ADAGE: A Generic Two-layer Framework for Adaptive Agent based Modelling

Extended Abstract

Benjamin Patrick Evans
JP Morgan AI Research
London, United Kingdom
benjamin.x.evans@jpmorgan.com

Sihan Zeng
JP Morgan AI Research
Palo Alto, USA
sihan.zeng@jpmchase.com

Sumitra Ganesh
JP Morgan AI Research
New York, USA
sumitra.ganesh@jpmorgan.com

Leo Ardon
JP Morgan AI Research
London, United Kingdom
leo.ardon@jpmorgan.com

ABSTRACT

Agent-based models (ABM) are valuable for modelling complex systems, however, they are often manually specified and lack behavioral and/or environmental adaptation. In this work, we develop a generic two-layer framework for ADaptive AGEnt based modelling (ADAGE) for addressing this. ADAGE formalises the bi-level problem of agent and environment adaptation as a Stackelberg game, providing a consolidated framework for adaptive ABM. We demonstrate how ADAGE encapsulates several modelling tasks, such as policy design, calibration, scenario generation, and robust behavioural learning under one unified framework. We provide example simulations on various environments, showing the flexibility of ADAGE while addressing long-standing critiques of ABMs.

KEYWORDS

Agent-based model; Multi-agent reinforcement learning; Stackelberg games; Adaptation; Bi-level optimisation; Lucas critique

ACM Reference Format:
Benjamin Patrick Evans, Sihan Zeng, Sumitra Ganesh, and Leo Ardon. 2025. ADAGE: A Generic Two-layer Framework for Adaptive Agent based Modelling: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 4 pages.

1 INTRODUCTION

Agent-based models (ABMs) have shown promise for modelling complex systems where outcomes deviate from equilibrium [3, 4]. However, ABMs crucially rely on the behavioural rules of the agents, which are typically fixed and/or manually specified. This fixed behaviour opens ABMs to the famed Lucas Critique: *Given that the model consists of decision rules of agents, and that decision rules vary systematically with changes in the environment, any change in policy will invalidate the model* (paraphrased from [24]) raising concerns about using fixed behaviour rules [26, 35]. While prior efforts have tackled agent adaptation in macroeconomics [19, 30] and AI [2, 9, 22, 38], the interplay between agent behavior and environmental dynamics remains under-explored [44]. Sufficient adaptation requires a framework for co-evolving agents and their environment. The AI Economist [34, 44] introduced a two-layer approach to adapt agent behavior and macroeconomic conditions, excelling in tax policy design. However, this is tailored to policy design and lacks a general formulation for broader modeling tasks.

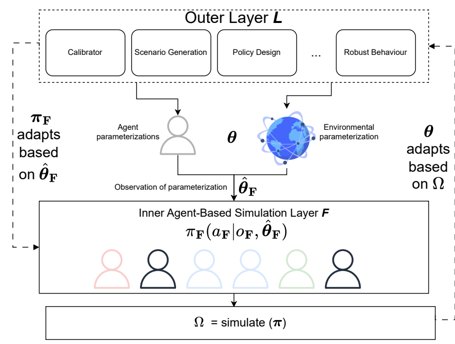

Figure 1: ADAGE: Two-layer framework.

In this work, we develop a generic two-layer framework for ADaptive AGEnt-based modeling (ADAGE), unifying diverse modeling tasks, such as model calibration, policy design, and scenario generation. ADAGE formalises agent and environmental co-adaptation as a Stackelberg game, where an outer layer updates the environmental parameterization $\boldsymbol{\theta}$, and an inner simulation layer conditions agent behavior on observations of $\boldsymbol{\theta}$. ADAGE addresses the Lucas critique by co-adapting agents and environments, while unifying diverse modeling tasks under a single versatile framework.

