© September 2025 | IJIRT | Volume 12 Issue 4 | ISSN: 2349-6002

AI-Driven Policy Stress Testing in Synthetic Economies

Nirmala Talape[1], Aryan Patil[2]
[1,2.] *Economics Department Fergusson College*

***Abstract*— Policy failures are costly, slow to detect, and politically difficult to reverse. This paper proposes a rigorous framework for AI-driven policy stress testing using synthetic economies. large-scale, agent-based, data-conditioned simulations populated by heterogeneous households, firms, intermediaries, and a learning government. We introduce a modular architecture in which (I) micro-behavior is calibrated to empirical moments; (ii) markets clear through adaptive price discovery; and (iii) policy instruments are controlled by a reinforcement-learning (RL) policy-maker that is constrained by legal, fiscal, and equity guardrails. We define a family of stress regimes (supply shocks, demand collapses, climate events, banking runs, geopolitical trade disruptions) and propose counterfactual rollouts that quantify distributional impacts, macro-stability, and long-run resilience. Key contributions include: (1) a Policy Wind-Tunnel Protocol (PWP) for ex-ante evaluation under extreme yet plausible scenarios; (2) a Welfare-Stability Frontier (WSF) that trades off social welfare and systemic volatility; (3) Bias-Aware Learning (BAL) to prevent the RL agent from exploiting spurious correlations; and (4) a Civic Oversight Layer (COL) that makes model behavior auditable and democratically steerable. We illustrate the framework on three stylized policies. VAT adjustment, targeted UBI, and carbon dividends. and show how stress testing reveals nonlinear regime vulnerabilities that would be invisible to static models. The result is a policy design pipeline that is fast, transparent, and safer than trial-and-error in the real world.**

***Index Terms*— policy stress testing, synthetic economies, agent-based modeling, reinforcement learning, distributional effects, macroprudential design, welfare optimization, transparency**

I. INTRODUCTION

Public policy is typically drafted on partial evidence, implemented once, and evaluated retrospectively. When conditions shift. pandemics, energy shocks, bank runs. static models and point estimates fail. We argue that governments need the equivalent of an aerodynamic wind tunnel for policies: a place to push candidate interventions through turbulent, high-dimensional, and adversarial environments *before* adoption. Advances in AI, simulation, and data availability make this feasible.

This paper develops a synthetic-economy stress testing paradigm: an AI-orchestrated, modular simulation in which households and firms follow boundedly rational rules grounded in micro-evidence; markets clear via adaptive mechanisms; and a policy-making agent learns to adjust instruments subject to constitutional and fiscal constraints. We formalize the evaluation problem, propose new metrics for robustness, and present a governance architecture for legitimacy and safety.

1.1 Problem statement

Given a set of candidate policies and a family of stress regimes, we seek to compute the robust Pareto set of policies that maximize social welfare while respecting inequality, stability, and budget constraints under all. This requires (I) realistic micro-foundations, (ii) general-equilibrium feedbacks, and (iii) learning-based adaptation to non-stationarity.

1.2 Research questions

- RQ1: How can we design synthetic economies that are *empirically grounded* yet *computationally tractable* for policy wind-tunneling?
- RQ2: What metrics best capture the welfare–stability–equity trade-space, and how do they behave under stress?
- RQ3: Can a constrained RL policy-maker improve policy robustness without overfitting to a particular scenario?
- RQ4: What transparency and governance mechanisms are necessary for democratic oversight of AI-assisted policy design?

1.3 Contributions (summary)

1. Policy Wind-Tunnel Protocol (PWP): standardized, reproducible stress-test procedures for fiscal, regulatory, and climate policies.
2. Welfare-Stability Frontier (WSF): a two-dimensional efficient frontier mapping social welfare against macro instability; extended to a tri-objective with equity.
3. Bias-Aware Learning (BAL): regularization and counterfactual training to prevent spurious policy shortcuts (e.g., targeting the wrong proxies).
4. Civic Oversight Layer (COL): model cards, audit logs, contestability hooks, and simulatable explanations for all policy actions.

