# Scaling Multi-Agent Systems with Prospect-State Propagation

**Anonymous ACL submission**


## Abstract

Current LLM-based multi-agent systems (MAS) periodically compress intermediate states to reduce inference-time token consumption, thereby attempting to incorporate more agents. However, this scaling strategy faces challenges in economic simulation, because it discards semantic-rich economic states, *i.e.*, agent behavioral trajectories, which are key drivers of macroeconomic fluctuations. In this paper, we reveal a phenomenon in which agent heterogeneity gradually decreases during simulation, and propose **P**rospect-**S**tate **P**ropagation for **M**ulti-**A**gent **S**ystems (PspMAS). Inspired by prospect theory, PspMAS decouples each agent's micro state into a compact Prospect State and an expressive Semantic State. The former records psychological traces through a lightweight, parallelizable propagator and continuously injects heterogeneity into the system. The latter leverages the strong perception, reasoning, planning, and decision-making abilities of LLMs. These two components work complementarily, reducing computational cost while improving the scalability of MAS. Our implementation shall be released upon acceptance.


## 1 Introduction

*Humans evaluate potential gains and losses relative to a specific reference point rather than focusing on absolute wealth, ..., is Prospect Theory.*
*– Daniel and Amos (1979)*

From Adam Smith's *"invisible hand"* in the 18th century to the *behavioral economics* revolution of the 20th century, the study of Macroeconomics has increasingly highlighted the importance of modeling the complexity of human behavior. Traditional representative-agent models (Kirman, 1992; Blanchard, 2017; Christiano et al., 2005) largely circumvent this difficulty by assuming a "perfect world" populated by a single average consumer or firm. Such assumptions, however, eliminate distributional heterogeneity and the non-linear aggregate fluctuations (Gabaix, 2011) it generates,

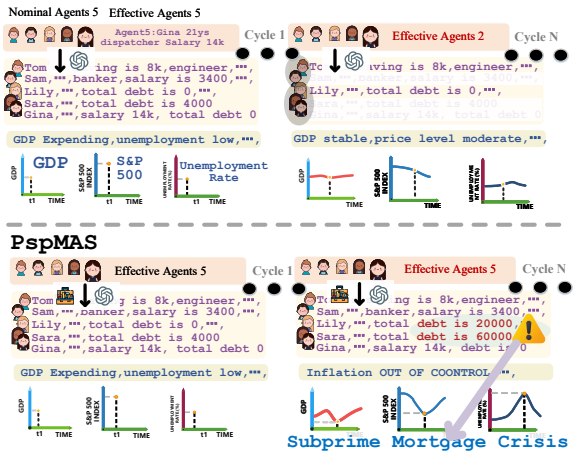


Figure 1: The top sub-figure shows a traditional simulation system. After long-horizon simulation, at cycle $N$, the heterogeneity of agents decreases, and the macroeconomic fluctuations are consequently weakened. The bottom sub-figure shows PspMAS. By preserving heterogeneity, the sharp increase in personal mortgage debt for some agents dominates the description of their micro states, thereby successfully simulating the subprime mortgage crisis.

limiting the ability to replicate typical macroeconomic events (*e.g.*, crises and recessions).

Agent-based modeling (ABM) offers a bottom-up solution (Farmer and Foley, 2009; Tesfatsion and Judd, 2006). By simulating the behaviors and interactions of heterogeneous agents, ABM allows stylized facts (*e.g.*, GDP, inflation) to emerge from the micro level. Rule-based ABMs maintain a large number of agents with their economic microstate (*e.g.*, income, savings, debt, employment status, wage) that are updated via hand-crafted functions, one representative microstate in the form of key-value pair is `'agent:1, income:3200, saving:9000, unemployment status:0, ...'`. Such micro states can be batched, parallelized, and vectorized, enabling hardware-efficient simulation. However, the pre-defined rules fail to capture real psychological processes, ultimately hindering the replication of key macroeconomic dynamics driven by complex human cognition.

Large language models (LLMs) have recently shown remarkable capabilities in natural language understanding (OpenAI, 2023; Liang et al., 2022; Zheng et al., 2023), role-playing (Wang et al., 2024b), and contextual reasoning (Yao et al., 2023; Zhu et al., 2024; Dalal et al., 2024), making them attractive for social simulation (Horton, 2023; Argyle et al., 2023; Li et al., 2024b; Jia et al., 2024; Binz et al., 2025; Wang et al., 2024b; Zou et al., 2025; Li et al., 2023). The microstate of an LLM-based agent is no longer a fixed set of numbers but a rich semantic description, for example: 👨‍🦰 `'You're Tom, an engineer aged 25. Your saving is 8000, your salary is 3200, and total debt is 0, ..., CPI indicates current economy is in a high-inflation cycle. What level of expenditure do you plan for next period?'` Such semantic-rich microstates enable sufficient behavioural heterogeneity, which is a key driver of macroeconomic fluctuations. But long descriptive profiles of agents during long-horizon simulation consume substantial inference-time token budgets, hindering scalability to large agent populations. Existing multi-agent systems resort to periodic reflection (Park et al., 2023; Shinn et al., 2023), compressing economic histories into high-level summaries. While effective in multi-turn conversations(Qian et al., 2024; Yang et al., 2024; Wu et al., 2024a; Wang et al., 2023) or code generation tasks, such semantic compression strategies discard the temporal order and intermediate states that are essential for economic decision-making. From the perspective of behavioral economics, human behavior is highly sensitive to early economic fluctuations, and even small changes can exert lasting influences on subsequent choices. Repeated summarization across simulation cycles tends to decrease the dynamic heterogeneity of agents, leading to behavioral homogenization—a phenomenon called *action convergence*. As illustrated in the top part of Fig. 1, a traditional system injects rich descriptions into all agents at system initialization, thereby creating a large nominal population of heterogeneous agents. However, as the simulation proceeds, the micro states of agents are compressed into high-level summaries on a recurring basis. This process gradually removes individual behavioral differences and greatly reduces diversity among agents. Consequently, macroeconomic fluctuations that should emerge from heterogeneous local distress become attenuated. In macroeconomics, risk signals with genuine predictive value often first appear in the local states of a small subset of agents. If state descriptions are dominated by lagging macro variables such as the unemployment rate, these localized abnormalities can be diluted or erased during semantic compression. As a result, most agents are guided toward a homogenized common context in which the economy still appears broadly healthy, suppressing the micro-to-macro propagation through which crises emerge. For example, in the buildup to the subprime mortgage crisis, early risks were not immediately reflected in aggregate labor-market indicators. Instead, stress first accumulated in agents with average wages but high debt burdens, such as `agent Sara` 👩🏾 shown in the bottom part of Fig. 1. Such agents can transmit early default risk and other key local signals that eventually develop into the subprime mortgage crisis. If these loan-related micro states are compressed away, the system loses the heterogeneous mechanisms through which localized distress propagates into broader macroeconomic fluctuations.

Based on the above observations, we propose **P**rospect-**S**tate **P**ropagation for **M**ulti-**A**gent **S**ystems (`PspMAS`). `PspMAS` decouples each agent's micro state into two complementary parts: a compact *Prospect State* and an expressive *Semantic State*. The design of *Prospect State* is inspired by prospect theory, a classical theory in behavioral economics for characterizing human decision-making under gains, losses, and uncertainty. It maintains key psychological variables, including reference point, loss aversion, and probability weighting, and updates them through a fast propagator. *Prospect State* provides the long-term micro-level heterogeneity necessary for economic emergence. *Semantic State* leverages the perception, reasoning, planning, and decision-making abilities of LLMs to generate human-like actions. By combining these two states, `PspMAS` avoids relying on long textual state descriptions, making the system easier to scale, while also mitigating the heterogeneity loss caused by long-horizon semantic compression. Our contributions are summarized as follows:

- We reveal an inherent heterogeneity-loss problem in MAS for economic simulation, where periodical semantic compression gradually reduces the effective number of behaviorally distinct agents.
- We incorporate prospect theory, preserving long-term heterogeneity.
- We enable scalable LLM-based multi-agent simulation under fixed inference-token budgets.

