ECONAI: DYNAMIC PERSONA EVOLUTION AND MEMORY-AWARE AGENTS IN EVOLVING ECONOMIC ENVIRONMENTS

Yijin Chen[1,*], **Ning Lyu**[2,*], **Shengning Lang**[3,*],
Hao Yan[3,*], Xiaotong Ding[4,*,†], Xiaotong Zhu[1,†], Zhiguo Tao[5]

[1]University of Southern California, Los Angeles, CA, USA
[2]Carnegie Mellon University, Pittsburgh, PA, USA
[3]Stevens Institute of Technology, Hoboken, NJ, USA
[4]Boston University, Boston, MA, USA
[5]University of Nottingham, Nottingham, UK

[*]Equal contribution
[†]Corresponding authors

ABSTRACT

The integration of large language models (LLMs) in economic simulations has significantly enhanced agent-based modeling, yet existing frameworks struggle to capture the interplay between short-term optimization and long-term strategic planning. Conventional approaches rely on static data-driven predictions, failing to incorporate adaptive behaviors influenced by economic sentiment, market volatility, and individual goals. To address these limitations, we introduce a novel EconAI framework, incorporating economic sentiment indexing (ESI), memory weighting, and dynamic decision-making mechanisms. By quantifying economic belief, adjusting historical data influence, and linking work-consumption behaviors, EconAI achieves a more human-like decision process, where agents adapt their actions based on both market signals and long-term objectives. It is the first LLM-powered simulation system that can simulate the macro/microeconomic environment and interactions in a unified framework. Empirical evaluations show that EconAI improves stability in economic responses, better replicates real-world employment-consumption cycles, and enhances overall decision robustness. This advancement marks a crucial step towards more realistic, adaptive economic agent simulations.

1 INTRODUCTION

Artificial intelligence has revolutionized economic research through advanced data processing, with Agent-based modeling (ABM) emerging as a powerful framework for simulating economic systems from the bottom up without predetermined equilibria (Farmer & Foley, 2009), though early rule-based approaches (Tesfatsion & Judd, 2006; Brock & Hommes, 1998) and recent learning-based techniques (Trott et al., 2021; Zheng et al., 2022; Mi et al., 2023) face challenges in tailoring individual agent decision-making due to extensive calibration requirements and computational costs (Windrum et al., 2007; Mi et al., 2023). While LLM integration has advanced agent-based modeling (Zhao et al., 2023b), existing frameworks fail to capture the interplay between short-term optimization and long-term strategic planning, particularly in integrating event memory and economic beliefs—leading to accumulating simulation errors (Yue et al., 2024). To address this, we introduce EconAI, a novel framework enhancing long-term LLM decision-making capabilities by categorizing agents into households (microeconomic) and firms (macroeconomic) while integrating government and financial institutions, incorporating a knowledge memory module with long-term banks storing vector representations of high-level event summaries and short-term banks retaining contextual information, plus a persona extraction module that models evolving preferences and economic beliefs stored in a long-term persona bank for personalized decision-making. Our contributions are three-fold: (1) EconAI is the first LLM-powered simulation integrating macroeconomic dynamics and micro-level interactions in a unified framework; (2) it incorporates adaptive mechanisms including the Economic Sentiment Index (ESI), memory weighting, and dynamic decision-making models; (3) extensive simulations demonstrate EconAI significantly outperforms existing models in simulating economic indicators while verifying different economic laws.

2 Related Work

2.1 Simulation in Macroeconomics

Agent-Based Modeling (ABM) has emerged as a particularly effective tool in macroeconomic research, surpassing traditional empirical statistical models (Hendry & Richard, 1982; Phelps, 1967; Kydland & Prescott, 1982) and DSGE models (Christiano et al., 2005) by enabling deep exploration of complex, non-linear behaviors through independent agents interacting via predefined rules without assuming economic equilibrium—making it valuable for simulating policy interventions. However, models utilizing fixed rules (Tesfatsion & Judd, 2006; Brock & Hommes, 1998) or neural network architectures (Trott et al., 2021; Zheng et al., 2022; Mi et al., 2023) often oversimplify agent behavior or rely heavily on large datasets, reducing effectiveness in capturing economic complexities. We introduce EconAI, equipped with cognitive and strategic reasoning capabilities to adaptively simulate both macroeconomic and microeconomic processes.

