

---

# LLM Economist: Large Population Models and Mechanism Design in Multi-Agent Generative Simulacra

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

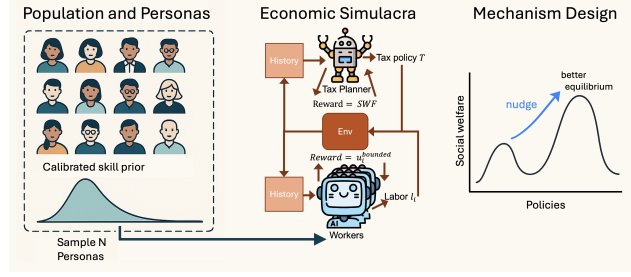
1 We present the *LLM Economist*, a novel framework that uses agent-based  
2 modeling to design and assess economic policies in strategic environments  
3 with hierarchical decision-making. At the lower level, bounded rational  
4 worker agents—instantiated as persona-conditioned prompts sampled from  
5 U.S. Census-calibrated income and demographic statistics—choose labor  
6 supply to maximize text-based utility functions learned *in-context*. At the  
7 upper level, a planner agent employs in-context reinforcement learning to  
8 propose piecewise-linear marginal tax schedules anchored to the current U.S.  
9 federal brackets. This construction endows economic simulacra with three  
10 capabilities requisite for credible fiscal experimentation: (i) optimization  
11 of heterogeneous utilities, (ii) principled generation of large, demographi-  
12 cally realistic agent populations, and (iii) mechanism design—the ultimate  
13 nudging problem—expressed entirely in natural language. Experiments with  
14 populations of up to one hundred interacting agents show that the planner  
15 converges near Stackelberg equilibria that improve aggregate social welfare  
16 relative to Saez solutions, while a periodic, persona-level voting procedure  
17 furthers these gains under decentralized governance. These results demon-  
18 strate that large language model-based agents can jointly model, simulate,  
19 and govern complex economic systems, providing a tractable test bed for  
20 policy evaluation at the societal scale to help build better civilizations.

## 21 1 Introduction

22 The rapidly expanding marketplace of autonomous language agents forms *economic simu-*  
23 *lacra*—synthetic societies whose allocation of effort and influence is governed by algorithmic  
24 code rather than legislation. As web-agents book tickets, draft briefs, and trade cryptocur-  
25 rencies, they adapt to digital incentives, creating complex economic ecosystems requiring  
26 governance to prevent early-mover exploitation.

27 Recent advances demonstrate remarkable potential for coherent multi-agent dynamics. Gen-  
28 erative Agents [24, 25] sustain believable interactions with thousands of persona-conditioned  
29 agents, Project Sid [1] scales toward "AI civilization" benchmarks, EconAgent [13] reproduces  
30 macroeconomic indicators with striking fidelity, and OASIS [32] explores large population  
31 simulacra of social media. These developments suggest LLMs exhibit sophisticated strategic  
32 reasoning [33], making them compelling substrates for policy experimentation.

33 In this work, we study designing tax mechanisms as a tractable and theoretically-grounded  
34 avenue for exploring governing agent societies. Classical optimal taxation faces two limitations



**Figure 1: LLM Economist Framework.** *Left:* Population of persona-conditioned agents. *Center:* Two-level economic simulacra. *Right:* Mechanism design via successive planner nudges.

in synthetic societies. First, solutions like the Saez formula [26, 27] assume fixed income elasticity, yet elasticity shifts dynamically with policy changes, making optimal rates moving targets requiring continuous recomputation. Second, human societies are heterogeneous and bounded rational [16, 18], while simulacra feature agents with text-specified motivations, requiring planners to reason over distributions of explicitly modeled personas.

We address both gaps by reframing optimal taxation as a repeated Stackelberg game optimized through two-level in-context reinforcement learning (ICRL) [10, 21, 22]. Building on the AI Economist’s deep RL approach [29, 34, 35], we replace value-function learning with interpretable language-based reasoning. Worker agents maximize persona-conditioned utilities via natural-language context encoding biographies, while a planner proposes tax schedules anchored to U.S. federal brackets through pure in-context optimization.

Our contributions: (i) *Large population models* [3] sampling Census-calibrated personas without manual utility engineering. (ii) In-context planners converging to Saez-level welfare through gradient-free optimization. (iii) Democratic turnover stabilizing outcomes and mitigating the Lucas critique [15] through synthetic counterfactuals.

