Munich Personal RePEc Archive

Adaptive Economic Governance Under Nonstationarity: A State-Contingent Alternative to Fixed Ideological Policy

Badreddine, Sadek

Independent Researcher

7 August 2026

Online at https://mpra.ub.uni-muenchen.de/130393/
MPRA Paper No. 130393, posted 03 Sep 2026 13:46 UTC

Adaptive Economic Governance Under Nonstationarity

A State-Contingent Alternative to Fixed Ideological Policy

Sadek Badreddine
Independent Researcher
sadekbader22@gmail.com
August 2026

Abstract

Economic government is often debated through persistent ideological labels even though the conditions that policy must address—demography, production, prices, migration, conflict, trade, housing, public finance, education, crime, and the distribution of income and wealth—change continually. This paper develops Adaptive Economic Governance System (AEGS), a constitutional and computational proposal that separates stable rights from revisable objectives and state-contingent policy. Under AEGS, fundamental rights operate as hard constraints. Elected, field-qualified representatives establish transparent social objectives. Auditable statistical models retrieve historical analogues, estimate policy outcomes and uncertainty, and generate recommendations; a language model may synthesize verified evidence and construct candidate explanations but cannot act as an outcome oracle. An elected expert congress approves, rejects, or remands recommendations. The system abstains when present conditions fall outside the support of its evidence. A synthetic proof of concept compares an adaptive historical-analogue rule with four fixed policy bundles over 400 paired, 180-period simulations. The adaptive rule reduces cumulative modeled social loss by 6.9 percent relative to the strongest fixed baseline (paired mean difference 64.68; 95 percent simulation interval, 63.52–65.85). It nevertheless performs worse than that baseline in several familiar regimes; its aggregate advantage is concentrated in a compound shock excluded from training, where loss is 18.5 percent lower. Review delays and removal of the abstention safeguard reduce the advantage. These results demonstrate a mechanism under stated assumptions, not real-world causal effectiveness. The paper concludes that adaptive economic governance is a testable institutional research program, but one that must solve difficult problems of democratic legitimacy, credentialism, model error, privacy, capture, and constitutional feasibility before deployment.

Keywords: adaptive governance; nonstationarity; economic policy; policy learning; artificial intelligence; technocracy; democratic accountability; robust decision-making

JEL classifications: C53, C54, D63, D78, E61, H11

AI-use disclosure. The author supplied the original thesis, normative commitments, and institutional design. OpenAI's ChatGPT was used substantially to assist with literature discovery, drafting, mathematical notation, Python simulation code, editing, and document formatting. The author is responsible for the paper's claims, interpretation, and errors. The computational evidence uses synthetic data.

1. Introduction

Public debate commonly asks whether capitalism, socialism, communism, or some hybrid is the best economic system. The question is too static. Modern economies are already mixed: markets, public provision, regulation, taxation, social insurance, monetary institutions, and industrial policy coexist in different proportions. The more useful question is therefore not which label should permanently govern, but which feasible policy mix is justified under the country's current state, uncertainty, objectives, and constitutional commitments.

The distinction matters because the environment facing government is nonstationary. Population ages; migration accelerates or reverses; a leading industry rises and declines; inflation changes the meaning of nominal programs; wars sever supply networks; new technologies alter labor demand; and the distribution of income, housing insecurity, educational attainment, and crime can shift at different speeds. A policy that is reasonable in a demand recession can be damaging in a supply shock. A trade strategy suited to a diversified peacetime economy can be fragile during conflict. A tax-transfer rule appropriate for one demographic structure may become fiscally unstable under another. Yet constitutions, statutes, agencies, party programs, and ideological commitments usually move more slowly.

This paper calls that tension the **constitutional-policy timescale mismatch**. Some slowness is valuable. Freedom of speech, privacy, due process, nondiscrimination, and other human rights should not fluctuate with the business cycle. Stable procedures also allow citizens and firms to plan. The mistake is to treat the economic instruments located inside those safeguards as if they must be equally permanent. A good constitutional order can be rigid about rights while remaining adaptive about tax rates, transfers, public investment, housing measures, trade resilience, and other economic tools.

The proposed Adaptive Economic Governance System (AEGS) makes that separation explicit. It has four layers. First, a constitutional layer fixes noncompensatory rights and places weapons and national-security command outside the economic recommendation system. Second, elected people with relevant advanced training set a public, multi-dimensional social objective. Third, an auditable policy engine uses historical analogues, causal evidence, forecasts, uncertainty estimates, and constrained optimization to recommend a policy bundle. Fourth, an elected expert congress—proposed here, for economic policy, as citizens holding doctorates in economics—approves, rejects, or remands the recommendation. The engine has no independent power to tax, spend, regulate, command armed forces, or suspend rights.

The proposal differs from both technocratic dictatorship and direct algorithmic government. Artificial intelligence is an advisory institution, not a sovereign. Elections remain the source of authority for the people who choose objectives and decide policy. The restriction of candidacy to experts is intended to raise competence, but it plainly narrows political equality and risks homogeneity. That is not a detail to be waved away; it is among the design's most serious normative and constitutional liabilities. The paper consequently treats credentialed eligibility as a hypothesis to test and supplies institutional counterweights rather than assuming that a doctorate produces wisdom, neutrality, or moral authority.

The paper makes four contributions. First, it translates a critique of fixed economic ideology into a layered institutional design: stable rights, slowly revised objectives, quickly revisable policies, and continuous audit. Second, it formulates a historical-analogue policy rule that makes model uncertainty

and out-of-distribution (OOD) conditions explicit. Third, it locates a limited role for large language models (LLMs): evidence retrieval, structured synthesis, candidate generation, and explanation, not unverified prediction or final decision. Fourth, it presents a reproducible synthetic experiment. The experiment does not validate AEGS for a real country. Its purpose is more modest: to show what would need to be measured, how a fixed rule can be compared with an adaptive one, how abstention affects performance, and where the proposed mechanism fails.

The principal result is deliberately mixed. Across the full simulated sequence, the adaptive rule outperforms the best fixed bundle by 6.9 percent in cumulative modeled loss. It loses, however, in several ordinary regimes and wins mainly in the compound novel shock. This is evidence against a triumphalist claim that adaptation always dominates. It instead supports a conditional claim: when states change materially, a governed ability to revise policy can be valuable, provided the system knows when its evidence is weak.

2. The Case for State-Contingent Governance

2.1 Ideologies as priors, not permanent programs

Economic traditions contain important knowledge. Market-oriented theories illuminate prices, decentralized information, entrepreneurship, and incentives. Socialist and social-democratic theories illuminate bargaining power, public goods, insurance, inequality, and the political construction of markets. Developmental-state traditions emphasize coordination, infrastructure, learning, and productive capacity. AEGS does not require erasing these traditions. It treats them as sources of hypotheses and candidate instruments rather than indivisible packages that must be accepted for all conditions.