2.2 LLM-empowered Agents

Large Language Models have demonstrated human-comparable performance, serving as foundations for advanced simulation agents (Wang et al., 2023; Xi et al., 2023) with autonomous adaptability (AutoGPT Team, 2022; Yoheinakajima, 2023), strategic planning capabilities, and human-agent interactions (Park et al., 2023; Gilbert & Troitzsch, 2005), spanning social sciences (Park et al., 2022; Kovač et al., 2023; Gao et al., 2023; Jinxin et al., 2023) and natural sciences (Boiko et al., 2023; Bran et al., 2023). In economics, LLM-powered agents model individual behavior (Horton, 2023; Chen et al., 2023b), support interactive planning (Guo, 2023; Akata et al., 2023), and simulate market systems (Zhao et al., 2023a; Chen et al., 2023a), yet existing studies (Li et al., 2024) predominantly focus on short-term optimal approaches, neglecting long-term learning aligned with human decision-making. Our research introduces simulation agents driven by long-term learning to bridge this gap.

3 EconAI Framework

In this section, we introduce EconAI, which simulates economic activity plans in the decision-making process. This plan comprises a series of abstract actions to be carried out across various scenarios with a long-term-based framework. According to economists (Falk et al., 2018), long-term decision-making rather than short-term counterparts is the base for a robust economic system. To this end, as illustrated in Figure 1, we develop a new approach to simulate the long-term decision-making process including event perception, dynamic personas extraction, and response generation for decision-making.

3.1 Event Perception

The core of simulating economic activities lies in managing the environment by focusing on four critical entities: households, firms, financial institutions, and government. Together, these entities create an economic system that includes both macroeconomic and microeconomic dimensions. Specifically, the simulation should represent key economic decisions, including labor supply and consumption, as fundamental drivers of government tax revenue and determinants of labor and consumer market dynamics (Gatti et al., 2011; Wolf et al., 2013; Dawid & Gatti, 2018), which influence government tax revenue (Zheng et al., 2022; Trott et al., 2021) and shape labor and consumer market dynamics (Lengnick, 2013; Deissenberg et al., 2008; Dawid et al., 2012). In response to these market conditions, banks adjust interest rates in line with inflationary or deflationary pressures (Wolf

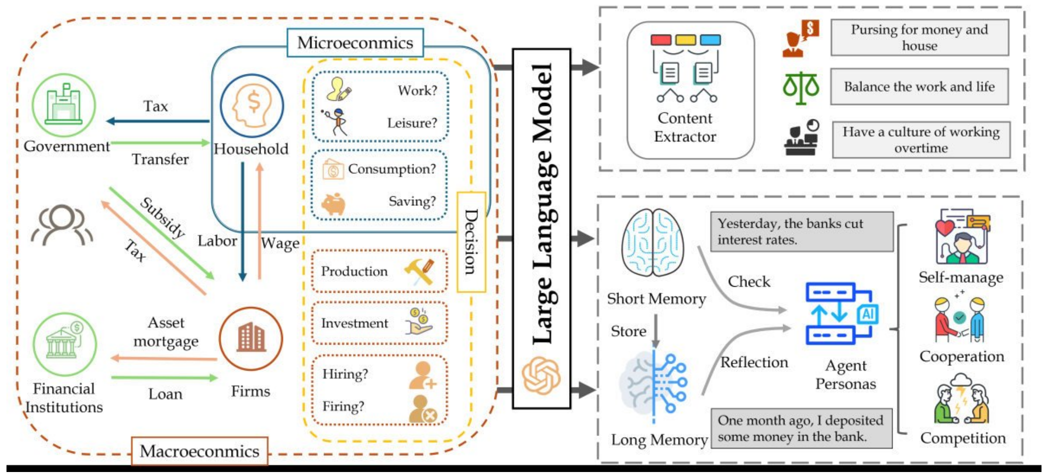

Figure 1: The illustration of the simulation for the microeconomic and macroeconomic environments (left) and our EconAI (right). The event module retains historical data in long-term memory and contextual information in short-term memory. The persona module dynamically extracts and updates agent personas, storing them in a dedicated personal bank.

et al., 2013; Dawid & Gatti, 2018). However, none of the previous work can achieve the goal of simulating both macroeconomics and microeconomics in a unified framework.

For an agent, such as a household, it has specified metadata such as the profession, specialty, skills, credentials, and experiences of the agent. The agent observes information from the environment, makes decisions, and conducts the appropriate action. In real-world economic activities, humans often make decisions based on heuristic rules and established customs derived from experience. Much like humans, the agent's brain serves as a central nucleus driven by an LLM, enabling the agent to exhibit sophisticated cognitive abilities critical for professional-grade performance, including memory, planning, and reasoning.

Memory Storage. The event memory module is designed to perceive historical events to generate coherent responses across interval time. This event memory module is segmented into two major sub-modules that focus separately on long-term and short-term memory. The memory module is designed to capture and encode events from previous interactions in economic activities. This process involves logging both the timestamps t and concise summaries o, which are transformed into representations and stored in a low-cost memory bank $M_L = \{\phi(t_j, o_j) | j \in \{1, 2, \ldots, l\}\}$. Here, $\phi(\cdot)$ denotes the text encoder (such as MiniLM (Wang et al., 2020)), while l represents the size of the memory bank. The encoded information is then made easily accessible via an embedding-based retrieval mechanism.

