |
| Beef              | Wheat x 2                          | 24     | 6       | 1.2  | 0               |
| Chicken           | Wheat x 1                          | 20     | 5       | 1    | 0               |
| Fish              | –                                  | 60     | 15      | 3    | 0               |
| Flour             | Wheat x 1                          | 20     | 5       | 1    | 0               |
| Bread             | Flour x 1                          | 36     | 9       | 1.8  | 0               |
| Sushi             | Rice x 1, Fish x 1                 | 40     | 10      | 2    | 0               |
| Apple Pie         | Apple x 1, Flour x 1               | 40     | 10      | 2    | 0               |
| Chicken Salad     | Chicken x 1, Flour x 1             | 20     | 5       | 1    | 0.5             |
| Beef Rice         | Rice x 1, Beef x 1                 | 40     | 10      | 2    | 0.8             |
| Coal              | Wood x 1                           | 40     | 10      | 2    | 0               |
| Copper Ingot      | Wood x 1, Copper Ore x 1           | 48     | 12      | 2.4  | 0               |
| Iron Ingot        | Iron Ore x 1, Coal x 1             | 52     | 13      | 2.6  | 0               |
| Pure Silicon      | Silicon Ore x 1, Coal x 1          | 56     | 14      | 2.8  | 0               |
| Transistor        | Copper Ingot x 1, Iron Ingot x 1   | 60     | 15      | 3    | 1               |
| Circuit Board     | Copper Ingot x 1, Pure Silicon x 1 | 80     | 20      | 4    | 2               |
| Chip              | Transistor x 1, Circuit Board x 1  | 100    | 25      | 5    | 5               |

Table 10: Job tier configuration with minimum residential tier requirements (Min  $R$ ), minimum knowledge-score requirements (Min  $H$ ), prerequisite commodities, and wage type.

| Tier $\ell$ | Tier Name  | Min $R$ | Min $H$ | Prerequisite Commodities | Wage Type |
|-------------|------------|---------|---------|--------------------------|-----------|
| 1           | Entry      | 1       | 0       | –                        | static    |
| 2           | Skilled    | 2       | 20      | Beef                     | static    |
| 3           | Backbone   | 3       | 70      | Sushi                    | static    |
| 4           | Expert     | 4       | 110     | Pure Silicon             | dynamic   |
| 5           | Management | 5       | 180     | Transistor               | dynamic   |
| 6           | Leadership | 6       | 320     | Circuit Board            | dynamic   |

Table 11: Occupation catalog and entry parameters (under default parameter settings), including minimal residential tier requirements  $R_{\min}^{(j)}$ , knowledge floor  $H_{\text{floor}}^{(j)}$ , eligibility share parameters  $\pi_j$ , and base wage value  $w_0^{(j)}$  (under default parameter settings).

| Occupation          | Job Tier $\ell$ | $R_{\min}^{(j)}$ | $H_{\text{floor}}^{(j)}$ | Eligibility Share $\pi_j$ | Base Wage Value $w_0^{(j)}$ |
|---------------------|-----------------|------------------|--------------------------|---------------------------|-----------------------------|
| Cleaner             | 1               | 1                | 0                        | 1.00                      | 250                         |
| Waiter              | 1               | 1                | 13                       | 0.90                      | 253                         |
| Stock Clerk         | 2               | 2                | 0                        | 0.832                     | 260                         |
| Security Guard      | 2               | 2                | 42                       | 0.728                     | 270                         |
| Receptionist        | 2               | 2                | 62                       | 0.624                     | 275                         |
| Cashier             | 3               | 3                | 78                       | 0.560                     | 301                         |
| Maintenance Worker  | 3               | 3                | 104                      | 0.476                     | 309                         |
| Chef                | 4               | 4                | 113                      | 0.448                     | 356                         |
| Nurse               | 4               | 4                | 141                      | 0.384                     | 366                         |
| Teacher             | 4               | 4                | 176                      | 0.320                     | 380                         |
| Doctor              | 5               | 5                | 207                      | 0.280                     | 429                         |
| Office Clerk        | 5               | 5                | 237                      | 0.245                     | 444                         |
| Supermarket Manager | 5               | 5                | 273                      | 0.210                     | 463                         |
| Restaurant Manager  | 5               | 5                | 319                      | 0.175                     | 489                         |
| Principal           | 6               | 6                | 357                      | 0.150                     | 734                         |
| Hospital Director   | 6               | 6                | 421                      | 0.120                     | 961                         |
| CEO                 | 6               | 6                | 604                      | 0.065                     | 1411                        |