This reformulation avoids a misleading history. Capitalism, socialism, and communism were not each invented at a single moment and then imposed unchanged. They developed through institutions, intellectual movements, political struggles, and country-specific mixtures. Nor are contemporary constitutions literally stationary. They are amended, interpreted, and supplemented. The relevant claim is comparative: structural and environmental conditions often change faster than foundational law, administrative routines, party coalitions, and public narratives.

Lucas (1976) demonstrated a related problem in economic evaluation: historical correlations can change when policy alters people’s expectations and behavior. Kydland and Prescott (1977) showed that discretionary optimization can be time-inconsistent, supporting credible rules. These arguments do not imply either frozen policy or unconstrained discretion. They imply that an adaptive system must model behavioral response, publish its rules of revision, bind itself to procedures, and distinguish precommitted constitutional constraints from state-contingent instruments. Robust-control approaches similarly ask decision-makers to account for model misspecification instead of optimizing as though a single model were certainly true (Hansen and Sargent 2001).

AEGS therefore uses ideologies as **priors**—structured expectations about mechanisms and trade-offs. Evidence should be able to update those priors. Some instruments may remain stable for long periods because repeated evaluation supports them. Others may change at predefined review points. Adaptation is not constant motion; unnecessary volatility imposes costs, destroys credibility, and makes policy itself a source of uncertainty.

2.2 Adaptation already exists

The broad intuition is not unprecedented. Progressive taxes and unemployment insurance act as automatic stabilizers, changing net taxes and transfers as income changes without a new vote (Auerbach and Feenberg 2000). Central banks condition decisions on inflation, employment, financial conditions, and forecasts. Fiscal rules often include escape clauses. Adaptive policy pathways set triggers for changing course as evidence accumulates (Walker, Rahman, and Cave 2001; Haasnoot et al. 2013).

These examples are a challenge to any claim of novelty. AEGS does not invent adaptive policy. It generalizes and constitutionalizes it. Existing mechanisms are often sector-specific, rely on a small set of indicators, and remain filtered through conventional partisan institutions. The proposal asks whether a wider policy portfolio can be governed by explicit state measurement, uncertainty-aware recommendation, rights constraints, expert review, and auditable revision. The burden is to show that this generalization adds value without producing unaccountable technocracy or destabilizing policy churn.

2.3 A multidimensional state and objective

No single indicator describes a country. Gross domestic product per capita can rise while housing insecurity, inequality, or preventable mortality worsens. The Stiglitz-Sen-Fitoussi Commission (2009) argued for moving beyond production toward household circumstances, distribution, quality of life, and sustainability. The OECD well-being framework likewise organizes multiple dimensions of current well-being and resources for the future. Sen's (1999) capabilities approach places substantive freedom —what people can actually be and do—at the center of evaluation.

AEGS begins with a state vector that includes, at minimum, population and age structure; immigration and emigration; real GDP and real GDP per capita; employment and unemployment; inflation; fiscal position; tax capacity and incidence; homelessness and housing burden; leading industries and concentration; corruption and state capacity; income and wealth distribution; educational access and attainment; imports, exports, and critical dependencies; violent conflict and national-security exposure; and crime and victimization. Definitions, lags, measurement error, and geographic distributions must be visible. The vector is not a leaderboard. It is a structured description used to frame causal questions.

The social objective is also plural. The baseline proposed here balances median material living standards, employment, price stability, housing security, education, inequality, public safety, fiscal sustainability, resilience, and individual freedom. The weights cannot be discovered from data alone. They are moral and political choices. Arrow's (1951) social-choice result is a warning against pretending that a uniquely neutral aggregation rule exists. Elected representatives must set and periodically reconsider weights through public reasons, within constitutional rights that cannot be traded away for aggregate welfare.

3. Related Literature

AEGS sits at the intersection of adaptive policy, statistical policy learning, political selection, and accountable artificial intelligence.

First, adaptive policy research treats uncertainty as a reason to establish monitoring, signposts, triggers, and pathways rather than a one-time optimal plan (Walker, Rahman, and Cave 2001; Haasnoot et al. 2013). Robust decision-making shifts attention from optimizing a point forecast to finding policies that perform acceptably across plausible models (Hansen and Sargent 2001). Nonstationary machine learning and dataset-shift research show why a model validated on one distribution may fail when the data-generating process changes (Quiñonero-Candela et al. 2009; Ditzler et al. 2015).

Second, policy learning estimates decision rules that maximize an explicit welfare criterion from observational or experimental data. Athey and Wager (2021) show how observational policy learning can use doubly robust methods under assumptions about identification. Synthetic control methods build comparison units from weighted combinations of unaffected cases and have become important for comparative case studies (Abadie and Gardeazabal 2003; Abadie, Diamond, and Hainmueller 2010). These methods are more disciplined than a slogan that "history repeats itself." Historical cases are conditionally relevant analogues, not templates. Their use requires comparable measurement, adjustment for confounding, sensitivity analysis, and a refusal to extrapolate when overlap is poor.

Third, research on political selection asks whether institutions attract and select competent, public-spirited officials (Besley 2005). Democracies appear more likely to select highly educated leaders (Besley and Reynal-Querol 2011), and cross-country evidence associates leader education with growth under particular research designs (Besley, Montalvo, and Reynal-Querol 2011). Credentials are at best noisy signals. Electoral accountability can reduce corruption: Ferraz and Finan (2011), using randomized audit disclosure in Brazilian municipalities, estimate substantially lower misappropriation where incumbents could seek reelection. Compensation may also affect candidate quality (Ferraz and Finan 2009). These findings support a combined emphasis on competence, incentives, disclosure, audit, and removal; they do not justify rule by degree-holders alone.

Fourth, economists have explored generative AI as a research tool and simulated agent (Korinek 2023; Horton, Filippas, and Manning 2023). LLMs can summarize a large evidence base, translate across specialties, and generate scenarios, but fluent output is not reliable causal inference. Foundation models inherit data bias, obscure provenance, and generate plausible falsehoods (Bender et al. 2021; Bommasani et al. 2021). Accountability consequently requires more than releasing code. Kroll et al. (2017) show that procedural mechanisms, testing, explanation, and institutional oversight can make automated systems accountable even when complete source-code transparency is insufficient.