## 2 LLM Economist

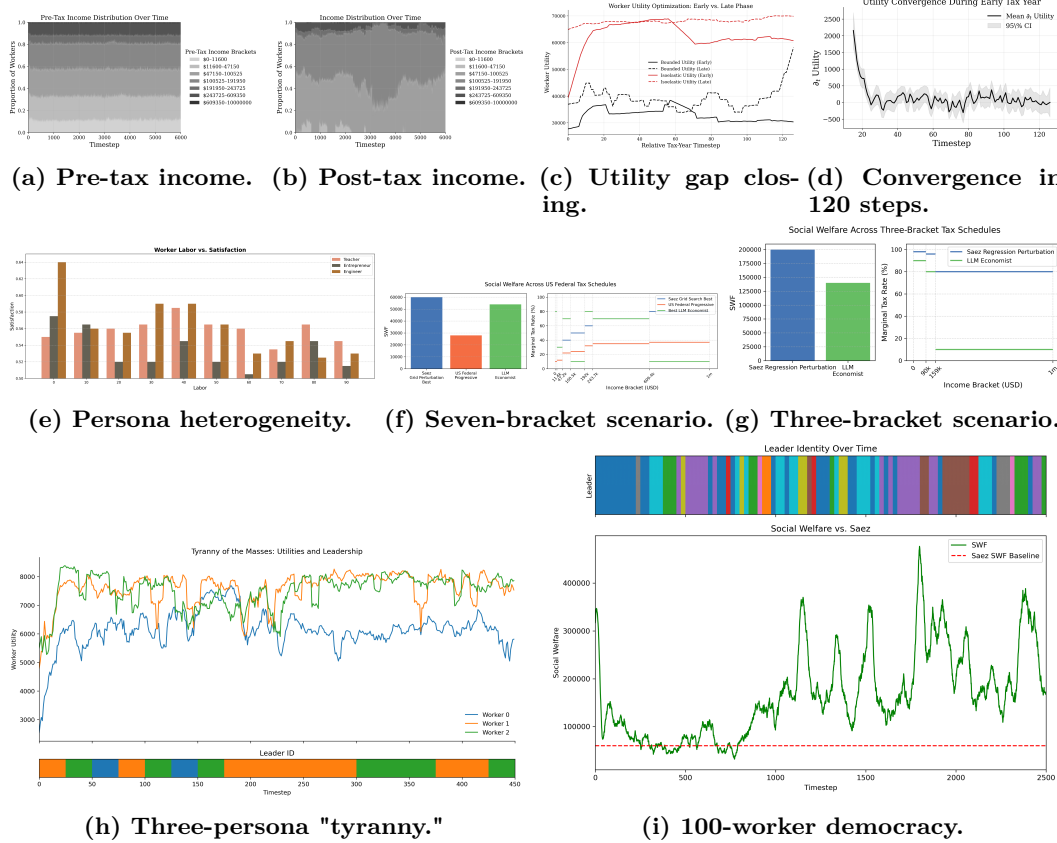
We model optimal taxation as a repeated Stackelberg game between a *planner*  $\mathcal{P}$  and workers  $\mathcal{W} = \{\mathcal{W}_1, \dots, \mathcal{W}_N\}$ . Time is divided into daily steps  $t = 0, \dots, T - 1$  and tax years of length  $K$ . Each worker  $i$  has latent skill  $s^i > 0$  and chooses labor  $l_t^i$ , yielding pre-tax income  $z_t^i = s^i l_t^i$ . The planner selects marginal tax schedule  $\tau_k$  at year start, giving post-tax income  $\hat{z}_t^i = z_t^i - T_{\tau_k}(z_t^i) + R_t$  where  $R_t$  is lump-sum rebate. Social welfare is  $\text{SWF} = \sum_{i=1}^N w(z_t^i) u_i(\hat{z}_t^i, l_t^i)$  with distributional weights  $w(z_t^i) = 1/z_t^i$ . A Stackelberg equilibrium satisfies:  $\tau^* \in \arg \max_{\tau} \mathbb{E}[\text{SWF}(\mathbf{1}, \tau)]$  and  $l^{i*}(\tau) \in \arg \max_{l^i} \mathbb{E}[u_i(\hat{z}^i, l^i)]$  for each worker.