2 PROPOSED APPROACH

ADAGE (fig. 1) is represented as a Stackelberg game [7, 8, 18], a type of Partially Observable Markov Game, with $n + 1$ agents: a leader $L = 0$ as the outer layer and n followers $\mathbf{F} = \{1, \ldots, n\}$ in the inner simulation layer, operating in a parameterised environment [16, 41] representing the ABM. ADAGE's generality lies in L's flexibility to

Extended version available online: https://arxiv.org/abs/2501.09429.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

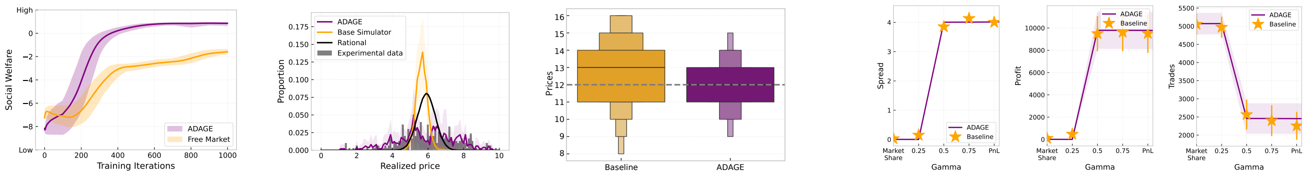

(a) Policy Design **(b) Calibration** **(c) Scenario Generation** **(d) Robust Behavioral Learning**

Figure 2: Experimental Results

serve multiple purposes depending on the task, while maintaining a unified representation and solution concept.

The game is defined by the tuple: (S, A, T, r, O, γ) where S is the state space, $A = (A_0, A_1, \ldots, A_n)$ the action spaces, $T : S \times A \longrightarrow S$ the transition function, $r : S \times A \longrightarrow \mathbb{R}^N$ the reward functions, $O = (O_0, O_1, \ldots, O_N)$ the observation spaces where agents have a (partial) observation of the state $o_i(s) : S \longrightarrow O_i$, and γ the discount rate. The leader acts first, optimizing a global objective, influencing followers via the parameterization $\boldsymbol{\theta}$, derived from L's policy π_L, parameterizing the ABM. Followers condition their behaviour $\pi_i(a \mid o_i, \hat{\boldsymbol{\theta}}_i)$ on observations of $\boldsymbol{\theta}$.At timestep t each active agent i takes an action $a_{i,t} \sim \pi_i(o_{i,t})$ based on their policy π_i and private observation $o_{i,t}$. The goal of agent i is to find a policy π_i to maximise their expected return: $R_i = \mathbb{E}[\sum_t \gamma^t r_{i,t}(s_t, a_{i,t}, a_{-i,t})]$ where $a_{-i,t}$ is the action of the other agents. A Stackelberg equilibrium (STE) is a solution $(\pi_L^*, \pi_{\mathbf{F}}^*)$ at which point no agent can improve their return while holding the behavior of others fixed: $\pi_L^* \in \arg\max_{\pi_L} [R_L | \pi_{\mathbf{F}} \sim \varepsilon(\pi_L)]$, $\pi_{\mathbf{F}} = \{\pi_i : i \in \mathbf{F}\}$ where ε represents a best-response oracle [18]. Generally, the task is to simultaneously optimise $\pi_L, \pi_{\mathbf{F}}$, as the oracle ε is unknown. Here, we focuses on approximating the STE and applying it to adaptive ABM, rather than convergence guarantees [15, 43]. To find an (approximate) STE, it suffices to solve [1] the following set of coupled non-linear equations: $\{\tilde{\nabla}_{\pi_i^*} R_i = 0, \forall i \in \mathbf{F}, \tilde{\nabla}_{\pi_L^*} R_L = 0\}$ ([1, 21, 40]).

3 EXPERIMENTS

To demonstrate ADAGEs flexibility, we provide illustrative (but non-exhaustive) examples spanning distinct modelling tasks and environments, showing how each is is encapsulated by ADAGE.

Policy Design, a key modelling use case [29, 31, 42, 44], is captured by ADAGE. To demonstrate this, we use a tax policy simulator based on TaxAI [25], with multiple households, firms, banks, and a central government (similar to [10, 44]). The government L selects taxation rates $\boldsymbol{\theta}$ to maximise social welfare of the households $\mathbf{F}$, who are maximising their individual reward subject to $\boldsymbol{\theta}$. ADAGE successfully maximises social welfare (fig. 2a), significantly improving upon the free market case, indicating the framework can perform policy design in complex economic systems.