Event Summary. Unlike previous agent methods (Park et al., 2023) that depend solely on the zero-shot capabilities of LLMs for event extraction and summarization, we incorporate instruction tuning (Wei et al., 2022) into our event summary module to enhance the quality of the summaries for economic activities. We remember the event in a structured format consisting of: (1) a task background introduction, (2) relevant economic activities to be comprehended, and (3) specific summarization instructions. These elements serve as input prompts and generate the output. The event summary module is fine-tuned using this data, effectively improving its summarization capabilities.

3.2 Economic Preference and Confidence

Memory information is objective, but human decision-making is often affected by emotions and confidence, such as being overly pessimistic during an economic recession and overly optimistic during an overheated economy. Equal weighting of short-term and long-term information cannot reflect people's different sensitivities to economic changes in different time frames. Then, we introduce the Economic Sentiment Index (ESI) to quantify the economic belief of agents and store this

ESI value in the memory module:

$$\text{ESI}_t = \lambda \cdot \text{ESI}_{t-1} + (1 - \lambda) \cdot \text{ESI}_{t,\text{LLM}} \tag{1}$$

where λ is the memory decay factor (0.8–0.95) and $\text{ESI}_{t,\text{LLM}}$ is the latest economic sentiment index generated by LLM.

In this way, the economic sentiment of each quarter can be quantified and stored for subsequent decision-making calculations. We can make the results of LLM reflection quantifiable, reducing computational overhead. Incorporating the dynamic changes in economic belief makes the decision-making of the intelligent agent not only data-dependent but also emotionally influenced. The impact of historical economic data will decay, avoiding excessive reactions by the intelligent agent due to short-term economic shocks.

Then, we use confidence to adjust the final decision. If confidence is low (e.g., below 0.5), we adjust the final decision in conjunction with the economic sentiment index:

$$p_w^i = \sigma(p_w^i - \theta \cdot \text{ESI}_t), \quad p_c^i = \sigma(p_c^i + \beta \cdot \text{ESI}_t) \tag{2}$$

If economic belief is low ($\text{ESI}_t < 0$), the agent reduces consumption and increases the willingness to work. If economic belief is high ($\text{ESI}_t > 0$), the agent increases consumption and reduces work.

3.3 Economic Environment Design

This section outlines the design and decision-making processes of intelligent economic agents, government-imposed taxation policies, production-consumption dynamics, and financial market mechanisms shaping macroeconomic behavior.

Agent Decision-Making Framework. The core decision-making process for each agent involves two key economic choices: work and consumption. Initially, these decisions are modeled as:

$$p_w^t, p_c^t \sim \text{LLM}(z_i, P, s_i, u, r, \text{ESI}_t) \tag{3}$$

where p_w^t is the labor market participation probability, and p_c^t is the consumption-to-income ratio. The variable z_i corresponds to the agent's individual income, while P reflects the prevailing market price. s_i accounts for the agent's current savings, u provides insight into labor market conditions, and r represents the bank interest rate. These inputs collectively allow the LLM to generate dynamic and context-aware decisions.

Production and Consumption Dynamics. Labor supply contributes to goods production following a Cobb-Douglas function:

$$G \leftarrow G + S = G + AK_t^{\beta_1} L_t^{\beta_2} e^{-\lambda K_t} \tag{4}$$

where A represents the overall productivity factor, K_t is the capital stock, and $L_t = \sum_{j=1}^{N} l_j \times 168$ denotes the total labor hours supplied. Capital evolves dynamically as:

$$K_t = (1 - \delta)K_{t-1} + I - \omega K_t \tag{5}$$

where δ represents depreciation, I is new investment, and ωK_t accounts for inefficiencies. The market imbalance is defined as $\hat{\phi} = \frac{D-G}{\max(D,G)}$, which drives price and wage adjustments:

$$w_i \leftarrow w_i(1 + \phi_i \cdot \kappa), \quad P \leftarrow P(1 + \phi_P \cdot \kappa) \tag{6}$$

Government Taxation and Redistribution. To enhance economic stability and equity, the government implements a progressive taxation system, where tax rates rise with income. Post-tax income incorporates redistributed tax revenues, promoting partial wealth reallocation and reducing extreme income inequality. In financial markets, agents' savings grow at an interest rate influenced by unemployment and inflation dynamics, following a modified Taylor rule.