The LLM Economist realizes this Stackelberg game through language-based agents acting purely *in-context*, where state, history, and objectives are rendered as text while actions are JSON snippets parsed by the environment. Skills  $s^i$  are drawn from generalized-Beta fits to 2023 American Community Survey data [30]. Each worker receives a persona prompt encoding demographics and preferences—*"You're a 32-year-old entrepreneur... You believe lower taxes let you reinvest in your company..."*—and uses bounded utility  $u_t^{\text{bounded}}(\hat{z}, l) = \frac{\hat{z}^{1-\eta}-1}{1-\eta} - \psi l^\delta - (1 - s_t^i)\phi$  where  $s_t^i \in \{0, 1\}$  is LLM-judged satisfaction and  $\phi$  is dissatisfaction penalty. Workers observe  $(z_t^i, \hat{z}_t^i, \tau(z_t^i), R_t, \text{history})$  and return  $\{\text{"LABOR": } \mathbf{X}\}$ .

The planner observes aggregate statistics and proposes bracket shifts  $\Delta \tau_k \in [-20, 20]^B$  via  $\{\text{"DELTA": } [\dots]\}$ . At each daily step  $t$ , the environment serializes joint state  $o_t$  into prompt  $\pi_t$ , following exploration-exploitation phases with broad search then convergence. Replay buffers maintain best state-action-welfare triples for token-level credit assignment across long horizons. Unlike the AI Economist’s value-function learning, this design eliminates task-specific reward shaping while exposing agents’ rationales, enabling interpretable policy optimization. This approach leverages LLMs’ ability to identify patterns in textual reward histories—a key advantage when preferences shift and causal links between individual utilities, policies, and outcomes must remain transparent.

(a) Tax-year length			(b) Test-time search			
Steps / yr	Total steps	%SWF*	Variant	Expl.+Expl.	No Explore	No Exploit
8	310	62.3	%SWF*	<b>84.9</b>	77.9	63.0
16	600	64.9				
64	2 000	84.9				
128	6 000	<b>90.0</b>				
256	6 000	90.0				

**Figure 2: In-context RL ablations.** (a) Welfare saturates at  $K = 128$ . (b) Exploitation and exploration are both critical.



**Figure 3: Experimental results.** (a-b) Income redistributes 15% downward. (c-d) Worker utilities adapt and converge. (e) Heterogeneous persona responses. (f-g) Tax schedules approach Saez optimum. (h-i) Democratic dynamics from tyranny to welfare-enhancing turnover.

### 3 Experiments

We evaluate: (i) design choices for planner and worker optimization, (ii) tax policy performance versus baselines, and (iii) emergent voting dynamics.

**Setup:** We use Llama-3.1-8B (though no discernable differences with other models, open-source and frontier, towards our hypotheses),  $N = 100$  workers,  $T = 3\,000$  steps with tax years  $K = 128$ . Skills follow GB2 calibrated to ACS 2023 [30]. Workers choose  $[0, 100]$  hours/week; planners optimize seven brackets with lump-sum rebates.

We compare against **Saez** (perturbed, intractable optimum) and **U.S. Fed** (2024 statutory rates).

### 85 3.1 Planner’s Social Welfare Optimization

86 The planner-worker interaction in the LLM Economist requires succesful ablation of two  
87 design choices to reach stable Stackelberg equilibria: *time-scale separation* between planner  
88 updates and worker adaptation, and balanced *exploration* and *exploitation* over tax years.  
89 Table 2a demonstrates that very short tax years ( $K \leq 16$ ) stall below 65% of optimal welfare  
90 because workers lack time to adapt, while performance plateaus at  $K = 128$  steps, capturing  
91 90% of the theoretical optimum. Meanwhile, Table 2b reveals that both exploration and  
92 exploitation are critical: LLM agents leverage their priors to reason about promising policies  
93 (exploitation) while requiring systematic search to discover optimal schedules (exploration),  
94 with exploitation being more impactful, but not sufficient. These results validate our  
95 hypothesis that in-context reinforcement learning can achieve near-optimal social welfare  
96 through careful design of temporal dynamics and test-time search.

### 97 3.2 Workers’ Utility Optimization

98 To test whether LLM workers optimize heterogeneous utilities under realistic income distribu-  
99 tions, we initialize skills using Generalized-Beta fitted to ACS 2023 microdata. Figure 3a-b  
100 show the learned policy redistributes 15% of workers to lower post-tax brackets while  
101 preserving aggregate labor. Figure 3c demonstrates bounded workers nearly close a 30k dis-  
102 satisfaction gap as the planner converges, while Figure 3e reveals persona-specific responses.  
103 These results validate that LLM workers coherently adapt labor choices under evolving tax  
104 incentives while maintaining realistic heterogeneity.