## 2 Method

This section introduces PspMAS, a framework for preserving agent heterogeneity while scaling LLM-based multi-agent systems over long-horizon simulation. The key idea is to decouple each agent's micro state into two complementary parts: a Markovian Prospect State updated by a lightweight Prospect-State Propagator, and a Semantic State updated periodically by an LLM. We first provide background on multi-agent systems for macroeconomic simulation (Section 2.1), then introduce the simulation loop, including economic state construction, prospect-state propagation, economic state read, and system scaling. Implementation details are given in Section 2.3.

### 2.1 Preliminary

In a Multi-agent Systems (MAS) (Li et al., 2024b; Feng et al., 2025), agents perceive the environment and take actions to achieve goals; their behavior is not hard-coded but emerges from situated interactions with other agents and the environment. We consider a MAS for macroeconomic simulation populated by $N$ heterogeneous agents, indexed by $\mathcal{I} = \{1, 2, \ldots, N\}$. The corresponding agent set is $\{A_i\}_{i\in\mathcal{I}}$. Each agent is equipped with LLM-powered perception, reasoning, planning, and decision-making abilities. During simulation, agents interact, collectively constructing a macroeconomy, which in turn updates their perceptions and influences subsequent actions.

### 2.2 PspMAS

**Economic State Construction.** Micro differences of agents are a key driver of macroeconomic fluctuations. The combination of agents' micro states and the environment's macro state constitutes the Economic State of the simulation.

*Micro State*. Each agent $A_i$ at time $t$ maintains an independent micro state capturing its internal behavioral state and economic history. In traditional rule-based ABMs, the micro state is a fixed-dimensional structured vector (Kuroki et al., 2025). In LLM-based simulations, the micro state typically comprises a compact numerical component and a rich semantic component (Li et al., 2024b). In PspMAS, we write the micro state as

$$\mathcal{M}_i^t = \left(s_i^t,\ \sigma_i^t,\ e_i^t\right), \tag{1}$$

where $s_i^t \in \mathcal{P}$ is the Prospect State, $\sigma_i^t \in \mathcal{M}_{\text{sem}}$ is the Semantic State, and $e_i^t$ denotes economic bookkeeping variables (e.g., wealth, debt, tax paid, wage, realized consumption, employment status) that are updated by the environment via market clearing. In Particular, the two central branches of PspMAS are $s_i^t$ and $\sigma_i^t$: the former preserves compact psychological traces, while the latter supports expressive natural-language reasoning.

*Macro State*. The overall condition of the economic environment at time $t$ is described by a macro state $\mathcal{S}^t$, typically including price level $P^t$, interest rate $r^t$, unemployment rate $u_{\text{emp}}^t$, and total GDP $Y^t$. Formally, $\mathcal{S}^t \in \mathcal{Y}$. The macro state reflects the aggregation of all agents' actions through market mechanisms, forming a closed feedback loop (Li et al., 2024b; Yang et al., 2026).

*Agent Action*. At each time step $t$, each agent $A_i$ selects an action $a_i^t$ from its action space $\mathbb{A}_i$ based on its current micro state $\mathcal{M}_i^t$ and the macro state $\mathcal{S}^t$. Following macroeconomic simulation literature, the action space includes two fundamental decisions: whether to work ($l_i^t \in \{0, 1\}$) and what fraction of available funds to spend on consumption ($p_i^t \in [0, 1]$) (Li et al., 2024b; Feng et al., 2025). The action space can be expressed as

$$a_i^t = (l_i^t, p_i^t) \in \mathbb{A}_i = \{0, 1\} \times [0, 1]. \tag{2}$$

**Prospect-State Propagation.** We incorporate prospect theory (Kahneman and Tversky, 2013) into the micro state to capture bounded rationality in decision-making. The theory highlights three key features: - *Reference dependence*: decisions are made based on gains and losses relative to a psychological reference point (e.g., initial wealth), rather than absolute wealth. - *Loss aversion*: losses are felt approximately twice as strongly as equivalent gains (i.e., $\lambda \approx 2.25$). - *Probability weighting*: people tend to overweight small probabilities and underweight moderate to high probabilities, leading to distorted subjective probabilities.

Based on these insights, we define the Prospect State of agent $i$ at time $t$ as a Prospect-State Vector:

$$s_i^t = \left(\lambda_i^t,\ \gamma_i^t,\ R_i^t\right) \in \mathcal{P} \subset \mathbb{R}^3. \tag{3}$$

Here $R_i^t$ is the dynamic reference point, $\lambda_i^t$ the loss-aversion coefficient, and $\gamma_i^t$ the probability-weighting parameter. Other economic quantities, such as wealth, income, realized consumption, and unmet demand, are stored as accounting variables or immediate outcomes rather than as dimensions of the Prospect-State Vector.

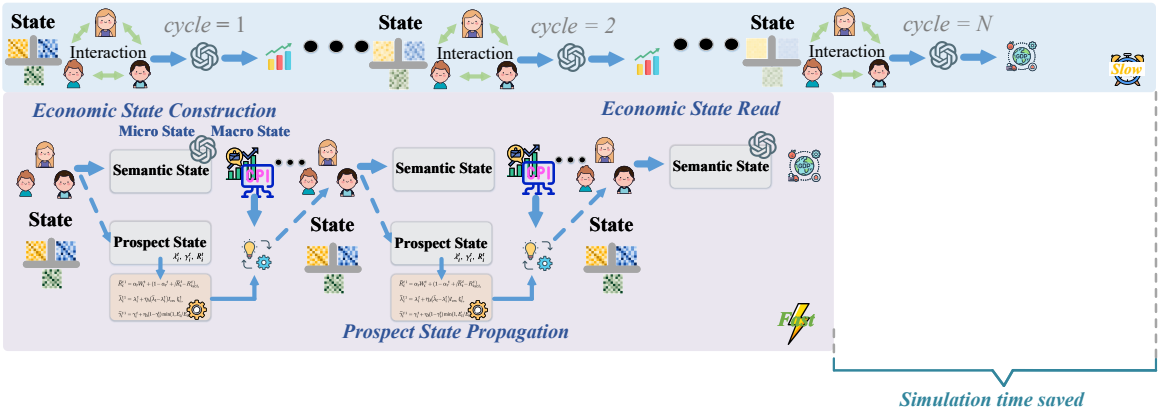


Figure 2: Overview of the major components of PspMAS. PspMAS decouples each agent's micro state into two complementary branches: a compact Prospect State and an expressive Semantic State. The Prospect State is continuously updated by a lightweight Prospect-State Propagator, preserving psychological traces and injecting long-term heterogeneity into the agent population. The Semantic State is periodically refreshed by the LLM to support perception, reasoning, planning, and human-like decision-making. PspMAS preserves heterogeneous micro dynamics while enabling scalable long-horizon macroeconomic simulation.