3.4 Response Generation

When the agent engages in economic activity, it will recall the memory module, which combines event extractor: retrieved relevant memories m, the short-term memory context M_S, and the personas P_a representing the agent. These inputs are passed to a response generator, which produces an appropriate response r. To improve the agent's capability to generate coherent and contextually relevant responses for economic activities, for each example spanning a series of times (such as several months in the simulation), we dynamically simulate the economic activity's progression. As new economic events arise, the system applies event summarization, persona extraction, and topic-aware memory retrieval to collect relevant context.

4 EXPERIMENTS

In this section, we conduct experiments to study the ability of EconAI, aiming to answer the following research questions (RQ). **RQ1**: How does EconAI perform in simulations compared to conventional models? **RQ2**: What role do the key components of EconAI play in influencing the simulation outcomes? **RQ3**: Is the decision-making process within EconAI interpretable, and do simulations effectively capture the impact of external interventions?

4.1 EXPERIMENTAL SETUP

We use LEN (Lengnick, 2013) and CATS (Gatti et al., 2011) as baselines because they partially capture the macroeconomic phenomena mentioned earlier within their simulation frameworks and have carefully designed decision-making rules for work and consumption; given the significance of agent heterogeneity in macroeconomic simulations, we introduce an additional baseline, Composite, by combining these two models, and include a learning-based approach, AI-Economist (Zheng et al., 2022) (AI-Eco), which uses reinforcement learning to optimize agent utility under rational decision-making assumptions, and compare with EconAgent (Li et al., 2024), which incorporates a perception module for modeling the macroeconomic environment and simulates heterogeneous agents.
Definition of Economic Indicators. The annual nominal GDP is calculated as the total of $S \times P$ over the course of a year. To compute real GDP, we set the first year of the simulation as the base year. The wage inflation rate follows a similar formula to price inflation. For households, disposable income refers to the total income left after taxes and essential expenditures. The savings rate is the proportion of disposable income that is saved rather than used for consumption. For firms, profit margin is defined as the ratio of net profit to total revenue.

Simulation Setup. To fully leverage the contextual understanding and extensive knowledge of LLM, each agent in the simulation is assigned specific real-world attributes, such as name, age, and occupation. The agents' age distribution follows the demographic pattern of the U.S. population between 18 and 60 years old, based on 2018 data (US Census Bureau, 2024). For the economic variables, the simulation scales the parameters of the Pareto distribution for hourly wages, ensuring that the resulting monthly wages align with actual U.S. economic data and tax brackets from 2018 (Zheng et al., 2022). The simulation was implemented using Python, utilizing GPT-4o-mini via the OpenAI API.

4.2 MACRO-LEVEL ANALYSIS (RQ1)

Economic Indicators. Figure 2 illustrates the annual fluctuations in inflation rate and nominal GDP. Notably, due to the unreasonably high unemployment rate (approximately 46%) and nominal GDP values in AI-Eco, these results are omitted from the report. Both rule-based and RL-driven baseline models exhibit irregular indicators and significant volatility. In contrast, decision-making in EconAI has been shown to produce more stable and realistic macroeconomic outcomes across various dimensions, even without extensive calibration. This indicates that EconAI's decision-making process is more coherent and closely mirrors real-world human behavior, thereby creating a more natural balance between supply and demand in the consumption market.

Economic Regularity. The Phillips Curve (Phelps, 1967), a widely recognized regularity in macroeconomic theory, illustrates the inverse relationship between the annual unemployment rate and wage inflation. As shown in Figure 3, the modeling performance of EconAgent and other models is relatively weak. In contrast, EconAI's decision-making process accurately captures the expected negative correlation between these two factors (Pearson correlation coefficient of -0.522, $p < 0.01$). Notably, the rule-based baseline exhibited an incorrect positive correlation on the Phillips Curve. Furthermore, EconAI is also capable of modeling Okun's law (Okun, 1963), which describes the empirical relationship between unemployment and economic output. We attribute EconAI's success

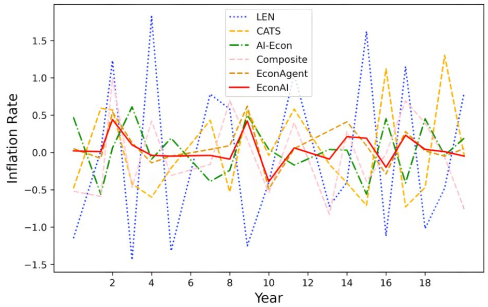

(a) Inflation rate.

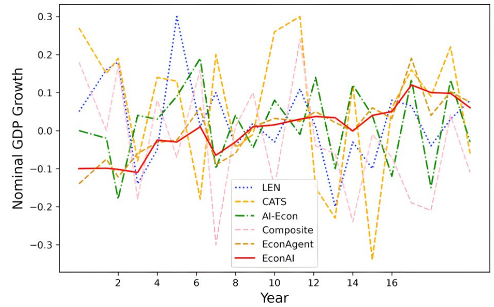

(b) Nominal GDP Growth.