The NIST AI Risk Management Framework identifies validity, reliability, safety, security, resilience, accountability, transparency, explainability, privacy enhancement, and management of harmful bias as dimensions of trustworthy AI (Tabassi 2023). UNESCO's Recommendation on the Ethics of Artificial Intelligence and the OECD AI Principles center human rights, democratic values, fairness, transparency, and human oversight. AEGS borrows this orientation but raises the standard because economic government is a high-impact public function: every recommendation, model version, data vintage, dissent, vote, and realized outcome should be reconstructible by independent auditors, subject to narrowly tailored protection for genuine security-sensitive information.

4. Institutional Architecture

4.1 Layer 1: constitutional rights and separated powers

The constitutional layer is deliberately hard to change. It protects freedom of speech and association, privacy, due process, equal protection, religious liberty, political and ideological pluralism, and nondiscrimination. The system may not recommend a policy that violates these guarantees, even if a model predicts higher output, lower crime, or lower fiscal cost. Rights are **noncompensatory constraints**: a predicted benefit elsewhere cannot offset their violation.

An independent judiciary reviews rights challenges and can enjoin a policy or the use of a data source. Citizens have notice, access, correction, appeal, and remedy rights when individual-level automated processing affects them. Data collection follows purpose limitation, minimization, retention limits, differential access, security controls, and independent privacy audit. Whenever policy can be evaluated with aggregate or privacy-preserving statistics, personally identifiable data should not enter the recommendation engine.

Weapons, military operations, emergency force, intelligence targeting, and direct national-security command remain under elected human authorities operating under ordinary constitutional and legal checks. Economic consequences of security events may enter the state vector, but AEGS cannot initiate force or acquire command authority through an “economic” recommendation.

4.2 Layer 2: publicly chosen objectives

The objective layer changes more slowly than ordinary policy. Citizens elect councils whose members have advanced qualifications in the relevant domain—economics for the economic objective, and corresponding expertise for health, environmental, technological, or legal components. These councils translate public commitments into observable objectives, ranges, priorities, and risk tolerances. They must publish the proposed weights, distributional effects, dissenting views, and reasons before adoption.

The user-proposed design excludes party or institutional affiliation and private benefit beyond salary. Taken literally, “unaffiliated with anyone” is impossible and would suppress ordinary civic life. A workable legal version focuses on conflicts of interest: no party office or party membership while serving; resignation from outside employment and boards; no outside compensation, paid consulting, gifts, or sponsored travel; public campaign financing; disclosure of prior employers, clients, donors, assets, and beneficial interests; divestment or a qualified blind trust; recusal; a post-service revolving-door period; random and risk-based audits; and civil and criminal enforcement by an independent integrity body. Past ties are disclosed, not erased.

Objective councils may revise weights only at scheduled intervals or after a declared, reviewable emergency. This prevents a governing majority from changing the scoreboard whenever results disappoint. Changes require public notice, an impact analysis, a recorded supermajority vote, and judicial review for conflict with rights. A citizen assembly selected by lot should issue a nonbinding public report on proposed objective changes. Although this addition does not give lay opinion a technical veto, it exposes expert reasoning to people who experience the consequences.

4.3 Layer 3: the adaptive policy engine

The policy engine maintains a versioned state database and a model registry. For each review, it:

1. validates and documents current measurements;
2. retrieves historically comparable country-periods;
3. assembles experimental, quasi-experimental, and structural evidence;
4. generates feasible policy bundles;
5. estimates outcomes, distributional effects, uncertainty, and implementation risk;
6. rejects bundles that violate rights, law, budget, or operational constraints;
7. stress-tests surviving bundles across alternative models and adversarial scenarios;
8. publishes a recommendation, confidence statement, assumptions, alternatives, and a plain-language explanation; and
9. specifies monitoring indicators, sunset dates, and reversal triggers.

Historical similarity is only a starting point. A case from another country may differ in institutions, culture, demographics, technology, monetary regime, external constraints, or measurement. The engine therefore uses analogues to define a local evidence neighborhood and generate questions; causal identification must come from appropriate designs. Analysts must disclose positivity or overlap failures and sensitivity to unobserved confounding. Where evidence does not support a policy comparison, the system says so.

The engine is an ensemble rather than one model. Components may include causal forests or doubly robust policy-learning estimators, structural macroeconomic models, microsimulation, input-output and network models, calibrated stress scenarios, and qualitative evidence coded under preregistered protocols. Model weights depend on out-of-sample performance and domain relevance, but minority models and expert dissent remain visible. A red team tests data poisoning, objective manipulation, proxy discrimination, distribution shift, cyber compromise, and strategically gamed indicators.

4.4 A limited role for language models

When the current state has close precedents, the LLM's role is retrieval-augmented synthesis: locate claims in a curated corpus of expert research and testimony, map disagreements, draft candidate bundles, and explain the verified quantitative analysis. Citations must link to source passages; unsupported claims are excluded. The LLM does not assign truth merely by consensus and does not estimate causal effects by verbal intuition.

When the state is unique or OOD, more LLM autonomy would be exactly backwards. Novelty is a reason for greater humility. The LLM may help experts enumerate mechanisms and scenarios, but the system must flag low support, present uncertainty, and abstain from a confident optimum. The default response is a reversible pilot, tighter monitoring, an expert hearing, and staged expansion only if preregistered outcomes justify it. Emergency action may be necessary, but it should have a narrow scope, short sunset, and expedited judicial and legislative review.

Every model interaction relevant to a recommendation is logged with model identity, system instructions, retrieved sources, tool outputs, stochastic settings where available, and human edits.

Sensitive personal information is excluded from general-purpose prompts. Models are evaluated for hallucination, ideological skew, disparate error, prompt injection, and stability before and during deployment. The public interface shows evidence and uncertainty, not an anthropomorphic "AI leader."

4.5 Layer 4: the elected expert congress

The economic congress is elected by universal adult suffrage, but candidacy is limited to people holding a recognized doctorate in economics. Districting or proportional representation should prevent geographic concentration, and terms should be staggered. For ordinary recommendations, a recorded majority approves, rejects, or returns the package with specified questions. Supermajorities are required for emergency powers, material changes to objective weights, overrides of independent audit findings, and policies with unusually large intergenerational or irreversible effects. Silence never counts as approval.

Committees include macroeconomics and public finance; labor and distribution; trade and industrial organization; housing and regional economics; data, methods, and model risk; and rights and distributional review. The last committee includes nonvoting constitutional lawyers, privacy specialists, statisticians, affected-community representatives, and ethicists. Members may commission independent replication. Dissenting congress members receive space in the official policy report, and all votes are public except where a court authorizes a narrow security exception.