II. RELATED WORK AND GAP

- Macroeconomic DSGE models provide internal consistency but struggle with heterogeneity, fat tails, and rare shocks. Our synthetic economies relax representative-agent assumptions and admit non-equilibrium dynamics.
- Agent-Based Models (ABM) capture heterogeneity and emergent phenomena but have lacked standardized calibration, validation, and policy interfaces. We operationalize an ABM that is *calibrated to empirical moments*, *validated by back-casts*, and *controlled by a constrained policy learner*.
- Financial stress testing (post-GFC) focuses on banks under stylized macro paths. We generalize stress testing to whole-economy policy with multi-sector interactions and behavioral frictions.
- AI for policy has emphasized prediction and targeting. We shift focus to ex-ante design under uncertainty with guardrails and civic oversight.

Gap: No unified framework currently combines empirically grounded ABM, learning-based policy control, and transparent stress protocols to evaluate broad public policies under adversarial regimes. This paper fills that gap.

III. SYNTHETIC ECONOMY ARCHITECTURE

3.1 Agents and state

Let the economy at time be described by state were

- households (heterogeneous in income, liquidity, preferences, skills, location).
- firms (sector, productivity, leverage, inventories, wage offers).
- banks/intermediaries (capital, risk appetite, liquidity buffers).
- government (budget, instrument vector).
- relative prices, wages, interest rates, exchange rates.

Each household solves a bounded rational consumption–saving–labor problem with liquidity constraints and reference-dependent utility:

$$U_h (c_t, \ell_t) = u (c_t - \eta \\, \hat{c}_{h, t}) - v(\ell_t) + \phi \cdot S_h(t)$$

Firms choose prices and vacancies via adaptive heuristics with menu costs and inventory targets; banks allocate credit under VaR-like constraints with contagion on interbank networks.

3.2 Markets and clearing

Labor, goods, credit, and energy markets clear through decentralized matching with frictions. Price discovery is adaptive: Wages follow staggered contracts; unemployment emerges endogenously.

3.3 Policy instruments

Vector may include: VAT rate, payroll tax, targeted transfers/UBI, carbon price and dividend, credit risk weights, capital buffers, import tariffs, green subsidies, price caps/floors, and countercyclical buffers. Instruments are subject to legal and fiscal constraints.

3.4 Data conditioning and calibration

Calibration targets empirical moments (e.g., income/wealth distributions, firm size distribution, Okun's law slope, Phillips curve semi-elasticities, credit elasticities). We use indirect inference and approximate Bayesian computation (ABC) to fit parameters. Where micro data are scarce, we employ synthetic data synthesis constrained by public aggregates and privacy budgets.

IV. POLICY WIND-TUNNEL PROTOCOL (PWP)

4.1 Stress regimes

We define canonical families of shocks

- Supply: energy/food price spikes; climate disasters disrupting production; input bottlenecks.
- Demand: precautionary savings surges; abrupt fiscal consolidation.
- Financial: bank run; margin call spiral; sovereign spread shock.
- External: terms-of-trade shock; export bans; currency crisis.
- Epidemiological: contact-rate collapse with sectoral asymmetries.

Shocks may be single-hit, persistent, or regime-switching (Markov)

4.2 Guardrails and constraints

- Budget constraint: where is net revenue.
- Equity constraint: Gini or poverty headcount.
- Stability constraint: output volatility; bank default probability.
- Legal constraint: instrument bounds and rate-of-changes limits.

4.3 Experimental design

1. Baseline fit: calibrate on pre-policy data.
2. Counterfactual rollouts: for each policy and stress, simulate stochastic paths.
3. Hyper-sweeps: grid/BO over policy parameter space; compute frontiers.
4. Robustness: sensitivity to agent heuristics, network topology, and shock persistence.