Calibration is another crucial modelling task, matching simulators to real-world dynamics [5, 12, 37]. We demonstrate how this fits within ADAGE using the Cobweb market game [20] (CMG). CMG explores price fluctuations, where with human participants, we see larger fluctuations than implied by perfectly rational participants, so we must calibrate the level of bounded rationality of the agents [13] through $\boldsymbol{\theta}$. ADAGE successfully calibrates the simulator to real-world data (from [20]), capturing the overall distribution and producing the best fit (fig. 2b), demonstrating that calibration and bounded rationality are compatible with ADAGE.

Scenario Generation is another common modelling task [11, 28]. Despite being distinct from the previous tasks, we demonstrate the compatibility with ADAGE. We utilise a market entrance game [27] and explore the impact of Tobins tax [33] on controlling market volatility [6, 32], where the goal is to generate scenarios for stabilizing the market. ADAGE discovers Tobin tax settings (paramaterised by $\boldsymbol{\theta}$) that reduce volatility (fig. 2c), aligning with prior manual findings [6, 32], demonstrating scenario generation capabilities.

Robust Behaviour Learning is another modelling task, learning behaviour across preferences, e.g., meta-learning [17]. We demonstrate how ADAGE can be used for this in a financial environment (adapted from [36, 39]) with a market maker (MM) trading with multiple zero-intelligence liquidity takers [14, 23]. The MM can have different preferences ω in terms of maximising market share $\omega \longrightarrow 0$ or PnL $\omega \longrightarrow 1$, and the goal is to learn generic behaviour across $\boldsymbol{\omega} = \{0, 0.25, 0.5, 0.75, 1\}$. Despite having a single behavioural policy, ADAGE reproduces fixed ω results (fig. 2d) that previously required retraining for each preference, demonstrating robust behavioural learning with a single policy through sampling preferences via $\boldsymbol{\theta}$.

4 DISCUSSION AND CONCLUSION

We introduced ADAGE, a novel adaptive ABM framework in which agents and their environment co-evolve, with agents adjusting their behaviour in response to environmental changes and vice versa, encompassing diverse modelling tasks (e.g., policy design, calibration, scenario generation, and robust behavioural learning). Given ADAGEs generality, we hope this will become a valuable framework for developing adaptive agent-based simulations.

[1] we use alternating gradient descent, iteratively updating $\pi_{i,t} \forall i$ as an estimate of π_i^* based on the approximate gradient $\tilde{\nabla}_{\pi_{i,t}} R_i$ with $\pi_j, \forall j \neq i$ fixed to their latest iterate.

REFERENCES

[1] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. 2021. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research* 22, 98 (2021), 1–76.

[2] Leo Ardon, Jared Vann, Deepeka Garg, Thomas Spooner, and Sumitra Ganesh. 2023. Phantom - A RL-driven Multi-Agent Framework to Model Complex Systems. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (London, United Kingdom) *(AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2742–2744.

[3] W Brian Arthur. 2021. Foundations of complexity economics. *Nature Reviews Physics* 3, 2 (2021), 136–145.

[4] Robert L Axtell and J Doyne Farmer. 2022. Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature* (2022), 1–101.

[5] Marco Benedetti, Gennaro Catapano, Francesco De Sclavis, Marco Favorito, Aldo Glielmo, Davide Magnanimi, and Antonio Muci. 2022. Black-it: A ready-to-use and easy-to-extend calibration kit for agent-based models. *Journal of Open Source Software* 7, 79 (2022), 4622.

[6] Ginestra Bianconi, Tobias Galla, Matteo Marsili, and Paolo Pin. 2009. Effects of Tobin taxes in minority game markets. *Journal of Economic Behavior & Organization* 70, 1-2 (2009), 231–240.

[7] Gianluca Brero, Darshan Chakrabarti, Alon Eden, Matthias Gerstgrasser, Vincent Li, and David C Parkes. 2022. Learning stackelberg equilibria and applications to economic design games. (2022).

[8] Gianluca Brero, Alon Eden, Darshan Chakrabarti, Matthias Gerstgrasser, Vincent Li, and David C Parkes. 2022. Stackelberg POMDP: A Reinforcement Learning Approach for Economic Design. *arXiv preprint arXiv:2210.03852* (2022).