An approved policy enters a time-limited implementation order with outcome ranges and automatic reconsideration triggers. Rejection preserves the lawful status quo unless an existing policy has expired. The congress can request a revised candidate but cannot secretly change the model output; its amendments are separately identified and scored. The executive branch implements approved law and reports administrative feasibility and realized outcomes.

4.6 Integrity and capture controls

Expert status does not eliminate self-interest. It can concentrate it. The proposal therefore treats anti-capture measures as constitutive, not decorative. Campaigns are publicly financed with equal qualifying grants and strict spending limits. Candidates disclose all funding and organizational support. Lobbying contacts are logged promptly. Officeholders receive a competitive, fixed salary and no other earned or unearned benefit connected to office. Procurement, grants, tax changes, and regulations receive machine-readable beneficiary analyses. An independent inspector general audits random recommendations and all unusually concentrated benefits.

The integrity body itself requires divided appointment, staggered terms, transparent procedures, judicial review, and protected funding. Whistleblowers receive confidential reporting channels and protection. Violations can trigger recusal, disgorgement, civil penalties, removal through due process, and prosecution under generally applicable law. Because no institution can remove all conflicts, the realistic objective is to identify, manage, expose, and sanction them.

5. Formal Model

Let the measured country state at review time t be $s_t \in \mathbb{R}^K$, and let $a_t \in A \subset \mathbb{R}^P$ be a policy bundle. In the proof of concept, the state is expressed as seven normalized social-loss components and the action has five intensities. A real implementation would preserve the underlying indicators and report distributional outcomes rather than compressing them prematurely.

The historical archive contains transitions $D = \{(s_i, a_i, s_{i+1}, z_i)\}_{i=1}^{N}$, where z_i records institutional and contextual features. For a candidate a, a local distance is

$$d_i(s_t, a) = \left[(s_t - s_i)' \Omega (s_t - s_i) + \eta \parallel a - a_i \parallel_2^2\right]^{1/2},$$

where Ω scales and weights state dimensions and η governs policy similarity. Within the k nearest observations, kernel weights w_i produce a one-step analogue forecast

$$\hat{s}_{t+1}(a) = \frac{\sum_{i \in N_k} w_i(s_t, a)\, s_{i+1}}{\sum_{i \in N_k} w_i(s_t, a)}.$$

The weighted local dispersion is $\hat{V}_t(a)$. This estimator is intentionally simple and interpretable. A production system would compare it with causal and structural estimates, correct for policy-selection bias, and use multi-period dynamics.

Let $x_j(s)$ be a normalized loss—for example, an unemployment gap, housing-insecurity rate, or education deficit—and $q_j \geq 0$ the publicly chosen weight. The modeled loss is

$$L_t(a) = \sum_{j=1}^{K} q_j x_j\left(\hat{s}_{t+1}(a)\right)^2 + \rho \parallel a \parallel_2^2 + \gamma \parallel a - a_{t-1} \parallel_2^2 + \lambda\, tr\{Q \hat{V}_t(a)\}.$$

The second term penalizes implementation intensity, the third discourages policy whiplash, and the fourth penalizes predictive uncertainty. The recommendation is

$$a_t^* \in arg \min_{a \in A_{feasible}(s_t)} L_t(a),$$

where $A_{feasible}$ is defined by constitutional rights, statutes, administrative capacity, fiscal constraints, and minimum service commitments. A rights-violating policy is absent from the feasible set; it is not assigned a finite penalty that can be outweighed by growth.

Define support distance D_t as the mean distance to neighboring state observations. If $D_t > \tau$, where τ is a validation-set threshold, the engine marks the recommendation OOD. It must then abstain from claiming an optimum and fall back toward a balanced, previously authorized policy while convening review. This creates an explicit answer to a central question: what happens when history does not contain a credible analogue?

The recommendation becomes law only through a human gate:

$$a_t^{law} = \begin{cases} a_t^*, & \text{if the expert congress approves;} \\ a_{t-1}^{law}, & \text{if it rejects and prior authority remains;} \\ \text{revised candidate}, & \text{if it remands.} \end{cases}$$

This gate makes responsibility traceable. The engine is responsible for the integrity of its recorded analysis; elected members are responsible for the decision.

6. Synthetic Proof of Concept

6.1 Purpose and design

The experiment is a computational thought experiment, not an empirical test of any country. Its purpose is to make the proposal falsifiable enough to inspect. The full code and output tables accompany the paper.

The simulation has seven bounded loss-state variables: income shortfall, unemployment gap, inflation deviation, wealth inequality, fiscal stress, housing-and-crime distress, and education deficit. The five policy instruments, each scaled from zero to one, are countercyclical support; progressive tax and transfer; public investment and education; housing and social services; and trade and industrial resilience. Policy effects vary by regime and include cross-effects and diminishing returns. State transitions also contain persistence and Gaussian disturbance.

An archive of 10,000 synthetic transitions is generated from five familiar regimes. Eighty-five percent is used to construct the analogue rule and the remainder to evaluate one-step forecasts and set the OOD threshold. The adaptive rule evaluates 68 candidate bundles using 30 nearest neighbors. It minimizes the weighted squared state losses plus implementation, adjustment, and predictive-uncertainty costs. If support distance exceeds the 97.5th percentile validation threshold, the rule moves only halfway from the previous action toward a balanced fallback.

The test sequence contains 180 periods: normal conditions, a demand recession, normal conditions, a supply-inflation shock, a migration-and-housing shock, a war-and-trade shock, a compound novel shock, and a final normal period. The compound regime is absent from the training archive. Four fixed bundles serve as comparators: market-leaning, redistributive, developmental, and balanced. They are labels for parameter vectors, not claims that real ideological systems can be summarized by five numbers.

The experiment runs 400 paired Monte Carlo paths. Each strategy on a path receives the same initial state and stochastic shocks, improving precision in their difference. The primary outcome is cumulative modeled social loss. A 95 percent simulation interval for the paired mean difference is calculated across paths. This interval describes Monte Carlo variation within the specified data-generating process; it is not a confidence interval for a real-world policy effect.

6.2 Predictive validation

Held-out root mean squared error (RMSE) is 0.080 for income shortfall, 0.041 for unemployment, 0.132 for inflation, 0.029 for inequality, 0.151 for fiscal stress, 0.031 for housing-and-crime distress, and 0.040

for the education deficit. The relatively weak inflation and fiscal forecasts are substantively important: an implementation should not hide aggregate accuracy behind dimensions that are easy to predict.

Across test paths, 12.7 percent of adaptive decisions are flagged OOD. In the deliberately novel compound regime, 93.6 percent are flagged. Thus, the safeguard generally recognizes the training gap, although false negatives and false positives remain possible.