Let $o_i^t$ denote the immediate individual outcome produced by market clearing at time $t$ (e.g., realized labor income, realized consumption, savings change, and unmet demand), and let $\mathcal{S}^{t+1}$ denote the realized macro state after clearing. The Prospect State update is Markovian:

$$\tilde{s}_i^{t+1} = \Phi(s_i^t, o_i^t, \mathcal{S}^{t+1}), \qquad s_i^{t+1} = \Pi_{\mathcal{P}}(\tilde{s}_i^{t+1}), \tag{4}$$

where $\Pi_{\mathcal{P}}$ projects the updated Prospect-State Vector onto a predefined valid domain $\mathcal{P}$. This projection step, rather than the coefficient ranges alone, keeps the Prospect State bounded throughout long-horizon simulation.

The Markovian Prospect-State Propagator $\Phi$ consists of deterministic, parallelizable lightweight rules:

$$\begin{aligned}
\tilde{R}_i^{t+1} &= \alpha_r W_i^t + (1-\alpha_r) R_i^t + \beta_r (\bar{R}_t^c - R_i^t) I_{\text{scl}}^t, \\
\tilde{\lambda}_i^{t+1} &= \lambda_i^t + \eta_\lambda (\bar{\lambda}_t - \lambda_i^t) I_{\text{crs}}^t + \xi_i^t, \\
\tilde{\gamma}_i^{t+1} &= \gamma_i^t + \eta_\gamma (1-\gamma_i^t) \min(1, E_i^t / E_{\text{thr}}) - \zeta I_{\text{shk}}^t,
\end{aligned} \tag{5}$$

where $W_i^t$ is total wealth, $\bar{R}_t^c$ the cohort-average reference point, and $I_{\text{scl}}^t$, $I_{\text{crs}}^t$, and $I_{\text{shk}}^t$ are indicator flags for social comparison, crisis, and shock, respectively. The variable $E_i^t$ denotes a recent experience intensity derived from the immediate outcome $o_i^t$, and $E_{\text{thr}}$ is a normalization threshold. After computing Eq. 5, the projection in Eq. 4 is then applied component-wise to obtain the next Prospect-State Vector, which is denoted by $s_i^{t+1}$.

The Semantic State $\sigma_i^t$ is a textual summary of recent events and agent reflections, updated *periodically* (every $K$ steps) by an LLM. The LLM receives the current Prospect State $s_i^t$, the current accounting variables $e_i^t$, a short history of the most recent $K$ raw events, and the previous Semantic State $\sigma_i^{t-K}$, then produces a new $\sigma_i^t$. This branch provides perception, reasoning, and human-like decision-making. The two branches operate in parallel: the Prospect State evolves continuously at low cost, preserving long-term path dependence; the Semantic State is refreshed on a longer timescale, injecting rich behavioral patterns without prohibitive token costs.

**Economic State Read.** After all agents have taken their actions, the environment updates the Economic State via a *market clearing function* $\Psi$:

$$(\mathbf{M}^{t+1}, \mathcal{S}^{t+1}, \mathbf{o}^t) = \Psi(\mathbf{a}^t, \mathbf{M}^t, \mathcal{S}^t), \tag{6}$$

where $\mathbf{M}^t = (\mathcal{M}_i^t)_{i\in\mathcal{I}}$, $\mathbf{a}^t = (a_i^t)_{i\in\mathcal{I}}$, and $\mathbf{o}^t = (o_i^t)_{i\in\mathcal{I}}$ contains the individual realized outcomes used by the Prospect-State Propagator. In our implementation, $\Psi$ aggregates intended consumption and labor supply. If total intended demand exceeds total production, goods are rationed proportionally; individual unmet demand $q_i^t$ is recorded as part of $o_i^t$ and can affect the next Prospect State update through the experience intensity $E_i^t$. Prices and wages adjust according to the imbalance (e.g., demand $>$ supply $\rightarrow$ price increase). The government collects progressive taxes and redistributes all revenue equally. Annually, the central bank sets the interest rate using a Taylor rule, and savings earn that interest. This transition closes the loop, feeding macro conditions back into agents' future states, thus cycling repeatedly.

**System Scaling.** PspMAS is designed to scale efficiently to large agent populations while preserving behavioral diversity.

*Computational Scalability*. The Prospect State is updated through lightweight and highly parallelizable propagation rules, avoiding expensive long-context reasoning during most simulation steps. Meanwhile, the Semantic State is refreshed only periodically rather than continuously, substantially reducing the number of costly LLM inference calls while keeping prompts compact over long simulation horizons. Since both state-update branches operate independently across agents, the overall system scales efficiently with the number of agents and simulation steps, enabling large-scale macroeconomic simulations with thousands of agents under practical computational budgets.

*Heterogeneity Quantification*. To measure whether diversity is maintained at scale, we construct three population-level matrices: a Trajectory Matrix, a Prospect Matrix, and a fused Heterogeneity Matrix. The Trajectory Matrix captures realized behavioral diversity from recent action sequences, while the Prospect Matrix captures latent path-dependent psychological diversity from Prospect-State Vectors. The fused matrix combines both views and is used to compute the effective number of distinct behavioral modes.

*Trajectory Matrix.* Let

$$\mathbf{z}_i^t = (p_i^{t-L+1}, l_i^{t-L+1}, \ldots, p_i^t, l_i^t) \tag{7}$$

be the action sequence of agent $i$ over the last $L$ steps ($L = 12$). After normalizing $\mathbf{z}_i^t$ over agents, we construct the Trajectory Matrix as an RBF similarity kernel:

$$H_{\text{traj},ij}^t = \exp\left(-\frac{\|\hat{\mathbf{z}}_i^t - \hat{\mathbf{z}}_j^t\|_2^2}{2\tau_{\text{traj}}^2}\right), \tag{8}$$

where $\hat{\mathbf{z}}_i^t$ is the normalized trajectory feature and $\tau_{\text{traj}}$ is a bandwidth parameter. This matrix is large when two agents have similar recent labor and consumption trajectories.

*Prospect Matrix.* The Prospect Matrix is constructed from the Prospect-State Vector

$$s_i^t = (\lambda_i^t, \gamma_i^t, R_i^t). \tag{9}$$

After normalizing $s_i^t$ over agents, we define

$$H_{\text{pros},ij}^t = \exp\left(-\frac{\|\hat{s}_i^t - \hat{s}_j^t\|_2^2}{2\tau_{\text{pros}}^2}\right), \tag{10}$$

where $\hat{s}_i^t$ is the normalized Prospect-State feature and $\tau_{\text{pros}}$ is a bandwidth parameter. This matrix captures whether two agents have similar path-dependent psychological states, even when their recent actions appear similar.

*Fused Heterogeneity Matrix.* We combine the two matrices with a balance coefficient $\alpha \in [0, 1]$:

$$H_{\text{fuse}}^t = \alpha H_{\text{traj}}^t + (1-\alpha) H_{\text{pros}}^t. \tag{11}$$

Because both $H_{\text{traj}}^t$ and $H_{\text{pros}}^t$ are RBF kernels, the fused matrix is positive semidefinite. Therefore, its spectrum can be used as a stable diversity measure.