V. THE WELFARE–STABILITY FRONTIER (WSF)

Define a composite objective for path

$$J_\omega(\pi) = \sum_{t=0} ^T \delta^t \left (W_t - \lambda_1 \cdot \text {Ineq}_t - \lambda_2 \cdot \text {Instab}_t - \lambda_3 \cdot \text {Emissions}_t \right)$$

We report:

- Welfare-robustness score: minimum-across-scenarios welfare (minimax).
- Tail risk index: CVaR of unemployment and inflation.
- Distributional impact: change in poverty and bottom-decile consumption.

VI. BIAS-AWARE LEARNING (BAL) FOR THE POLICY AGENT

We model the policy-maker as a constrained RL agent observing (with privacy-preserving aggregation). The agent selects to maximize subject to guardrails.

BAL components:

1. Counterfactual data augmentation: interleave on-policy rollouts with synthetic histories under swapped shocks to break brittle correlations.
2. Causal regularization: penalize policy moves that exploit non-causal predictors (measured via do-calculus proxies / invariant risk minimization across regimes).
3. Fairness constraints: demographic parity/benefit-incidence bounds on transfer policies.
4. Smoothness prior: total-variation penalty to avoid whiplash policy changes.

Optimization uses safe policy iteration with Lagrange multipliers for constraints and a shielding layer that projects proposed back into the feasible set.

VII. CIVIC OVERSIGHT LAYER (COL)

Legitimacy requires that policy suggestions be contestable and understandable.

- Model Cards for Policies: scope, assumptions, calibration targets, data lineage.
- Action Logs: every policy change is logged with (state, rationale, expected impact distribution, uncertainty bands).
- Simula table Explanations: any stakeholder can replay the simulation around a decision point with sliders.
- Red-Team Scenarios: user-submitted worst-case shocks are executable as first-class tests.
- Plural Objectives: citizens can re-weight to inspect trade-offs.

VIII. CASE STUDIES (STYLIZED)

8.1 Targeted UBI under a demand slump

- Setup: 2-sector (services, manufacturing); negative demand shock; liquidity-constrained households.
- Policies: (i) flat UBI; (ii) targeted UBI to bottom 40%; (iii) wage subsidy.

- Findings (stylized): targeted UBI dominates on WSF for same fiscal cost; wage subsidy stabilizes employment but can entrench low-productivity equilibria; mixing (ii)+(iii) yields best tail-risk reduction.

8.2 VAT reduction during a supply shock

- Setup: energy price spike with sectoral pass-through heterogeneity.
- Policies: (i) VAT cut; (ii) time-limited energy voucher; (iii) carbon dividend financed by windfall tax.
- Findings: VAT cut leaks to profits in constrained sectors; vouchers + dividend protect bottom deciles and dampen inflation persistence.

8.3 Carbon price with dividend in a trade shock

- Setup: import disruption raises input costs; green tech learning curves.
- Policies: carbon price ramp + equal dividend; green capex subsidy; border adjustment.
- Findings: dividend maintains political acceptability; capex subsidy accelerates learning, reducing medium-term CPI drag.

(These are scenario blueprints; empirical results will depend on calibration to the target economy.)

IX. VALIDATION, ROBUSTNESS, AND UNCERTAINTY

- Back-casting: fit to a historical period; test whether the synthetic economy replicates observable responses to known shocks.
- Moment-matching: ensure income/wealth distributions, Okun/Phillips elasticities, firm exit rates, and bank default frequencies match empirical ranges.
- Sensitivity: tornado charts over key parameters (habit strength, price rigidity, credit multipliers, network sparsity).
- Network stress: vary supply-chain topology from random to scale-free; assess contagion.
- Uncertainty quantification: ensemble over calibrated parameter sets; report Bayesian predictive intervals.

X. ETHICS, SAFETY, AND GOVERNANCE

- Privacy: use differentially private statistics for calibration; no raw microdata exposure.
- Value alignment: explicit objective weights are policy choices; enable citizen control of weightings.
- Non-deployment clause: simulation recommendations are decision aids, not auto-law. Human oversight mandatory.
- Failure modes: model misspecification, Goodhart's law, incentive gaming by stakeholders; mitigated by COL and periodic red-teaming.