7. Results

7.1 Main comparison

Table 1 reports the primary results. The market-leaning fixed bundle is the strongest fixed comparator, with mean cumulative loss of 939.37. The adaptive analogue rule has mean loss of 874.69, a reduction of 64.68 units or 6.9 percent. The adaptive loss is lower in all 400 paired paths; the 95 percent simulation interval for the mean paired reduction is 63.52 to 65.85. The balanced, redistributive, and developmental bundles perform substantially worse under these particular coefficients and regime durations. Their poor performance should not be generalized to actual countries or doctrines.

Table 1. Cumulative modeled social loss over 400 paired 180-period paths

Strategy	Mean	5th percentile	95th percentile	Adaptive advantage over comparator
Adaptive historical analogue	874.69	853.71	897.30	—
Market-leaning fixed	939.37	909.52	969.98	6.9%
Balanced fixed	1,199.75	—	—	27.1%
Redistributive fixed	1,554.77	—	—	43.7%
Developmental fixed	1,556.20	—	—	43.8%

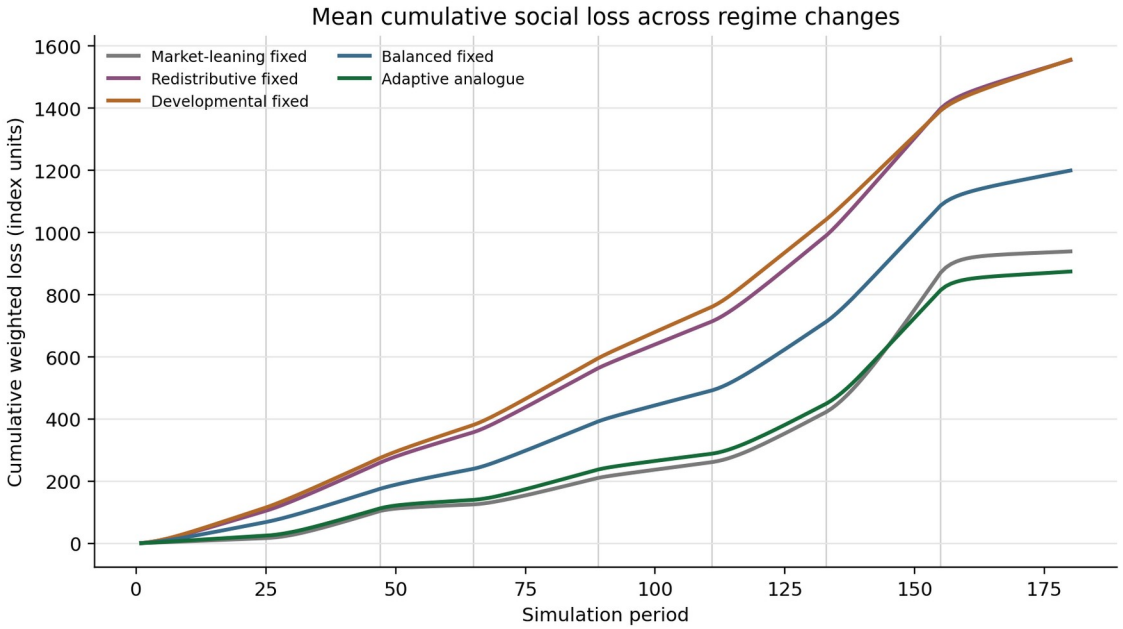

Figure 1. Mean cumulative modeled loss by strategy.

7.2 Regime heterogeneity

The aggregate result conceals a more instructive pattern. Relative to the strongest fixed rule, the adaptive rule is 5.3 percent worse in normal conditions, 0.7 percent worse in the demand recession, 14.6 percent worse in the supply-inflation shock, and 0.2 percent worse in the migration-and-housing regime. It is 0.3 percent better in the war-and-trade shock and 18.5 percent better in the compound novel shock.

Table 2. Mean per-period loss: adaptive versus strongest fixed comparator

Regime	Fixed	Adaptive	Adaptive difference
Normal	1.561	1.644	5.3% worse
Demand recession	3.964	3.991	0.7% worse
Supply-inflation shock	3.549	4.068	14.6% worse
Migration-and-housing shock	2.319	2.324	0.2% worse
War-and-trade shock	7.350	7.331	0.3% better
Compound novel shock	20.370	16.594	18.5% better

This is not evidence that an analogue model magically solves novelty. The improvement occurs because the fallback and state feedback happen to be safer than the fixed bundle under the simulated compound dynamics. The OOD flag correctly communicates that the evidence base is weak. The result supports a governed adaptation option, not blind trust in extrapolation.