Let $\mu_1^t, \mu_2^t, \ldots, \mu_N^t$ be the eigenvalues of $H_{\text{fuse}}^t$, and normalize them as

$$\tilde{\mu}_k^t = \frac{\mu_k^t}{\sum_{j=1}^N \mu_j^t}. \tag{12}$$

We define the spectral entropy as

$$\mathcal{E}^t = -\sum_{k=1}^N \tilde{\mu}_k^t \log(\tilde{\mu}_k^t + \epsilon), \tag{13}$$

and the effective diversity as

$$\mathcal{D}_{\text{eff}}^t = \exp(\mathcal{E}^t). \tag{14}$$

The normalized diversity is

$$D_{\text{norm}}^t = \frac{\mathcal{D}_{\text{eff}}^t}{N}. \tag{15}$$

When agents collapse into highly similar behaviors and similar Prospect States, $H_{\text{fuse}}^t$ approaches a low-rank matrix and $\mathcal{D}_{\text{eff}}^t$ becomes correspondingly small. When agents remain behaviorally and psychologically diverse, the effective diversity remains high. These metrics are computed periodically and used to assess heterogeneity preservation over long-horizon simulation.

**Long-Horizon Simulation Loop.** The simulation proceeds over a finite horizon $T$. The full simulation trajectory is

$$\tau = \left(\mathbf{M}^0, \mathcal{S}^0, \mathbf{a}^0, \mathbf{M}^1, \mathcal{S}^1, \ldots, \mathbf{M}^T, \mathcal{S}^T\right). \tag{16}$$

Algorithm 1 summarizes the entire simulation loop. At each time step, for each agent we optionally update the Semantic State every $K$ steps, then call the LLM to generate actions, execute market clearing, and update Prospect States with the Prospect-State Propagator. Heterogeneity is quantified by the spectral entropy $\mathcal{E}^t$ and effective diversity $\mathcal{D}_{\text{eff}}^t$ of the fused Heterogeneity Matrix $H_{\text{fuse}}^t$, which combines the Trajectory Matrix and the Prospect Matrix.

**Algorithm 1** Prospect-State Propagation Simulation Loop

**Input:** $N$ agents, horizon $T$, update interval $K$, initial macro state $\mathcal{S}^0$
Initialize micro states $\{\mathcal{M}_i^0\}_{i=1}^N$ with profiles (basic, economic, psychological)
**for** $t = 0$ **to** $T-1$ **do**
  **for** each agent $i$ **in parallel do**
    **if** $t \bmod K = 0$ **then**
      Update Semantic State $\sigma_i^t \leftarrow \mathrm{LLM}(s_i^t,\ e_i^t,\ \mathrm{history}_{t-K+1:t},\ \sigma_i^{t-K})$
    **end if**
    Generate economic action $a_i^t \sim \pi_\theta(\cdot \mid s_i^t,\ \sigma_i^t,\ e_i^t,\ \mathcal{S}^t)$
  **end for**
  Execute market clearing $\Psi$: compute $\mathcal{S}^{t+1}$ and record outcomes $\{o_i^t\}_{i=1}^N$
  **for** each agent $i$ **in parallel do**
    Update Prospect State $s_i^{t+1} = \Pi_{\mathcal{P}}(\Phi(s_i^t,\ o_i^t,\ \mathcal{S}^{t+1}))$
  **end for**
  Construct $H_{\text{traj}}^t$, $H_{\text{pros}}^t$, and $H_{\text{fuse}}^t$; compute $\mathcal{E}^t$, $\mathcal{D}_{\text{eff}}^t$, and $\mathcal{D}_{\text{norm}}^t$
**end for**

### 2.3 Implementation Details

We test a series of LLMs including Qwen3-8B (Bai et al., 2025), Qwen3-32B (Bai et al., 2025), Llama3.1-8B (Dubey et al., 2024), Llama3.1-70B (Dubey et al., 2024) for both decision-making and semantic state updates. The temperature is set to 0.7 for decision calls and 0.2 for summarization to balance diversity and coherence. Based on preliminary experiments, the update interval $K$ is set to 5. The Prospect-State Propagator coefficients are $\alpha_r = 0.1$, $\beta_r = 0.05$, $\eta_\lambda = 0.001$, $\eta_\gamma = 0.005$, $\zeta = 0.1$, $\sigma_\lambda = 0.02$. The threshold $E_{\text{thr}} = 10$. Unless otherwise specified, each Prospect-State update is projected onto the predefined valid domain $\mathcal{P}$ after Eq. 5. For heterogeneity evaluation, we set the trajectory window to $L = 12$ and use $\alpha = 0.5$ in the fused Heterogeneity Matrix unless otherwise specified. The bandwidths $\tau_{\text{traj}}$ and $\tau_{\text{pros}}$ are selected by the median heuristic over pairwise distances at each evaluation step. Further details are shown in the Appendix.

## 3 Experiments

Our experiments answer the following research questions:

**RQ1:** How does PspMAS compare with traditional LLM-based agents on economic simulation fidelity (*i.e.*, simulation realism)?

**RQ2:** Does the proposed Prospect-State Propagation preserve agent heterogeneity at scale?

### 3.1 Experimental Setting

*Simulation Environment.* We simulate a closed economy where each agent makes labor supply and consumption decisions at each step. The macro state includes GDP, inflation, and unemployment.

*Baseline Methods.* We compare three agent architectures:

- Plain Summary (PS): Agent histories are periodically compressed into LLM-generated summaries ($K = 5$). The summary replaces raw event history in the context window (Park et al., 2023; Qian et al., 2024).
- Periodic Reflection (PR): Agents maintain a memory stream and generate high-level reflections ("insights") at regular intervals ($K = 5$). Reflections are stored alongside recent events (Shinn et al., 2023; Wang et al., 2024a).
- PspMAS: Agents maintain a compact Prospect-State Vector $s_i^t \in \mathbb{R}^3$ updated via deterministic propagator at every step; Semantic State refreshed every $K = 5$ steps.

*Fairness of Heterogeneity Evaluation.* Since PS and PR do not maintain an internal Prospect State, we compute the Prospect Matrix for all methods from the same logged simulation outcomes. Specifically, for each method we run an evaluation-side observer using the same update rule as Eq. 4. For PS and PR, this observer is used only for post-hoc analysis and never exposed to agents during decision making. This allows the fused diversity metric to compare all methods using the same trajectory logs and the same transformation, without assuming baselines possess an internal state.

*Common Settings.* Agent counts $N \in \{20, 50, 100, 200, 500\}$. Each configuration runs with 3 independent random seeds. We report mean values with error bars in lighter gray.

### 3.2 RQ1: Comparison on Economic Simulation Fidelity

We first compare PspMAS against traditional LLM-based agents on standard economic simulation metrics. Following prior work (Chen and Chen, 2026),

Table 1: Comparison on economic simulation metrics. Higher is better ↑.

| Method | VR (%) ↑ | HR (%) (d/m) ↑ |
|---|---|---|
| PS | 75.8 ±1.2 | 35.6 ±1.5 / 54.7 ±1.3 |
| PR | 78.2 ±1.1 | 38.9 ±1.4 / 59.3 ±1.2 |
| PspMAS | **83.4** ±0.9 | **44.2** ±1.2 / **67.2** ±1.1 |

Table 2: Normalized fused diversity $D_{\text{norm}}$ (%) under varying agent count $N$. Higher is better.

| $N$ | Summary | Reflection | **PspMAS** |
|---|---|---|---|
| 20 | 68.2 ±1.2 | 71.3 ±1.1 | **73.8** ±1.0 |
| 50 | 40.3 ±1.5 | 55.6 ±1.3 | **67.5** ±1.1 |
| 100 | 24.5 ±1.8 | 39.8 ±1.4 | **62.1** ±1.2 |
| 200 | 21.2 ±1.6 | 33.5 ±1.3 | **59.6** ±1.1 |
| 500 | 20.1 ±1.4 | 30.5 ±1.1 | **56.5** ±1.0 |