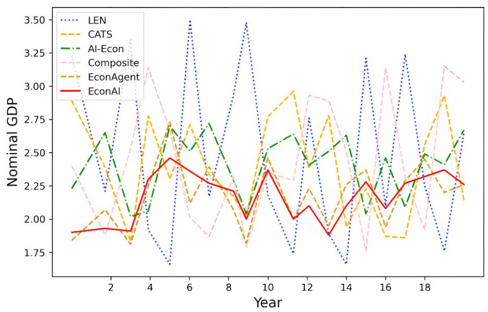

(c) Normal GDP.

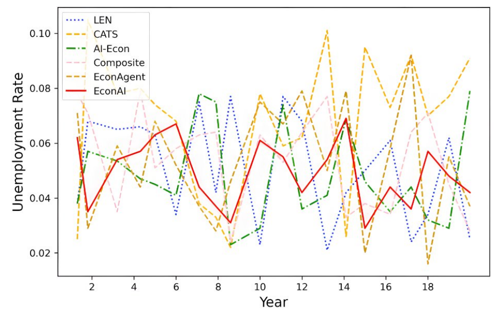

(d) Unemployment rate.

Figure 2: Annual variations of macroeconomic indicators, where the simulation based on EconAI shows more stable and numerically plausible indicators.

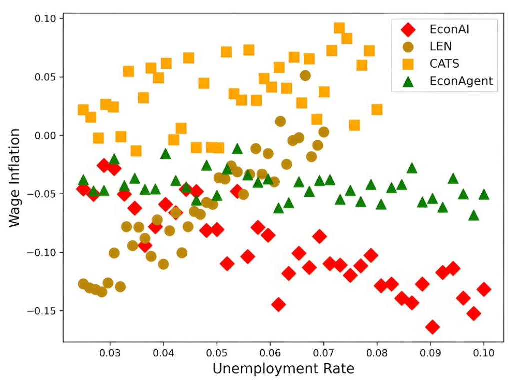

(a) Phillips curve .

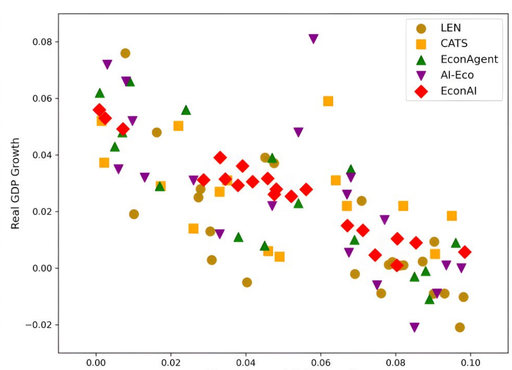

(b) Okun curve.

Figure 3: Economic regularity study. EconAI accurately captures the Phillips curve (negative correlation, $r = -0.522$, $p < 0.01$) and Okun's law.

to its correct understanding that consumption should be reduced during periods of high unemployment, thereby capturing these key economic dynamics effectively.

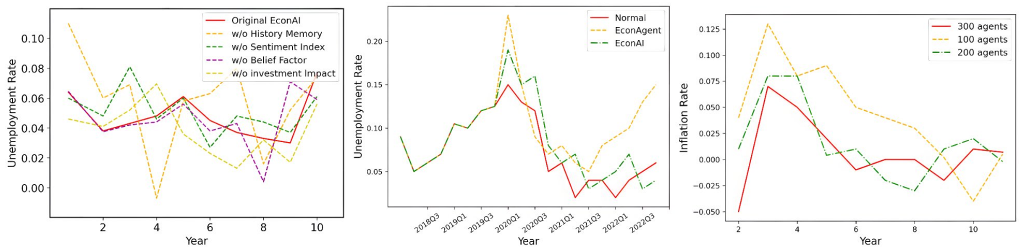

(a) Ablation study.

(b) COVID-19 brings the surge of simulated unemployment rates.

(c) Inflation rate under different number of agents.

Figure 5: Additional experimental studies including ablation study, external intervention (COVID-19), and sensitivity analysis with varying agent counts.

4.3 Micro-level Analysis (RQ1)

In the economic environment modeled by EconAI, simulations of firms reveal classic market strategies and economic phenomena such as competition and cooperation, as illustrated in Figure 4: one scenario where one firm's revenue stagnates or declines significantly suggests destructive competition occurring when firms engage in aggressive price wars, excessive resource allocation to outcompete rivals, or strategic behaviors that undermine industry stability, resembling a zero-sum or negative-sum game in game theory where excessively aggressive competition without value creation causes the industry to suffer from reduced profitability and instability; in contrast, the second scenario where both firms experience revenue growth reflects principles of healthy competition, where firms compete in ways that drive innovation, improve efficiency, and expand market demand so both can thrive, aligning with the game-theoretic concept of positive-sum competition where firms engage in strategic differentiation rather than direct conflict.