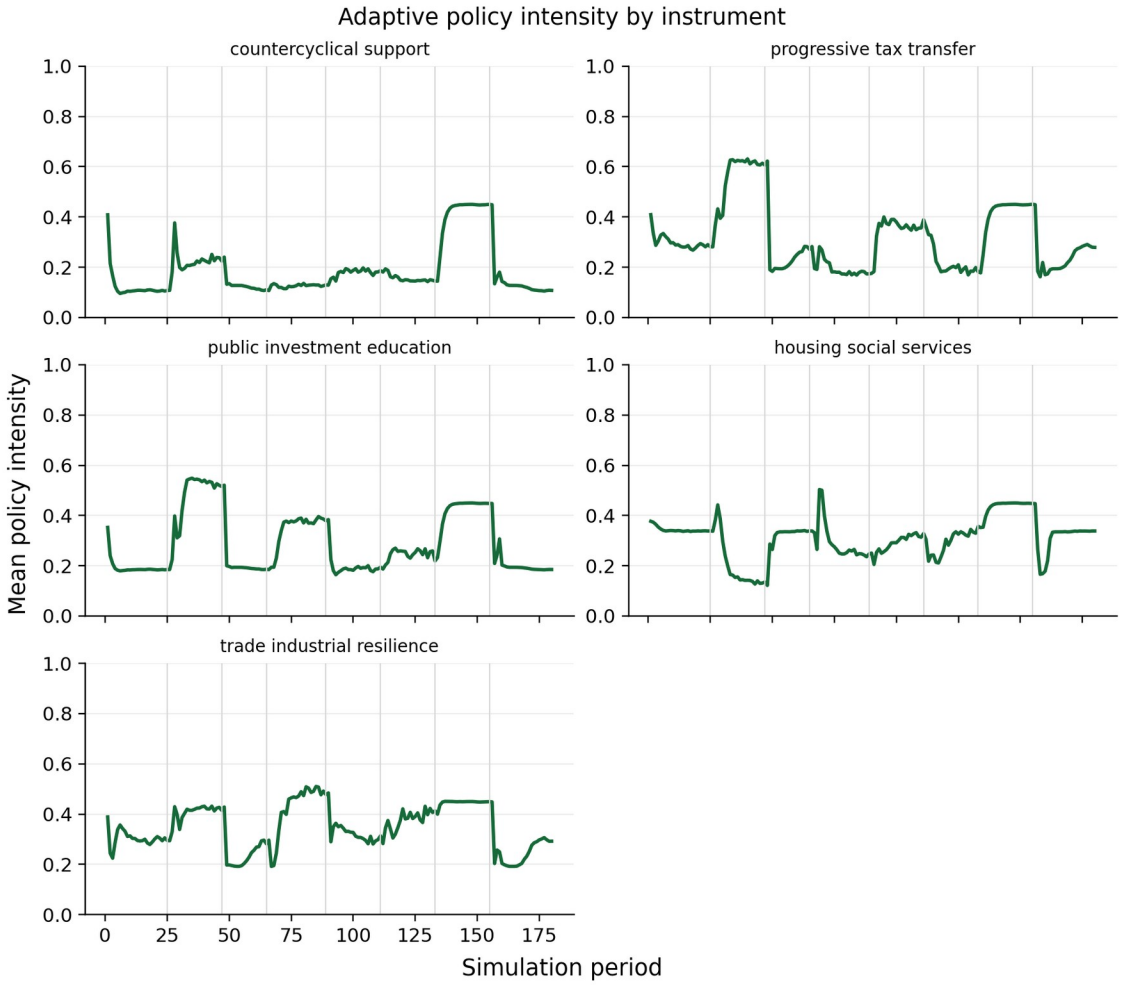

Figure 2. Average adaptive policy intensities across the simulated sequence.

7.3 Governance-oriented robustness checks

The baseline assumes review every period. To approximate legislative delay, the policy is alternatively reconsidered every three or six periods. A second variation removes OOD abstention. In 100 independently seeded paths, the adaptive rule with review every period and abstention reduces loss 5.0 percent relative to the market-leaning fixed comparator. Review every three periods reduces it 4.2 percent; review every six periods, 2.3 percent. Without abstention, the reduction is 2.6 percent and the adaptive rule wins in 85 percent rather than 100 percent of paths.

Table 3. Robustness to review timing and abstention (100 paths)

Adaptive specification	Mean cumulative loss	Reduction vs. fixed	Share of paths lower than fixed
Review every period; abstention	891.02	5.0%	100%
Review every 3 periods; abstention	898.22	4.2%	100%
Review every period; no abstention	913.69	2.6%	85%
Review every 6 periods; abstention	916.89	2.3%	89%
Market-leaning fixed	938.08	—	—

The smaller advantage under delayed review illustrates an institutional trade-off. Human deliberation and due process consume time but confer legitimacy, error correction, and rights protection. The proper conclusion is not to remove the human gate. It is to design review capacity, emergency procedures, and policy durations around the speed of the underlying risk.

8. U.S. Constitutional and Administrative Feasibility

AEGS as described is not implementable in the United States by ordinary statute. Article I establishes qualifications for Representatives and Senators, and the Supreme Court held in *U.S. Term Limits, Inc. v. Thornton* (1995) that neither states nor Congress may add qualifications to those specified in the Constitution. A doctorate requirement for membership in the existing Congress would therefore require constitutional amendment, or the body would have to be advisory rather than exercise Congress's legislative power.

A prohibition on party membership also creates a substantial First Amendment problem because political association is protected. Conflict-of-interest restrictions, disclosure, recusals, limits on outside compensation, and some forms of public financing are more familiar and more readily tailored than a categorical ban on association. The proposal should consequently be understood as a constitutional redesign for debate, not a description of currently lawful reform. Any amendment would need precise definitions of recognized degrees, equivalent expertise, candidacy, affiliation, enforcement, and judicial review.

There is also a structural danger in delegating legislative choices. Congress may obtain expert recommendations and establish conditional rules, but a system that effectively makes law without constitutionally accountable approval would raise nondelegation, due-process, and separation-of-powers concerns. The approval gate, published objective, bounded feasible set, and judicial review are meant to keep politically accountable humans responsible for law. Even with those safeguards, constitutional scholars would need to examine the design instrument by instrument.

A realistic U.S. research path would begin below the constitutional threshold:

1. create a nonbinding federal or state adaptive-policy laboratory with a multidisciplinary, politically balanced review board;
2. publish open benchmark datasets and prospective evaluation protocols;
3. test the system retrospectively on historical episodes without allowing it to act;
4. run narrow, reversible randomized or phased pilots where lawful and ethical;
5. compare recommendations with existing agency and legislative forecasts;
6. audit distributional, privacy, and ideological effects; and
7. decide only after evidence accumulates whether a new elected expert institution is desirable.

This path can test the analytical premise without prematurely entrenching the most controversial political premise.

9. Objections, Failure Modes, and Safeguards

9.1 Democratic legitimacy and credentialism

The strongest objection is that equal citizenship includes a meaningful opportunity to stand for office, not merely a vote among a certified elite. A doctorate is correlated neither perfectly with judgment nor with honesty. It can reward conformity, exclude practitioners and citizens with situated knowledge, and create a congress narrower in class, race, geography, ideology, and life experience than the population. Economics itself contains legitimate theoretical and methodological disagreement.

The proposal's reply is limited. Complex decisions benefit from demonstrated technical capacity, and voters retain the choice among qualified candidates. Public financing, plural electoral rules, transparent dissent, citizen review panels, and aggressive conflict controls can broaden accountability. Yet these safeguards do not dissolve the democratic cost. Comparative pilots should evaluate alternatives: competency examinations open to multiple educational pathways; mixed chambers of experts and ordinary citizens; a conventional legislature advised by an expert council; or randomly selected citizens supported by independent technical staff. The PhD-only congress remains the author's proposed baseline, not a proven optimum.

9.2 No unbiased objective or model

Calling AI "unbiased" would be false. Data encode past institutions and unequal treatment; measurement choices privilege some experiences; models omit variables; loss weights reflect values; and designers exercise discretion. AEGS can be less personally self-interested and more auditable than a single ruler while still producing systematic injustice. The safeguards are model cards and data sheets, protected-class performance analysis, distributional reporting, adversarial audits, public reasons, independent replication, and a right of appeal. None creates neutrality. They make bias more contestable.

9.3 Lucas critique, causal identification, and policy feedback

Historical responses may change once people anticipate the rule. A tax recommendation changes labor, investment, reporting, avoidance, and political organization. Cross-country associations may be confounded, and treatment effects may not transport. The engine must use structural reasoning where expectations matter, update after implementation, seek experiments and natural experiments, and report sensitivity. Policies should roll out gradually when learning value is high and irreversible harm is low. A superficially accurate predictor is insufficient if it predicts rather than causes the desired outcome.