### 105 3.3 Tax Policy Evaluation

106 To test whether in-context reinforcement learning approach theoretically optimal policies, we  
107 compare against Saez baselines in two settings: bounded-utility workers (seven U.S. brackets)  
108 and isoelastic workers (three brackets). In the bounded case (Figure 3f), LLM Economist  
109 achieves +93% welfare versus U.S. baseline while approaching perturbed grid search Saez  
110 (+114%), with slightly less smooth schedules than the perturbed optimal. In the isoelastic  
111 case (Figure 3g), perturbed Saez outperforms but LLM Economist preserves labor supply  
112 through lower rates. The in-context RL planner achieves close to the Saez optimum without  
113 gradient information in a sample efficient manner, demonstrating that agent-based modeling  
114 approaches first-order optimal design—validating our hypothesis that LLMs can serve as  
115 effective mechanism designers.

### 116 3.4 Voting Simulacra

117 To test whether LLM agents reproduce political-economy phenomena, we introduce demo-  
118 cratic elections where agents elect planners by majority vote each tax year. Figure 3h shows  
119 classic "tyranny of masses" in a three-agent society: two workers exploit the minority without  
120 hard-coded rules. Figure 3i demonstrates that 100-agent leadership turnover enhances welfare  
121 through electoral exploration, sometimes outperforming static optimal taxation.

## 122 4 Discussion

123 This work introduces the *LLM Economist*, an in-context reinforcement learning framework  
124 that embeds a population of persona-conditioned agents and a tax planner in a two-tier  
125 Stackelberg game. Our results show that LLM agents can (i) recover the Mirrleesian trade-off  
126 between equity and efficiency, (ii) approach Saez-optimal schedules in heterogeneous settings  
127 where analytical formulas are unavailable, and (iii) reproduce political phenomena—such as  
128 majority exploitation and welfare-enhancing leader turnover—without any hand-crafted rules.  
129 Taken together, the experiments suggest that the LLM Economist can serve as tractable test  
130 beds for policy design long before real-world deployment, providing a bridge between modern  
131 generative AI and classical economic theory. While our approach assumes static skills and  
132 fixed population, and the framework could be misused to craft biased policies, it offers a  
133 controlled environment for exploring economic mechanisms before real-world deployment.

## References

- [1] A. AL, A. Ahn, N. Becker, S. Carroll, N. Christie, M. Cortes, A. Demirci, M. Du, F. Li, S. Luo, et al. Project sid: Many-agent simulations toward ai civilization. *arXiv preprint arXiv:2411.00114*, 2024. 1
- [2] T. B. Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [3] A. Chopra. Large population models. *arXiv preprint arXiv:2507.09901*, 2025. 1
- [4] A. Chopra, S. Kumar, N. Giray-Kuru, R. Raskar, and A. Quera-Bofarull. On the limits of agency in agent-based models. *arXiv preprint arXiv:2409.10568*, 2024.
- [5] P. A. Diamond and J. A. Mirrlees. Optimal taxation and public production i: Production efficiency. *The American economic review*, 61(1):8–27, 1971.
- [6] P. Duetting, V. Mirrokni, R. Paes Leme, H. Xu, and S. Zuo. Mechanism design for large language models. In *Proceedings of the ACM on Web Conference 2024*, pages 144–155, 2024.
- [7] X. Feng, Z. Wan, H. Fu, B. Liu, M. Yang, G. A. Koushik, Z. Hu, Y. Wen, and J. Wang. Natural language reinforcement learning. *arXiv preprint arXiv:2411.14251*, 2024.
- [8] S. Karten, A. L. Nguyen, and C. Jin. Pokéchamp: an expert-level minimax language agent. *arXiv preprint arXiv:2503.04094*, 2025.
- [9] A. Korinek. Generative ai for economic research: Llms learn to collaborate and reason. Technical report, National Bureau of Economic Research, 2024.
- [10] M. Laskin, L. Wang, J. Oh, E. Parisotto, S. Spencer, R. Steigerwald, D. Strouse, S. Hansen, A. Filos, E. Brooks, et al. In-context reinforcement learning with algorithm distillation. *arXiv preprint arXiv:2210.14215*, 2022. 1
- [11] J. Z. Leibo, V. Zambaldi, M. Lanctot, J. Marecki, and T. Graepel. Multi-agent reinforcement learning in sequential social dilemmas. *arXiv preprint arXiv:1702.03037*, 2017.
- [12] J. Z. Leibo, A. S. Vezhnevets, W. A. Cunningham, S. Krier, M. Diaz, and S. Osindero. Societal and technological progress as sewing an ever-growing, ever-changing, patchy, and polychrome quilt. *arXiv preprint arXiv:2505.05197*, 2025.
- [13] N. Li, C. Gao, M. Li, Y. Li, and Q. Liao. Econagent: large language model-empowered agents for simulating macroeconomic activities. *arXiv preprint arXiv:2310.10436*, 2023. 1
- [14] J. Liu, D. Shen, Y. Zhang, B. Dolan, L. Carin, and W. Chen. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*, 2021.
- [15] R. E. Lucas Jr. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. North-Holland, 1976. 1
- [16] R. D. Luce et al. *Individual choice behavior*, volume 4. Wiley New York, 1959. 1
- [17] N. G. Mankiw, M. Weinzierl, and D. Yagan. Optimal taxation in theory and practice. *Journal of Economic Perspectives*, 23(4):147–174, 2009.
- [18] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995. 1
- [19] J. A. Mirrlees. An exploration in the theory of optimum income taxation. *The review of economic studies*, 38(2):175–208, 1971.
- [20] J. A. Mirrlees. Optimal tax theory: A synthesis. *Journal of public Economics*, 6(4):327–358, 1976.