4.4 Ablation Study (RQ2)

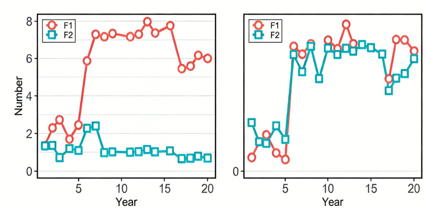

(a) Competition. (b) Cooperation.

Figure 4: Micro-level analysis showing competition and cooperation dynamics between firms.

We conduct an ablation study by independently removing key economic factors, with results over a 10-year span shown in Figure 5. We analyze five versions: the original model (EconAI), No History Memory, No Sentiment Index, No Belief Factor, and No Investment Impact. The absence of historical memory causes sharp unemployment rate fluctuations, indicating past economic trends crucially stabilize employment levels. Removing the sentiment index increases volatility, highlighting public sentiment's influence on labor market dynamics. The absence of the belief factor introduces inconsistencies, reinforcing its role in guiding employment expectations. Removing investment impact leads to lower overall unemployment variability, suggesting investment plays a stabilizing role. These findings emphasize the importance of incorporating these components.

4.5 External Intervention (RQ3)

We investigate the impact of external shocks on agent decision-making, a key focus in agent-based economic modeling (Dawid & Gatti, 2018). Using the COVID-19 pandemic as a case study in Figure 5(b), we integrate pandemic-related scenarios into EconAI's prompts to simulate its economic effects. From March 2020 onward, simulations include a directive modeling the pandemic's economic repercussions, such as the sharp rise in unemployment during early 2020 (Organization for Economic Co-operation and Development, 1970). While the model's quantitative results may not fully align with real-world data, it qualitatively captures human-like decision-making and macroeconomic trends, including sustained high unemployment in the absence of government intervention.

Table 1: Quantitative comparison of economic simulation metrics across different methods. Lower values indicate better performance for volatility metrics (↓), while higher values are better for correlation metrics (↑).

Method	Inflation Vol.↓	Unemp. Vol.↓	Phillips↑	Okun↑
LEN	0.082	0.045	0.12	0.08
CATS	0.075	0.038	0.18	0.15
Composite	0.068	0.042	0.21	0.19
AI-Eco	0.156	0.089	-0.05	-0.12
EconAgent	0.054	0.031	0.35	0.28
EconAI	**0.041**	**0.024**	**-0.52**	**-0.48**

Agent reflections during this period highlight adaptive strategies, such as skill development and income diversification, under prolonged economic uncertainty.

Sensitivity and Robustness. We increase the number of agents to 300 and rerun the simulation. As illustrated in Figure 5(c), the inflation rates remain consistently stable and realistic, closely resembling the results obtained with 100 agents. This trend is also observed across other economic indicators, indicating that the simulation outcomes are not significantly affected by changes in the number of agents.

4.6 Interpretability Analysis

To validate EconAI's decision-making process, we analyze agent reflections during economic scenarios. During COVID-19 simulation, an agent reflected: 'When facing long-term economic uncertainty, individuals should focus on building career stability and resilience through strategic transitions and skill development. Without long-term government intervention, individuals should adjust expectations, diversify income sources, and make cautious investments to maintain stability.' This demonstrates that EconAI agents exhibit human-like reasoning patterns incorporating both short-term adaptations and long-term strategic considerations, show awareness of external economic conditions, and respond with realistic behavioral adjustments, validating the effectiveness of our persona extraction and memory modules.

4.7 Comparison with Additional Metrics

Table 1 presents a comprehensive comparison of EconAI against baseline methods across multiple economic metrics. The results demonstrate consistent improvements across all indicators. As shown in Table 1, EconAI achieves the lowest volatility in both inflation and unemployment rates, indicating more stable economic simulations. Importantly, EconAI accurately captures the negative correlations described by the Phillips Curve (Pearson $r = -0.52$) and Okun's Law (Pearson $r = -0.48$), while other methods either show weak correlations or incorrect positive correlations. Table 1 shows EconAI achieves the lowest volatility in inflation and unemployment rates, accurately captures negative correlations in the Phillips Curve (Pearson $r = -0.52$) and Okun's Law (Pearson $r = -0.48$) while baselines show weak or incorrect correlations, and demonstrates significant policy analysis implications as policymakers can test interventions before implementation and the COVID-19 simulation shows EconAI models external shocks and their economic propagation for crisis response planning.