Table 3: Normalized fused diversity $D_{\text{norm}}$ (%) under varying LLM backbone ($N = 200$). Higher is better.

| Model | Reflection | **PspMAS** |
|---|---|---|
| Qwen 3 - 8B | 32.1 ±1.3 | **58.2** ±1.1 |
| Qwen 3 - 32B | 36.5 ±1.2 | **62.5** ±1.0 |
| Llama 3.1 - 8B | 31.5 ±1.4 | **57.5** ±1.2 |
| Llama 3.1 - 70B | 37.2 ±1.1 | **63.8** ±0.9 |

Table 4: Normalized fused diversity $D_{\text{norm}}$ (%) under varying simulation horizon ($N = 200$). Higher is better.

| Horizon | Summary | Reflection | **PspMAS** |
|---|---|---|---|
| 50 | 28.5 ±1.6 | 42.5 ±1.3 | **61.2** ±1.1 |
| 100 | 23.1 ±1.5 | 36.2 ±1.2 | **60.1** ±1.0 |
| 150 | 20.5 ±1.4 | 32.8 ±1.2 | **58.9** ±1.0 |
| 200 | 18.5 ±1.3 | 30.2 ±1.1 | **57.8** ±0.9 |

we evaluate Volatility Realism (VR, similarity between simulated macro variables and real-world data) as well as daily and monthly Turning-Point Hit Rates (HR-d and HR-m).

Table 1 presents the results. PS achieves moderate performance but suffers from heterogeneity loss over long horizons. PR improves upon PS by maintaining richer memory structures, yet still exhibits gradual diversity decay. PspMAS outperforms both baselines across all metrics, demonstrating that preserving long-term psychological traces is crucial for capturing business cycle dynamics.

### 3.3 RQ2: Heterogeneity Preservation at Scale

To quantify whether PspMAS preserves agent heterogeneity as the system scales, we compute the normalized diversity $D_{\text{norm}} = \mathcal{D}_{\text{eff}}/N$ from the fused Heterogeneity Matrix $H^t_{\text{fuse}}$ defined in Section 2.2. This matrix combines the Trajectory Matrix, which measures recent action diversity, and the Prospect Matrix, which measures path-dependent psychological diversity. Higher $D_{\text{norm}}$ indicates more behaviorally distinct agents.

**Varying agent count.** Table 2 shows $D_{\text{norm}}$ as a function of $N$. PS saturates at $N \approx 200$, confirming that repeated semantic compression homogenizes both recent action trajectories and prospect-state traces regardless of population size. PR shows sublinear scaling. In contrast, PspMAS maintains high fused diversity even at $N = 500$, demonstrating stronger heterogeneity scaling.

**Varying model size.** Table 3 examines the impact of LLM backbone on heterogeneity preservation ($N = 200$). Larger models consistently achieve higher $D_{\text{norm}}$ for both PR and PspMAS. Notably, PspMAS with smaller models (e.g., Qwen3-8B) outperforms PR with larger models (e.g., Llama3.1-70B), suggesting that the Prospect-State mechanism helps preserve diversity beyond what is provided by model scale alone.

**Varying simulation horizon.** Table 4 tracks $D_{\text{norm}}$ over increasing simulation lengths ($N = 200$). PS degrades rapidly as repeated compression accumulates information loss. PR shows more graceful degradation but still suffers from cumulative drift. PspMAS maintains remarkably stable fused diversity, indicating that the Prospect-State Propagator preserves both behavioral trajectories and path-dependent prospect-state differences over long horizons.

### 3.4 Summary of Findings

Our experimental results provide clear answers to the research questions:

**RQ1:** Economic simulation fidelity: PspMAS achieves superior performance, with 83.4% VR and daily/monthly hit rates of 44.2%/67.2%, outperforming both PS (75.8%/35.6%/54.7%) and PR (78.2%/38.9%/59.3%).

**RQ2:** Heterogeneity preservation at scale: PspMAS maintains $D_{\text{norm}} > 0.56$ even at $N = 500$, while PS saturates at $N \approx 200$ ($D_{\text{norm}} \approx 0.21$). Since

$D_{\text{norm}}$ is computed from the fused Heterogeneity Matrix combining the Trajectory Matrix and the Prospect Matrix, this result indicates that PspMAS preserves both realized behavioral diversity and latent prospect-state diversity. The advantage persists across different model sizes and various simulation horizons.

## 4 Related Work

**Macroeconomic Simulation** Traditional macroeconomic models such as DSGE (Smets and Wouters, 2007; Clarida et al., 1999; Christiano et al., 2005) and VAR (Sims, 1980; Negro et al., 2015) primarily analyze economic fluctuations and policy transmission in a top-down manner through mathematical modeling of the relationship between "representative agents" and macroeconomic variables. Agent-based Modeling adopts a "bottom-up" paradigm for social simulation: by simulating the behaviors and interactions of micro-level individuals within a specific environment, it reproduces macroeconomic phenomena, thereby overcoming the limitation of the "representative agent" assumption inherent in traditional models (Acemoglu et al., 2012; Poledna et al., 2023; Caiani et al., 2016; Axtell, 2001; Geanakoplos et al., 2012; Dawid and Delli Gatti, 2018). In recent years, LLM-driven agents have endowed ABM with scenario-based interaction capabilities, allowing for the simulation of more complex economic behaviors and enabling more sophisticated economic simulations (Yang et al., 2026; Li et al., 2024b; Chen et al., 2024b; Li et al., 2025a; Hagendorff, 2025; Argyle et al., 2025; Jia et al., 2024; Widler et al., 2026). Despite these advances, existing approaches largely rely on periodic semantic summarization or reflection to compress long economic histories, a practice that risks attenuating fine-grained behavioral trajectories. This compression, though token-efficient, tends to homogenize agents over time by progressively smoothing away their early fluctuations.

**LLM-driven Agentic System** AI Agent is an autonomous system that perceives its environment and takes actions to achieve goals. These actions are not hard-coded but rather emerge from situated interactions, continuously evolving through feedback from interactions with the environment and other agents. This paradigm has achieved notable success in multi-turn conversation systems (Zheng et al., 2023; Wu et al., 2024a), code generation (Jimenez et al., 2024; Yang et al., 2024), and general-purpose assistant tasks (Wang et al., 2024b; Chen et al., 2024b). Furthermore, the same paradigm has advanced social reasoning (Gandhi et al., 2023; Sap et al., 2019; Mankowitz et al., 2023; Zhang et al., 2024), policy optimization (Silver et al., 2017, 2016), and behavioral simulation (Chen et al., 2024a; Park et al., 2023; Jia et al., 2024; Li et al., 2024a). This paradigm excels at leveraging LLMs for flexible, context-aware reasoning, enabling agents to adapt to novel situations without hand-crafted rules. The feedback-driven, interactive design also allows agents to continuously refine their strategies, producing emergent behaviors that static models cannot replicate. As individual agent capabilities improve, LLM-driven Multi-agent Systems have come into sharper focus, and recent research has begun to further investigate the collective behavior and emergent dynamics arising from multi-agent interactions. Recent work demonstrates that LLM-based multi-agent systems can spontaneously form social norms, collaborative structures, and complex group behaviors through ongoing interaction (Li et al., 2025b; Schneider et al., 2025; Wu et al., 2024b; Riedl, 2025). Moreover, recent efforts have leveraged these systems for macroeconomic forecasting (Jin et al., 2024; Ansari et al., 2024; Rasul et al., 2023; Garza et al., 2023), revealing intricate emergent macro-level phenomena.