[9] M Ale Ebrahim Dehkordi, JM Lechner, Amineh Ghorbani, Igor Nikolic, EJL Chappin, and PM Herder. 2023. Using machine learning for agent specifications in agent-based models and simulations: A critical review and guidelines. *Journal of Artificial Societies and Social Simulation* 26, 1 (2023), 9.

[10] Kshama Dwarakanath, Svitlana Vyetrenko, Peyman Tavallali, and Tucker Balch. 2024. ABIDES-Economist: Agent-Based Simulation of Economic Systems with Learning Agents. *arXiv preprint arXiv:2402.09563* (2024).

[11] Joel Dyer, Arnau Quera-Bofarull, Nicholas Bishop, J. Doyne Farmer, Anisoara Calinescu, and Michael Wooldridge. 2024. Population Synthesis as Scenario Generation for Simulation-based Planning under Uncertainty. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) *(AAMAS '24)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 490–498.

[12] Joel Dyer, Arnau Quera-Bofarull, Ayush Chopra, J Doyne Farmer, Anisoara Calinescu, and Michael Wooldridge. 2023. Gradient-assisted calibration for financial agent-based models. In *Proceedings of the Fourth ACM International Conference on AI in Finance*. 288–296.

[13] Benjamin Patrick Evans and Sumitra Ganesh. 2024. Learning and Calibrating Heterogeneous Bounded Rational Market Behaviour with Multi-agent Reinforcement Learning. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) *(AAMAS '24)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 534–543.

[14] J Doyne Farmer, Paolo Patelli, and Ilija I Zovko. 2005. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences* 102, 6 (2005), 2254–2259.

[15] Tanner Fiez, Benjamin Chasnov, and Lillian Ratliff. 2020. Implicit Learning Dynamics in Stackelberg Games: Equilibria Characterization, Convergence Analysis, and Empirical Study. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 3133–3144.

[16] Carmel Fiscko, Brian Swenson, Soummya Kar, and Bruno Sinopoli. 2019. Control of parametric games. In *2019 18th European Control Conference (ECC)*. IEEE, 1036–1042.

[17] Matthias Gerstgrasser and David C Parkes. 2022. Meta-RL for Multi-Agent RL: Learning to Adapt to Evolving Agents. In *Sixth Workshop on Meta-Learning at the Conference on Neural Information Processing Systems*.

[18] Matthias Gerstgrasser and David C Parkes. 2023. Oracles & followers: Stackelberg equilibria in deep multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 11213–11236.

[19] Cars Hommes, Mario He, Sebastian Poledna, Melissa Siqueira, and Yang Zhang. 2024. CANVAS: A Canadian behavioral agent-based model for monetary policy. *Journal of Economic Dynamics and Control* (2024), 104986.

[20] Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk Van De Velden. 2007. Learning in cobweb experiments. *Macroeconomic Dynamics* 11, S1 (2007), 8–33.

[21] Mingyi Hong, Hoi-To Wai, Zhaoran Wang, and Zhuoran Yang. 2023. A two-timescale stochastic algorithm framework for bilevel optimization: Complexity analysis and application to actor-critic. *SIAM Journal on Optimization* 33, 1 (2023), 147–180.

[22] Franziska Klügl and Hildegunn Kyvik Nordås. 2023. Modelling Agent Decision Making in Agent-based Simulation-Analysis Using an Economic Technology Uptake Model. In *22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29–June 2, 2023*. International Foundation for Autonomous Agents and Multiagent Systems, 1903–1911.

[23] Dan Ladley. 2012. Zero intelligence in economics and finance. *The Knowledge Engineering Review* 27, 2 (2012), 273–286.

[24] Robert E Lucas Jr. 1976. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, Vol. 1. North-Holland, 19–46.

[25] Qirui Mi, Siyu Xia, Yan Song, Haifeng Zhang, Shenghao Zhu, and Jun Wang. 2024. TaxAI: A Dynamic Economic Simulator and Benchmark for Multi-agent Reinforcement Learning. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) *(AAMAS '24)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1390–1399.

[26] Mauro Napoletano. 2018. A short walk on the wild side: Agent-based models and their implications for macroeconomic analysis. *Revue de l'OFCE* 3 (2018), 257–281.