Comparison with Traditional ABM. While traditional agent-based models rely on hand-crafted rules that require extensive domain expertise, EconAI leverages the general reasoning capabilities of LLMs to generate more nuanced and contextually appropriate decisions. This reduces the need for extensive manual calibration while improving the realism of agent behaviors. Furthermore, EconAI's modular architecture allows easy integration of new economic factors or agent types without redesigning the entire system.

5 CONCLUSION

We propose EconAI consisting of three independently tunable modules: event perception, long-term action based on memory, and response generation. By integrating event memory with economic confidence through dynamic persona modeling, our approach addresses the limitations of traditional rule-based and machine-learning agents, which often fail to adapt to evolving market complexities. Empirical evaluations demonstrate that EconAI effectively captures key macroeconomic and microeconomic phenomena, including inflation, employment, and market shocks, producing results that closely align with established economic theories.

REFERENCES

Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. Playing repeated games with large language models. *arXiv preprint arXiv:2305.16867*, 2023.

AutoGPT Team. Autogpt: the heart of the open-source agent ecosystem. https://github.com/Significant-Gravitas/AutoGPT, 2022.

Daniil A Boiko, Robert MacKnight, and Gabe Gomes. Emergent autonomous scientific research capabilities of large language models. *arXiv preprint arXiv:2304.05332*, 2023.

Andres M Bran, Sam Cox, Andrew D White, and Philippe Schwaller. Chemcrow: Augmenting large-language models with chemistry tools. *arXiv preprint arXiv:2304.05376*, 2023.

William A Brock and Cars H Hommes. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control*, 22(8-9):1235–1274, 1998.

Jiangjie Chen, Siyu Yuan, Rong Ye, Bodhisattwa Prasad Majumder, and Kyle Richardson. Put your money where your mouth is: Evaluating strategic planning and execution of llm agents in an auction arena. *arXiv preprint arXiv:2310.05746*, 2023a.

Yiting Chen, Tracy Xiao Liu, You Shan, and Songfa Zhong. The emergence of economic rationality of gpt. *arXiv preprint arXiv:2305.12763*, 2023b.

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

Herbert Dawid and Domenico Delli Gatti. Agent-based macroeconomics. In *Handbook of computational economics*, volume 4, pp. 63–156. Elsevier, 2018.

Herbert Dawid, Simon Gemkow, Philipp Harting, Sander Van der Hoog, and Michael Neugart. The eurace@unibi model: An agent-based macroeconomic model for economic policy analysis. 2012.

Christophe Deissenberg, Sander Van Der Hoog, and Herbert Dawid. Eurace: A massively parallel agent-based model of the european economy. *Applied mathematics and computation*, 204(2): 541–552, 2008.

Armin Falk, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. Global evidence on economic preferences. *The Quarterly Journal of Economics*, 133(4):1645–1692, 2018.

J Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460(7256): 685–686, 2009.

Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. S3: Social-network simulation system with large language model-empowered agents. *arXiv preprint arXiv:2307.14984*, 2023.

Domenico Delli Gatti, Saul Desiderio, Edoardo Gaffeo, Pasquale Cirillo, and Mauro Gallegati. *Macroeconomics from the Bottom-up*, volume 1. Springer Science & Business Media, 2011.

Nigel Gilbert and Klaus Troitzsch. *Simulation for the social scientist*. McGraw-Hill Education (UK), 2005.

Fulin Guo. Gpt agents in game theory experiments. *arXiv preprint arXiv:2305.05516*, 2023.

David F Hendry and Jean-Francois Richard. On the formulation of empirical models in dynamic econometrics. *Journal of Econometrics*, 20(1):3–33, 1982.

John J Horton. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research, 2023.

Shi Jinxin, Zhao Jiabao, Wang Yilei, Wu Xingjiao, Li Jiawen, and He Liang. Cgmi: Configurable general multi-agent interaction framework. *arXiv preprint arXiv:2308.12503*, 2023.

Grgur Kovač, Rémy Portelas, Peter Ford Dominey, and Pierre-Yves Oudeyer. The socialai school: Insights from developmental psychology towards artificial socio-cultural agents. *arXiv preprint arXiv:2307.07871*, 2023.

Finn E Kydland and Edward C Prescott. Time to build and aggregate fluctuations. *Econometrica: Journal of the Econometric Society*, pp. 1345–1370, 1982.

Matthias Lengnick. Agent-based macroeconomics: A baseline model. *Journal of Economic Behavior & Organization*, 86:102–120, 2013.

Nian Li, Chen Gao, Mingyu Li, Yong Li, and Qingmin Liao. Econagent: Large language model-empowered agents for simulating macroeconomic activities. In *ACL*, pp. 15523–15536, 2024.