## 5 Conclusion

We present PspMAS, the first macroeconomic simulation MAS that incorporates prospect theory into economic state updating. To achieve this, we first decouple each agent's micro state into two complementary components. We then leverage the Prospect State to continuously inject heterogeneity into the system, while using the Semantic State to generate human-like actions. By combining the two, PspMAS not only significantly reduces LLM inference token consumption, making effective scaling possible, but also achieves strong simulation performance. More broadly, our work provides new insights into addressing the challenges of LLM-driven MAS through classic theories from behavioral economics. Future work could extend PspMAS to other complex socio-economic domains such as financial market regulation, climate policy negotiation, and organizational behavior simulation, where bounded rationality and heterogeneous decision-making play critical roles.

## Limitations

This study has several limitations. First, the current simulations are still conducted within a simplified closed economy environment, and thus may fail to capture open economy effects, institutional constraints, or real-world policy frictions. Second, decisions based on large language models may inherit biases from their underlying models; future research should examine robustness across different model families, prompting strategies, and calibration settings. Third, during the training process, large language models absorb cultural, political, and ideological biases from the internet, academic literature, and policy documents. It is currently impossible to completely disentangle these deep-seated value orientations, posing a serious challenge for applications seeking "value-neutral" policy insights in practice.

## References


Daron Acemoglu, Vasco Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, and 1 others. 2024. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*.

Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*.

Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2025. Position: Llm social simulations are a promising research method. *Proceedings of the International Conference on Machine Learning (ICML)*.

Robert Axtell. 2001. Zipf distribution of us firm sizes. *Science*, 293(5536):1818–1820.

Jinze Bai, Shuai Bai, Yunfei Chu, and 1 others. 2025. Qwen3: A next-generation large language model. *arXiv preprint arXiv:2503.12345*.

Marcel Binz, Elif Akata, Matthias Bethge, Franziska Brändle, Fred Callaway, Julian Coda-Forno, Peter Dayan, Can Demircan, Maria K Eckstein, Noémi Éltető, and 1 others. 2025. A foundation model to predict and capture human cognition. *Nature*, 644(8078):1002–1009.

Olivier Blanchard. 2017. Do dsge models have a future? *Peterson Institute for International Economics Policy Brief*, (PB17-11).

Alessandro Caiani, Antoine Godin, Eugenio Caverzasi, Mauro Gallegati, Stephen Kinsella, and Joseph E Stiglitz. 2016. Agent based-stock flow consistent macroeconomics: Towards a benchmark model. *Journal of Economic Dynamics and Control*, 69:375–408.

Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, and 1 others. 2024a. From persona to personalization: A survey on role-playing language agents. *arXiv preprint arXiv:2404.18231*.

Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, and 1 others. 2024b. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *International Conference on Learning Representations*.

Zhimei Chen and Mu Chen. 2026. Empowering economic simulation through situation-aware llm-driven generative system. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 20007–20011. IEEE.

Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans. 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

Richard Clarida, Jordi Gali, and Mark Gertler. 1999. The science of monetary policy: a new keynesian perspective. *Journal of Economic Literature*, 37:1661–1707.

Dhairya Dalal, Marco Valentino, Andre Freitas, and Paul Buitelaar. 2024. Inference to the best explanation in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Herbert Dawid and Domenico Delli Gatti. 2018. Agent-based macroeconomics. In *Handbook of Computational Economics*, volume 4, pages 63–156. Elsevier.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

J. Doyne Farmer and Duncan Foley. 2009. The economy needs agent-based modelling. *Nature*, 460:685–686.

Yeqi Feng, Yucheng Lu, Hongyu Su, Yixin Tao, and Tianxing He. 2025. Simcity: Multi-agent urban development simulation with rich interactions. *arXiv preprint arXiv:2510.01297*.

Xavier Gabaix. 2011. The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772.

Kanishk Gandhi, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah Goodman. 2023. Understanding social reasoning in language models with language models. *Advances in Neural Information Processing Systems*, 36:13518–13529.

Azul Garza, Cristian Challu, and Max Mergenthaler-Canseco. 2023. Timegpt-1. *arXiv preprint arXiv:2310.03589*.

John Geanakoplos, Robert Axtell, Doyne J Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M Palmer, and Chun-Yi Yang. 2012. Getting at systemic risk via an agent-based model of the housing market. *American Economic Review*, 102(3):53–58.

Thilo Hagendorff. 2025. Do large language models exhibit spontaneous rational deception? In *Findings of the Association for Computational Linguistics: ACL 2025*.

John J. Horton. 2023. Large language models as simulated economic agents: What can we learn from homo silicus? *Journal of Economic Behavior & Organization*. Forthcoming.

Feiran Jia, Ziyu Ye, Shiyang Lai, Kai Shu, Jindong Gu, Adel Bibi, Ziniu Hu, David Jurgens, James Evans, Philip H Torr, and 1 others. 2024. Can large language model agents simulate human trust behavior? *Advances in neural information processing systems*, 37:15674–15729.

Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2024. Swe-bench: Can language models resolve real-world github issues? In *International Conference on Learning Representations*.

Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, and 1 others. 2024. Time-llm: Time series forecasting by reprogramming large language models. In *International conference on learning representations*, volume 2024, pages 23857–23880.

Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific.

Alan P. Kirman. 1992. Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2):117–136.

So Kuroki, Yingtao Tian, Kou Misaki, Takashi Ikegami, Takuya Akiba, and Yujin Tang. 2025. Reimagining agent-based modeling with large language model agents via shachi. *arXiv preprint arXiv:2509.21862*.

Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for mind exploration of large language model society. In *Advances in Neural Information Processing Systems*.

Junjie Li, Yang Liu, Weiqing Liu, Shikai Fang, Lewen Wang, Chang Xu, and Jiang Bian. 2025a. Mars: A financial market simulation engine powered by generative foundation model. In *International Conference on Learning Representations (ICLR)*.

Junkai Li, Yunghwei Lai, Weitao Li, Jingyi Ren, Meng Zhang, Xinhui Kang, Siyu Wang, Peng Li, Ya-Qin Zhang, Weizhi Ma, and 1 others. 2024a. Agent hospital: A simulacrum of hospital with evolvable medical agents. *arXiv preprint arXiv:2405.02957*.

Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. 2024b. Econagent: large language model-empowered agents for simulating macroeconomic activities. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Zhengyang Li, Sawyer Campos, and Nana Wang. 2025b. Language-driven coordination and learning in multi-agent simulation environments. *arXiv preprint arXiv:2506.04251*.

Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, and 1 others. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*.

Daniel J Mankowitz, Andrea Michi, Anton Zhernov, Marco Gelmi, Marco Selvi, Cosmin Paduraru, Edouard Leurent, Shariq Iqbal, Jean-Baptiste Lespiau, Alex Ahern, and 1 others. 2023. Faster sorting algorithms discovered using deep reinforcement learning. *Nature*, 618(7964):257–263.

Marco Del Negro, Marc P. Giannoni, and Frank Schorfheide. 2015. Inflation in the great recession and new keynesian models. *American Economic Journal: Macroeconomics*, 7:168–196.

OpenAI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*.

Sebastian Poledna, Michael Gregor Miess, Cars Hommes, and Katrin Rabitsch. 2023. Economic forecasting with an agent-based model. *European Economic Review*, 151:104306.

Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, and 1 others. 2024. Chatdev: Communicative agents for software development. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 15174–15186.

Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Hena Ghonia, Rishika Bhagwatkar, Arian Khorasani, Mohammad Javad Darvishi Bayazi, George Adamopoulos, Roland Riachi, Nadhir Hassen, and 1 others. 2023. Lag-llama: Towards foundation models for probabilistic time series forecasting. *arXiv preprint arXiv:2310.08278*.