9.4 Goodhart's law and data manipulation

When an indicator becomes a target, officials and private actors may game it. Crime can appear to fall because reporting changes; homelessness can be displaced outside a survey boundary; educational metrics can improve through teaching to a test; fiscal liabilities can be moved off balance sheet. AEGS mitigates this by using indicator families, independent data collection, rotating audits, qualitative checks, distributional decomposition, anomaly detection, and penalties for manipulation. Objective revisions must not be used to retroactively declare success.

9.5 Privacy and surveillance

An adaptive state could become a surveillance state if it treats more data as automatically better. The constitutional layer rejects that premise. Most macroeconomic decisions should use aggregate, sampled, or privacy-preserving data. Individual records require statutory purpose, necessity, proportionality, access controls, deletion schedules, and enforceable remedies. Sensitive traits must never become bases for religious, racial, ideological, or other unlawful discrimination. Fairness testing may require protected-class data, but such data should be segregated and governed for audit rather than used to ration rights.

9.6 Capture of experts and infrastructure

An expert congress may be captured by industries, universities, funders, vendors, or a shared professional ideology. The policy engine can be captured through procurement, training data, benchmark design, model access, or cloud infrastructure. Open standards, multiple vendors, public-interest compute, reproducible builds, secure logging, procurement transparency, independent teams, and regular rotation reduce common-mode failure. Critical decisions should not depend on one proprietary model or one organization.

9.7 Cybersecurity and adversarial manipulation

State-level recommendations create a valuable target. Attackers may poison data, alter a model, inject false documents, exfiltrate sensitive information, or disrupt availability during crisis. The system needs authenticated data provenance, segmented networks, least privilege, hardware-backed signing, independent replicas, incident response, manual fallback, and public notification rules. A compromised engine is immediately suspended; its prior authorization does not permit autonomous recovery or action.

9.8 Policy instability and administrative capacity

Frequent optimization can create whiplash. Households, firms, and agencies need predictable rules, and implementation takes time. Adjustment penalties, minimum policy durations, scheduled review, grandfathering, sunsets, and predefined triggers limit instability. Recommendations must include capacity constraints and transition plans. Sometimes the optimal recommendation is to preserve a reasonably robust policy rather than chase a small estimated gain.

9.9 Emergency power

War, pandemic, or financial panic compresses decision time. Emergency procedures should be constitutionally predefined, narrowly scoped, time-limited, and reviewable. AI can accelerate scenario analysis, but elected humans retain decision authority. Rights that permit no emergency limitation remain absolute; any legally limitable right receives necessity and proportionality review. Emergency actions expire automatically unless affirmatively renewed.

9.10 Limits of the proof of concept

The simulation is built by the same research process that proposes the adaptive rule. Its regimes, loss weights, policy effects, and comparator vectors are assumptions. The historical archive is synthetic and

does not contain strategic behavior, endogenous institutions, election cycles, heterogeneous households, general-equilibrium expectations, measurement manipulation, administrative failure, or international retaliation. The adaptive rule searches more actions than any fixed baseline and thus has an option-value advantage. The fixed vectors are not optimized separately for the known test distribution, and the compound shock occupies enough periods to influence the aggregate result.

Accordingly, the numerical result is not evidence that AEGS would raise welfare by 6.9 percent. It shows that under a transparent environment where the best action changes by state, adaptation with adjustment costs and abstention can beat fixed bundles overall while losing locally. Future work should preregister environments, invite adversarial benchmark design, compare against optimized robust rules and human committees, and use real historical backtests with information restricted to what was available at the time.

10. Empirical Research Agenda

An empirical program should proceed in four stages.

Stage 1: measurement and retrospective forecasting. Construct a country-period panel with versioned data, harmonized definitions, institutional covariates, and vintages. Forecast outcomes using only information available at each historical date. Compare analogue, structural, machine-learning, and expert forecasts. Report calibration and errors by regime and population group.

Stage 2: policy evaluation. Select bounded policies with credible identification, such as transfer formulas, tax credits, housing interventions, training programs, or procurement changes. Use randomized trials where ethical, regression discontinuity, difference-in-differences, synthetic controls, and doubly robust observational methods as appropriate. Evaluate transportability before using foreign analogues.

Stage 3: shadow governance. Run AEGS alongside actual decision processes without authority. Preregister objectives and recommendations. Independent teams score accuracy, realized welfare, distributional consequences, explanation quality, novelty detection, and the rate at which human reviewers correctly reject harmful recommendations. Red-team the data and models.

Stage 4: reversible pilots. Permit narrow recommendations to affect policy only where law authorizes experiments, rights risks are low, affected communities have recourse, and treatment can be reversed. Use sunset clauses and external evaluation. Expansion requires evidence of net benefit over existing institutions, not merely technical accuracy.

Table 4 maps the proposed state to plausible public sources. These datasets differ in coverage, delay, revision, and quality. Their presence does not establish causal comparability.

Table 4. Illustrative state variables and data sources

Domain	Measures	Illustrative sources
Population and migration	age structure, fertility, net migration	World Bank WDI; national statistical agencies
Production and prices	real GDP, GDP per capita, inflation, unemployment	IMF World Economic Outlook; World Bank WDI; national accounts

Domain	Measures	Illustrative sources
Public finance and tax	revenue, expenditure, debt, incidence	IMF fiscal databases; OECD; national finance ministries
Trade and industry	imports, exports, product concentration, critical dependencies	UN Comtrade; input-output tables; national industry data
Distribution	income and wealth shares, median income	World Inequality Database; household surveys; tax records under privacy safeguards
Housing and homelessness	rent burden, vacancies, sheltered and unsheltered homelessness	HUD AHAR/PIT for the United States; census and housing surveys
Education	attainment, access, learning and completion	UNESCO Institute for Statistics; OECD; WDI; national education data
Corruption and institutions	control of corruption, government effectiveness, audit findings	Worldwide Governance Indicators; justice and audit agencies
Crime and conflict	victimization, reported offenses, organized violence	UNODC; UCDP; national victimization surveys and administrative data

Evaluation should use a public scorecard: forecast calibration; causal-estimation validity; average and worst-group welfare; rights violations; privacy incidents; recommendation reversals; policy volatility; administrative cost; capture indicators; OOD detection; public comprehensibility; and appeal outcomes. AEGS should be abandoned or redesigned if it cannot outperform simpler, accountable alternatives.