[27] Benjamin Patrick Evans and Mikhail Prokopenko. 2023. Bounded strategic reasoning explains crisis emergence in multi-agent market games. *Royal Society Open Science* 10, 2 (2023), 221164.

[28] Pisit Praiwattana and Abdennour El Rhalibi. 2016. Survey: Development and analysis of a games-based crisis scenario generation system. In *E-Learning and Games: 10th International Conference, Edutainment 2016, Hangzhou, China, April 14-16, 2016, Revised Selected Papers 10*. Springer, 85–100.

[29] Arnau Quera-Bofarull, Ayush Chopra, Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Anisoara Calinescu, Ramesh Raskar, and Michael Wooldridge. 2023. Don't Simulate Twice: One-Shot Sensitivity Analyses via Automatic Differentiation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (London, United Kingdom) *(AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1867–1876.

[30] Isabelle L Salle. 2015. Modeling expectations in agent-based models—An application to central bank's communication and monetary policy. *Economic Modelling* 46 (2015), 130–141.

[31] Tuomas Sandholm. 2003. Automated mechanism design: A new application area for search algorithms. In *International Conference on Principles and Practice of Constraint Programming*. Springer, 19–36.

[32] Dipyaman Sanyal. 2019. *Effect of Tobin Tax on Trading Decisions in an Experimental Minority Game*. Springer International Publishing, Cham, 71–79. https://doi.org/10.1007/978-3-030-11364-3_6

[33] James Tobin. 1978. A proposal for international monetary reform. *Eastern economic journal* 4, 3/4 (1978), 153–159.

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

[35] Arthur Turrell. 2016. Agent-based models: understanding the economy from the bottom up. *Bank of England Quarterly Bulletin* (2016), Q4.

[36] Nelson Vadori, Leo Ardon, Sumitra Ganesh, Thomas Spooner, Selim Amrouni, Jared Vann, Mengda Xu, Zeyu Zheng, Tucker Balch, and Manuela Veloso. 2023. Towards multi-agent reinforcement learning-driven over-the-counter market simulations. *Mathematical Finance* (2023).

[37] Nelson Vadori, Sumitra Ganesh, Prashant Reddy, and Manuela Veloso. 2020. Calibration of shared equilibria in general sum partially observable Markov games. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada). Curran Associates Inc., Red Hook, NY, USA, Article 1184, 11 pages.

[38] Xiaohan Wang, Lin Zhang, Yuanjun Laili, Kunyu Xie, Han Lu, and Chun Zhao. 2021. MADES: A Unified Framework for Integrating Agent-Based Simulation with Multi-Agent Reinforcement Learning. In *2021 Annual Modeling and Simulation Conference (ANNSIM)*. 1–12.

[39] Mason Wright and Michael P Wellman. 2018. Evaluating the stability of non-adaptive trading in continuous double auctions. In *17th International Conference on Autonomous Agents and Multiagent Systems*.

[40] Sihan Zeng, Thinh T Doan, and Justin Romberg. 2024. A two-time-scale stochastic optimization framework with applications in control and reinforcement learning. *SIAM Journal on Optimization* 34, 1 (2024), 946–976.

[41] Brian Zhang, Gabriele Farina, Ioannis Anagnostides, Federico Cacciamani, Stephen McAleer, Andreas Haupt, Andrea Celli, Nicola Gatti, Vincent Conitzer, and Tuomas Sandholm. 2024. Computing optimal equilibria and mechanisms via learning in zero-sum extensive-form games. *Advances in Neural Information Processing Systems* 36 (2024).

[42] Edwin Zhang, Sadie Zhao, Tonghan Wang, Safwan Hossain, Henry Gasztowtt, Stephan Zheng, David C. Parkes, Milind Tambe, and Yiling Chen. 2024. Position: Social Environment Design Should be Further Developed for AI-based Policy-Making. 235 (21–27 Jul 2024), 60527–60540.

[43] Liyuan Zheng, Tanner Fiez, Zane Alumbaugh, Benjamin Chasnov, and Lillian J Ratliff. 2022. Stackelberg actor-critic: Game-theoretic reinforcement learning algorithms. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 36. 9217–9224.

[44] Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C Parkes, and Richard Socher. 2022. The AI Economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science Advances* 8, 18 (2022), eabk2607.