Christoph Riedl. 2025. Emergent coordination in multi-agent language models. *arXiv preprint arXiv:2510.05174*.

Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Social iqa: Commonsense reasoning about social interactions. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 4463–4473.

Philipp J Schneider, Lin Tian, and Marian-Andrei Rizoiu. 2025. Learning to make friends: Coaching llm agents toward emergent social ties. *arXiv preprint arXiv:2510.19299*.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, and 1 others. 2016. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, and 1 others. 2017. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359.

Christopher A. Sims. 1980. Macroeconomics and reality. *Econometrica*, 48(1):1–48.

Frank Smets and Rafael Wouters. 2007. Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review*, 97(3):586–606.

Leigh Tesfatsion and Kenneth L. Judd. 2006. *Handbook of computational economics: agent-based computational economics*, volume 2. Elsevier.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2024a. Voyager: An open-ended embodied agent with large language models. *Nature Communications*, 15.

Noah Wang, Zy Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, and 1 others. 2024b. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*.

Zihao Wang, Shaofei Cai, Guanzhou Chen, Anji Liu, Xiaojian Shawn Ma, and Yitao Liang. 2023. Describe, explain, plan and select: interactive planning with llms enables open-world multi-task agents. *Advances in Neural Information Processing Systems*, 36:34153–34189.

Valeria Widler, Barbara Kamińska, André CR Martins, and Ivan Puga-Gonzalez. 2026. An investigation into the causal mechanism of political opinion dynamics: a model of hierarchical coarse-graining with community-bounded social influence. In *Computational Social Science of Social Cohesion and Polarization*, pages 225–256. Springer.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, and 1 others. 2024a. Autogen: Enabling next-gen llm applications via multi-agent conversations. In *First conference on language modeling*.

Zengqing Wu, Run Peng, Shuyuan Zheng, Qianying Liu, Xu Han, Brian I Kwon, Makoto Onizuka, Shaojie Tang, and Chuan Xiao. 2024b. Shall we team up: Exploring spontaneous cooperation of competing llm agents. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 5163–5186.

John Yang, Carlos E Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37:50528–50652.

Yuzhe Yang, Yifei Zhang, Minghao Wu, Kaidi Zhang, Yunmiao Zhang, Honghai Yu, Yan Hu, and Benyou Wang. 2026. Twinmarket: A scalable behavioral and social simulation for financial markets. *Advances in Neural Information Processing Systems*, 38:63469–63519.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*.

Yadong Zhang, Shaoguang Mao, Tao Ge, Xun Wang, Adrian de Wynter, Yan Xia, Wenshan Wu, Ting Song, Man Lan, and Furu Wei. 2024. Llm as a mastermind: A survey of strategic reasoning with large language models. *arXiv preprint arXiv:2404.01230*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.

Yilun Zhu, Joel Ruben Antony Moniz, Shruti Bhargava, Jiarui Lu, Dhivya Piraviperumal, Site Li, Yuan Zhang, Hong Yu, and Bo-Hsiang Tseng. 2024. Can large language models understand context? In *Findings of the Association for Computational Linguistics: EACL 2024*.

Henry Peng Zou, Wei-Chieh Huang, Yaozu Wu, Yankai Chen, Chunyu Miao, Hoang Nguyen, Yue Zhou, Weizhi Zhang, Liancheng Fang, Langzhou He, and 1 others. 2025. A survey on large language model based human-agent systems. *arXiv preprint arXiv:2505.00753*.

# A Appendix

This appendix provides supplementary material for PspMAS. It is organized into five parts:

- **A. Broader Impact** discusses the positive societal implications and potential risks of our framework.
- **B. Experimental Settings** defines the evaluation metrics (VR, HR) and lists simulation environment parameters.
- **C. Token Efficiency and Heterogeneity Preservation** compares token consumption and normalized diversity under matched budgets using Qwen3-32B.
- **D. Ablation Study** isolates the effect of the prospect-state propagator via a rule-based PSP-only baseline.
- **E. Discussion** explains the theoretical grounding of our design, including Markovian dynamics and fast–slow complementarity.

## A. Broader Impact

Our work introduces PspMAS, a scalable multi-agent simulation framework that preserves agent heterogeneity via prospect-state propagation. On the positive side, PspMAS enables more realistic macroeconomic modeling by maintaining long-term behavioral diversity under fixed token budgets, which can improve the fidelity of policy simulations, stress testing, and economic forecasting. The decoupled architecture—compact Markovian prospect states plus periodic LLM-based semantic updates—provides a general blueprint for scaling LLM-driven agent systems beyond economics, potentially benefiting domains such as urban planning, supply chain management, and epidemiological modeling. Moreover, the spectral diversity metric we introduce offers a principled way to quantify and monitor agent heterogeneity, which can be adopted as a diagnostic tool in any multi-agent simulation.

On the negative side, several risks warrant consideration. First, although PspMAS reduces reliance on long-text compression, the semantic branch still depends on LLMs, which may inherit or amplify biases present in their training data (e.g., stereotypes about income groups or consumption behavior). If the initial agent profiles or economic scenarios are not carefully designed, the simulation could produce systematically skewed macro outcomes that reinforce real-world inequalities. Second, the prospect-state propagator, while lightweight and parallelizable, is based on parameterized rules (e.g., loss-aversion update, reference-point adaptation). Inappropriate parameter settings or unrecognized feedback loops could lead to emergent behaviors that are economically unrealistic or unstable, potentially misleading policymakers if the system is deployed for real-world guidance. Third, as with any agent-based economic simulator, there is a risk of dual use: malicious actors might adapt the framework to model manipulative market strategies or to stress-test financial systems in ways that create systemic vulnerabilities. However, these risks are not unique to PspMAS; they are inherent to any high-fidelity economic simulation tool. We emphasize that our contribution is methodological, and we strongly advocate for transparent reporting of model parameters, initial distributions, and LLM prompts, as well as routine diversity audits using the proposed spectral metrics, to mitigate these concerns without compromising the core innovation.

## B. Experimental Settings

**Evaluation Metrics.** Following recent economic simulation literature (Li et al., 2024b; Yang et al., 2026), we evaluate simulation realism using two complementary metrics.

**Volatility Realism (VR).** VR measures the similarity between the statistical properties of simulated macroeconomic variables and those observed in real-world data. Specifically, we compute the Pearson correlation coefficient between each simulated variable (e.g., GDP, inflation, unemployment) and its real-world counterpart. VR is then defined as the average correlation across all macro variables. Higher VR indicates more realistic volatility

Table 5: Key simulation environment parameters.

| Parameter | Value |
|---|---|
| Number of agents $N$ | $\{20, 50, 100, 200, 500\}$ |
| Simulation horizon $T$ | 120 months |
| Semantic update interval $K$ | 5 steps |
| Initial wealth | $\text{Uniform}(10, 50)$ kUSD |
| Initial wage | 3.2 kUSD/month |
| Price adjustment speed | $\kappa = 0.1$ |
| Taylor-rule coefficients | $\phi_\pi = 1.5,\ \phi_y = 0.5$ |
| Progressive tax rate | 15% base tax + 5% surcharge on top-income agents |

patterns.