XI. IMPLEMENTATION ROADMAP

1. Phase I. Prototype (0–3 months): minimal synthetic economy; calibrate to public aggregates; implement PWP v1 and WSF visualizations.
2. Phase II. Calibration (3–6 months): ABC/indirect inference; add credit and labor frictions; implement COL v1.
3. Phase III. Policy RL (6–9 months): integrate BAL; safe policy iteration with guardrails; stress library build-out.
4. Phase IV. External Audit (9–12 months): publish model cards; organize red-team challenge; iterate on governance.

Resources: GPU/CPU cluster for parallel rollouts; reproducible containerized code; open-source license where possible.

XII. REPRODUCIBILITY PACKAGE (SPECIFICATION)

- Code: modular ABM in Python/Julia; RL via JAX/PyTorch; experiment runner with seed control.
- Data: public aggregates; synthetic microdata with DP; calibration notebooks.
- Artifacts: WSF plots, policy action logs, model cards, stress libraries.
- Documentation: API for adding policies and stressors; governance checklist.

XIII. CONCLUSION

AI-driven policy stress testing in synthetic economies offers a practical path to safer, faster, and more democratic policy design. By combining empirically grounded simulation, constrained learning, and civic oversight, the framework uncovers nonlinear vulnerabilities, quantifies distributional stakes, and surfaces robust policy mixes before real-world adoption. The proposed PWP, WSF, BAL, and COL constitute a coherent pipeline that governments and multilateral institutions can adapt to their contexts.

REFERENCES

[1] Agent-based macroeconomics texts and survey papers.
[2] Literature on macro stress testing and systemic risk.
[3] Safe reinforcement learning and constrained MDPs.
[4] Differential privacy and synthetic data generation in economics.
[5] Policy transparency, model cards, and governance frameworks.
[6] Appendices. Appendix A Notation
[7] household; firm; bank; government.
[8] policy instrument vector; state vector; observation.
[9] objective; welfare; Instab instability; Ineq inequality.
[10] Appendix B: Pseudocode. Safe Policy Iteration with Shielding
[11] Initialize policy π_0 within legal/fiscal bounds. fork in 0.K
[12] for scenario s in stress library. for seed in Seeds. rollout economy E under π_k with shocks s. log (state, action, rewards, constraint usage). estimate value function V̂^{π_k} with conservative baselines. update policy via constrained optimization:
[13] π_{k+1} = argmax_π E[V̂^{π}(x)] - λ_1·Ineq - λ_2·Instab. subject to: budget ≥ 0, legal bounds, smoothness. apply shielding: project proposed actions back to feasible set. audit: produce explanations + counterfactual replays return Pareto-efficient set of policies on WSF
[14] Appendix C: Example Metrics (Operational Definitions)
[15] Equivalent Variation (EV): money metric of welfare at household level, aggregated by population weights.
[16] Tail Unemployment CVaR: expected unemployment above 95th percentile over horizons.
[17] Price-Stability Loss: sum of squared inflation gaps with asymmetric penalty on spikes.
[18] Fiscal Risk Score: probability that intertemporal budget constraint is violated under shocks. Appendix D: Visualization Checklist
[19] WSF chart (welfare vs. instability) per policy; highlight robust set.
[20] Poverty headcount trajectories with confidence bands.
[21] Sectoral output heatmaps across stress regimes.
[22] Action-log timeline for RL policy with explanations.
[23] Appendix E: Policy Templates (Parameterization Examples)
[24] Targeted UBI: benefit schedule by decile; phase-out rates; funding mix.
[25] VAT: baseline rate; temporary cut; sunset clause; compensatory transfers.
[26] Carbon Dividend: price path; dividend schedule; border adjustment parameters.
[27] Author Notes (delete before submission)
[28] Replace placeholders with your name/affiliation.
[29] Localize calibration targets to your country/region.
[30] Build a minimal prototype (even a toy ABM) and include one figure to demonstrate the WSF concept.
[31] Keep an audit trail for any data used; include a model card.