Qirui Mi, Siyu Xia, Yan Song, Haifeng Zhang, Shenghao Zhu, and Jun Wang. Taxai: A dynamic economic simulator and benchmark for multi-agent reinforcement learning. *arXiv preprint arXiv:2309.16307*, 2023.

Arthur M Okun. Potential gnp: Its measurement and significance. *Cowles Foundation for Research in Economics at Yale University*, 1963.

Organization for Economic Co-operation and Development. Infra-annual labor statistics: Unemployment rate total: From 15 to 64 years for united states, 1970.

Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–18, 2022.

Joon Sung Park, Joseph C O'Brien, Carrie J Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–22, 2023.

Edmund S Phelps. Phillips curves, expectations of inflation and optimal unemployment over time. *Economica*, pp. 254–281, 1967.

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

Alexander Trott, Sunil Srinivasa, Douwe van der Wal, Sebastien Haneuse, and Stephan Zheng. Building a foundation for data-driven, interpretable, and robust policy design using the ai economist. *arXiv preprint arXiv:2108.02904*, 2021.

US Census Bureau. Access demographic, economic and population data from the u.s. census bureau, 2024.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents. *arXiv preprint arXiv:2308.11432*, 2023.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *NeurIPS*, 2020.

Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *ICLR*, 2022.

Paul Windrum, Giorgio Fagiolo, and Alessio Moneta. Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10(2):8, 2007.

Sarah Wolf, Steffen Fürst, Antoine Mandel, Wiebke Lass, Daniel Lincke, Federico Pablo-Marti, and Carlo Jaeger. A multi-agent model of several economic regions. *Environmental modelling & software*, 44:25–43, 2013.

Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*, 2023.

Yoheinakajima. Babyagi. `https://github.com/yoheinakajima/babyagi`, 2023.

Zihao Yue, Liang Zhang, and Qin Jin. Less is more: Mitigating multimodal hallucination from an eos decision perspective. *ACL*, pp. 11766–11781, 2024.

Qinlin Zhao, Jindong Wang, Yixuan Zhang, Yiqiao Jin, Kaijie Zhu, Hao Chen, and Xing Xie. Competeai: Understanding the competition behaviors in large language model-based agents. *arXiv preprint arXiv:2310.17512*, 2023a.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023b.

Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C Parkes, and Richard Socher. The ai economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science advances*, 8(18):eabk2607, 2022.

A Appendix

A.1 Government Taxation and Redistribution

To maintain economic stability and equity, the government implements a progressive taxation system. The tax burden for an agent earning z_i is computed as:

$$T(z_i) = \sum_{k=1}^{B} \tau_k \left((b_{k+1} - b_k)\mathbf{1}[z_i > b_{k+1}] + (z_i - b_k)\mathbf{1}[b_k < z_i \leq b_{k+1}]\right) \tag{7}$$

where b_k represents tax bracket thresholds, τ_k denotes the corresponding tax rates, and $\mathbf{1}[\cdot]$ is an indicator function. After taxation, post-tax income is updated as:

$$\tilde{z}_i = z_i - T(z_i) + \frac{T_{\text{total}}}{N}, \quad T_{\text{total}} = \sum_{j=1}^{N} T(z_j), \quad s_i \leftarrow s_i + \tilde{z}_i \tag{8}$$

A.2 Financial Market Mechanisms

Agents' savings evolve with an interest rate function: $s_i \leftarrow s_i \times (1 + r - \rho u)$, where unemployment u negatively impacts savings. The interest rate follows a modified Taylor rule:

$$r = \max(r^*, \pi_t + \gamma_\pi(\pi - \pi_t) + \gamma_u(u^* - u) - \nu\Delta G) \tag{9}$$

where r^* is the natural rate, π_t the target inflation, π the observed inflation, and u^* the natural unemployment rate. Inflation and unemployment dynamics are given by:

$$\pi = \frac{\tilde{P}_n - \tilde{P}_{n-1}}{\tilde{P}_{n-1}}, \quad u = \frac{\sum_{m=1}^{12}\sum_{j=1}^{N}(1 - l_j)}{12N} + \zeta\Delta P \tag{10}$$

where $\tilde{P}_n$ is the price index for year n, and $\zeta\Delta P$ accounts for inflation-induced unemployment effects.

A.3 SIMULATION PARAMETERS

We simulate $N = 100$ agents. The productivity is set as $A = 1$ for simplicity. The initial goods price is the average hourly wage across all agents. For labor and consumption dynamics, $\alpha_w = 0.05$ and $\alpha_P = 0.10$. For the financial market, $r^n = 0.01$, $\pi_t = 0.02$, $u^n = 0.04$, and $\alpha_\pi = \alpha_u = 0.5$. Note that our results and conclusions are not sensitive to these parameters.