**Turning-Point Hit Rate (HR).** HR quantifies the model's ability to correctly identify the timing of economic cycle turning points (peaks and troughs). Following common practice, we evaluate HR at three temporal resolutions: daily (HR-d), monthly (HR-m), and yearly (HR-y). For each resolution, a turning point is considered correctly predicted if the simulated turning date falls within a tolerance window of the actual turning date (e.g., ±5 days for daily, ±1 month for monthly, ±2 months for yearly). The final HR is the ratio of correctly predicted turning points to the total number of real turning points, expressed as a percentage.

**Simulation Environment Parameters.** Table 5 lists the key parameters of the closed-economy simulation environment used in all experiments. The market clearing mechanism follows EconAgent (Li et al., 2024b) with proportional rationing and Taylor-rule based interest rate adjustment.

## C. Token Efficiency and Heterogeneity Preservation

To verify that the diversity gain of PspMAS is not an artifact of increased token usage, we compare total token consumption and normalized diversity $D_{\text{norm}}$ under identical settings ($N = 500, T = 120, K = 5$, Qwen3-32B). Table 6 reports the results. All three methods consume nearly identical token budgets (within $\pm 2\%$), confirming a fair comparison. Despite using almost the same number of tokens, PspMAS achieves a $D_{\text{norm}}$ of 56.5%, which is 2.3× higher than PS (20.1%) and 1.9× higher than PR (30.5%). This directly supports the claim in the abstract.

Figure 3 further visualizes that normalized diversity remains stable throughout the simulation horizon, and the final diversity is achieved under nearly identical cumulative token budgets. The dashed vertical line indicates the total token consumption at the end of the horizon.

Table 6: Token consumption and normalized diversity under matched token budgets (Qwen3-32B).

| Method | Tokens (M) | $D_{\text{norm}}$ (%) ↑ |
|---|---|---|
| Plain Summary (PS) | $162.3 \pm 4.1$ | $20.1 \pm 1.4$ |
| Periodic Reflection (PR) | $164.7 \pm 3.9$ | $30.5 \pm 1.1$ |
| PspMAS | $161.8 \pm 4.3$ | $\mathbf{56.5 \pm 1.0}$ |

All results are reported under $N = 500$, $T = 120$, $K = 5$, using Qwen3-32B. Token counts are in millions.

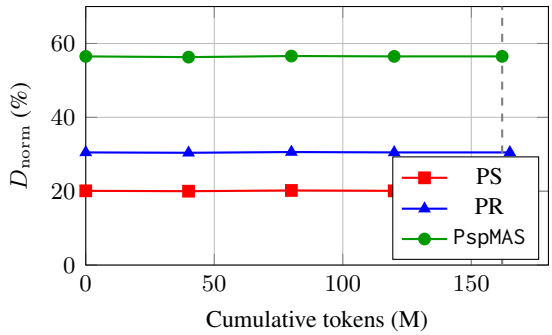


Figure 3: Normalized diversity as a function of cumulative token consumption (Qwen3-32B). The dashed vertical line marks the total tokens used at $T = 120$.

## D. Ablation Study

To disentangle the contribution of the Prospect-State branch from that of the LLM-based decision maker, we introduce an additional baseline: **PSP-only**. In this variant, each agent maintains the same Prospect-State Vector $s_i^t = (\lambda_i^t, \gamma_i^t, R_i^t)$ and updates it via the same Markovian propagator (Eq. 5). However, instead of calling an LLM to generate actions, the agent uses a simple rule-based policy that directly maps its current Prospect State and macro state to labor supply $l_i^t \in \{0, 1\}$ and consumption propensity $p_i^t \in [0, 1]$. The rule-based policy is defined as:

- **Labor supply:** $l_i^t = 1$ if unemployment rate $u_{\text{emp}}^t < 0.1$ (jobs available) or if $(W_i^t - R_i^t) < 0$ (loss domain, risk-averse seeking income); otherwise $l_i^t = 0$.
- **Consumption propensity:** $p_i^t = \sigma\big(0.5 - 0.2\lambda_i^t + 0.3\gamma_i^t + 0.1(W_i^t - R_i^t)\big)$, where $\sigma(x) = 1/(1 + e^{-x})$ is the logistic function.

The PSP-only baseline consumes no LLM tokens after initial profile generation. We evaluate its normalized diversity $D_{\text{norm}}$ under the same settings

as Table 3 in the main paper ($N = 200$, $T = 120$). Table 7 presents the results.

Table 7: Ablation study: contribution of the Prospect-State branch vs. LLM-based decision making. $D_{\text{norm}}$ (%) at $N = 200$, $T = 120$. Higher is better.

| Method | $D_{\text{norm}}$ (%) ↑ |
|---|---|
| PSP-only (no LLM) | $48.3 \pm 1.2$ |
| Plain Summary (PS) | $23.1 \pm 1.5$ |
| Periodic Reflection (PR) | $36.2 \pm 1.2$ |
| PspMAS (full) | $\mathbf{60.1 \pm 1.0}$ |

The PSP-only baseline achieves a $D_{\text{norm}}$ of $48.3\%$, which is substantially higher than PS ($23.1\%$) and PR ($36.2\%$), even without any LLM inference. This confirms that the Prospect-State mechanism alone preserves a significant amount of agent heterogeneity. Adding the LLM-based decision branch (full PspMAS) further boosts diversity to $60.1\%$, demonstrating that the LLM contributes additional behavioral richness beyond what the lightweight propagator can capture. Notably, PSP-only outperforms both PS and PR, highlighting the severe heterogeneity loss caused by iterative semantic compression in those baselines. These results support our claim that the Prospect-State branch is the primary driver of diversity preservation, while the LLM branch provides complementary fine-grained expressiveness.

## E. Discussion: Why Prospect Theory Preserves Heterogeneity

The design of PspMAS is grounded in a behavioral-economic insight: *the most path-dependent, heterogeneity-driving components of economic decision-making are exactly those that can be updated Markovianly with low-dimensional psychological variables.* Prospect theory provides precisely such a structure.

**Markovian dynamics do not contradict prospect theory.** A common concern is that a Markovian prospect state discards historical information. However, dynamic extensions of prospect theory (e.g., reference-point adaptation) are themselves Markovian: the next reference point depends only on the current reference point and recent outcomes. Loss aversion and probability weighting, being slowly evolving traits, can also be captured by first-order updates. Hence, the Markovian nature of the prospect state is a faithful implementation, not a violation.

**How each parameter preserves heterogeneity.**

- *Reference point $R_i^t$:* Different histories lead to different reference points, causing the same macro shock to be perceived as a "loss" for some agents and a "reduced gain" for others, directly preventing behavioral homogenization.
- *Loss aversion $\lambda_i^t$:* Agents that experience crises or financial stress increase $\lambda_i^t$, amplifying the psychological weight of negative outcomes and creating persistent differences under identical conditions.
- *Probability weighting $\gamma_i^t$:* High $\gamma_i^t$ leads to overweighting of small probabilities, distorting labor supply decisions and generating heterogeneity that no semantic summary can capture.

**Fast–slow complementarity.** The prospect state updates every step ($\mathcal{O}(1)$ FLOPs), capturing fast, path-dependent psychological adjustments. The semantic state updates every $K$ steps via an LLM, providing slow, reflective reasoning. This aligns with dual-process theory: System 1 (fast, automatic) is modeled by the prospect propagator, while System 2 (slow, analytical) is modeled by the LLM. Their integration is key to scalable, heterogeneity-preserving economic simulation.

In summary, prospect theory works because its core variables are low-dimensional, Markovian, and directly tied to the mechanisms that generate persistent individual differences. Experiments (§3) and the ablation study in the Appendix (§A) confirm that the prospect branch is the primary driver of heterogeneity preservation.