- [21] A. Moeini, J. Wang, J. Beck, E. Blaser, S. Whiteson, R. Chandra, and S. Zhang. A survey of in-context reinforcement learning. *arXiv preprint arXiv:2502.07978*, 2025. 1
- [22] G. Monea, A. Bosselut, K. Brantley, and Y. Artzi. LLMs are in-context reinforcement learners, 2024. URL <https://openreview.net/forum?id=YW791AHBUF>. 1
- [23] D. Paglieri, B. Cupiał, S. Coward, U. Piterbarg, M. Wolczyk, A. Khan, E. Pignatelli, Ł. Kuciński, L. Pinto, R. Fergus, et al. Balrog: Benchmarking agentic llm and vlm reasoning on games. *arXiv preprint arXiv:2411.13543*, 2024.
- [24] J. S. Park, J. O’Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023. 1
- [25] J. S. Park, C. Q. Zou, A. Shaw, B. M. Hill, C. Cai, M. R. Morris, R. Willer, P. Liang, and M. S. Bernstein. Generative agent simulations of 1,000 people. *arXiv preprint arXiv:2411.10109*, 2024. 1
- [26] E. Saez. Using elasticities to derive optimal income tax rates. *The review of economic studies*, 68(1):205–229, 2001. 1
- [27] E. Saez and S. Stantcheva. Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(01):24–45, 2016. 1
- [28] P. Srivastava, S. Golechha, A. Deshpande, and A. Sharma. Nice: To optimize in-context examples or not? *arXiv preprint arXiv:2402.06733*, 2024.
- [29] A. Trott, S. Srinivasa, D. van der Wal, S. Haneuse, and S. Zheng. Building a foundation for data-driven, interpretable, and robust policy design using the ai economist. *arXiv preprint arXiv:2108.02904*, 2021. 1
- [30] U.S. Census Bureau. American community survey, 2023 public-use microdata sample (pums). <https://www.census.gov/programs-surveys/acs>, 2023. Accessed May 14, 2025. 2, 3
- [31] R. Willis, Y. Du, J. Z. Leibo, and M. Luck. Will systems of llm agents cooperate: An investigation into a social dilemma. *arXiv preprint arXiv:2501.16173*, 2025.
- [32] Z. Yang, Z. Zhang, Z. Zheng, Y. Jiang, Z. Gan, Z. Wang, Z. Ling, J. Chen, M. Ma, B. Dong, et al. Oasis: Open agents social interaction simulations on one million agents. *arXiv preprint arXiv:2411.11581*, 2024. 1
- [33] Y. Zhang, S. Mao, T. Ge, X. Wang, A. de Wynter, Y. Xia, W. Wu, T. Song, M. Lan, and F. Wei. Llm as a mastermind: A survey of strategic reasoning with large language models. *arXiv preprint arXiv:2404.01230*, 2024. 1
- [34] S. Zheng, A. Trott, S. Srinivasa, N. Naik, M. Gruesbeck, D. C. Parkes, and R. Socher. The ai economist: Improving equality and productivity with ai-driven tax policies. *arXiv preprint arXiv:2004.13332*, 2020. 1
- [35] S. Zheng, A. Trott, S. Srinivasa, D. C. Parkes, and R. Socher. The ai economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science advances*, 8(18):eabk2607, 2022. 1