11. Conclusion

The central choice is not between a permanently capitalist, socialist, communist, or technocratic state. It is between institutions that conceal adaptation inside partisan discretion and institutions that make adaptation explicit, constrained, evidence-based, and contestable. Economic conditions change. A serious system should be able to recognize that a policy fitted to one state may not fit another.

AEGS proposes stable constitutional rights, slowly revised public objectives, fast but bounded policy recommendations, elected expert approval, and continuous audit. It uses historical analogues carefully, treats novel conditions as a reason to abstain, and confines LLMs to evidence synthesis and candidate generation. Its synthetic demonstration shows a possible aggregate advantage over fixed bundles, but also losses in familiar regimes and sensitivity to governance delay and abstention. That mixed evidence is the appropriate starting point.

The proposal’s hardest questions are political, not computational. Who may stand for office? Who sets the objective? How are minority rights protected? How can experts be prevented from becoming a self-dealing class? How can citizens challenge a model they did not build? No optimization procedure answers these questions. The answer must be constitutional: AI recommends; elected humans decide; courts protect rights; auditors reconstruct the evidence; citizens retain speech, privacy, equality, association, and the power to replace decision-makers.

AEGS should therefore be judged as a research program, not a ready-made constitution. Its next test is whether transparent backtesting, shadow operation, adversarial evaluation, and reversible pilots can demonstrate benefits over simpler institutions without sacrificing democratic equality and human freedom. If it cannot pass that test, it should not govern.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program.” *Journal of the American Statistical Association* 105 (490): 493–505. https://doi.org/10.1198/jasa.2009.ap08746.

Abadie, Alberto, and Javier Gardeazabal. 2003. “The Economic Costs of Conflict: A Case Study of the Basque Country.” *American Economic Review* 93 (1): 113–132. https://doi.org/10.1257/000282803321455188.

Arrow, Kenneth J. 1951. *Social Choice and Individual Values*. New York: Wiley.

Athey, Susan, and Stefan Wager. 2021. “Policy Learning with Observational Data.” *Econometrica* 89 (1): 133–161. https://doi.org/10.3982/ECTA15732.

Auerbach, Alan J., and Daniel Feenberg. 2000. “The Significance of Federal Taxes as Automatic Stabilizers.” *Journal of Economic Perspectives* 14 (3): 37–56. https://doi.org/10.1257/jep.14.3.37.

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. https://doi.org/10.1145/3442188.3445922.

Besley, Timothy. 2005. “Political Selection.” *Journal of Economic Perspectives* 19 (3): 43–60. https://doi.org/10.1257/089533005774357761.

Besley, Timothy, José G. Montalvo, and Marta Reynal-Querol. 2011. “Do Educated Leaders Matter?” *Economic Journal* 121 (554): F205–F227. https://doi.org/10.1111/j.1468-0297.2011.02448.x.

Besley, Timothy, and Marta Reynal-Querol. 2011. “Do Democracies Select More Educated Leaders?” *American Political Science Review* 105 (3): 552–566. https://doi.org/10.1017/S0003055411000281.

Bommasani, Rishi, et al. 2021. “On the Opportunities and Risks of Foundation Models.” arXiv:2108.07258. https://doi.org/10.48550/arXiv.2108.07258.

Ditzler, Gregory, Manuel Roveri, Cesare Alippi, and Robi Polikar. 2015. “Learning in Nonstationary Environments: A Survey.” *IEEE Computational Intelligence Magazine* 10 (4): 12–25. https://doi.org/10.1109/MCI.2015.2471196.

Ferraz, Claudio, and Frederico Finan. 2009. “Motivating Politicians: The Impacts of Monetary Incentives on Quality and Performance.” NBER Working Paper 14906. https://doi.org/10.3386/w14906.

Ferraz, Claudio, and Frederico Finan. 2011. “Electoral Accountability and Corruption: Evidence from the Audits of Local Governments.” *American Economic Review* 101 (4): 1274–1311. https://doi.org/10.1257/aer.101.4.1274.

Haasnoot, Marjolijn, Jan H. Kwakkel, Warren E. Walker, and Judith ter Maat. 2013. “Dynamic Adaptive Policy Pathways: A Method for Crafting Robust Decisions for a Deeply Uncertain World.” *Global Environmental Change* 23 (2): 485–498. https://doi.org/10.1016/j.gloenvcha.2012.12.006.

Hansen, Lars Peter, and Thomas J. Sargent. 2001. “Robust Control and Model Uncertainty.” *American Economic Review* 91 (2): 60–66. https://doi.org/10.1257/aer.91.2.60.

Horton, John J., Apostolos Filippas, and John R. Manning. 2023. “Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?” NBER Working Paper 31122. https://doi.org/10.3386/w31122.

Korinek, Anton. 2023. “Generative AI for Economic Research: Use Cases and Implications for Economists.” *Journal of Economic Literature* 61 (4): 1281–1317. https://doi.org/10.1257/jel.20231736.

Kroll, Joshua A., Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. 2017. “Accountable Algorithms.” *University of Pennsylvania Law Review* 165 (3): 633–705. https://scholarship.law.upenn.edu/penn_law_review/vol165/iss3/3.

Kydland, Finn E., and Edward C. Prescott. 1977. “Rules Rather than Discretion: The Inconsistency of Optimal Plans.” *Journal of Political Economy* 85 (3): 473–491. https://doi.org/10.1086/260580.

Lucas, Robert E., Jr. 1976. “Econometric Policy Evaluation: A Critique.” *Carnegie-Rochester Conference Series on Public Policy* 1: 19–46. https://doi.org/10.1016/S0167-2231(76)80003-6.

OECD. 2024. “OECD AI Principles.” Paris: Organisation for Economic Co-operation and Development. https://oecd.ai/en/ai-principles.

Quiñonero-Candela, Joaquin, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence, eds. 2009. *Dataset Shift in Machine Learning*. Cambridge, MA: MIT Press.

Sen, Amartya. 1999. *Development as Freedom*. New York: Knopf.

Stiglitz, Joseph E., Amartya Sen, and Jean-Paul Fitoussi. 2009. *Report by the Commission on the Measurement of Economic Performance and Social Progress*. https://ec.europa.eu/eurostat/documents/8131721/8131772/Stiglitz-Sen-Fitoussi-Commission-report.pdf.

Tabassi, Elham. 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD: National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1.

UNESCO. 2021. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO. https://unesdoc.unesco.org/ark:/48223/pf0000381137.

Walker, Warren E., S. Adnan Rahman, and Jonathan Cave. 2001. “Adaptive Policies, Policy Analysis, and Policy-Making.” *European Journal of Operational Research* 128 (2): 282–289. https://doi.org/10.1016/S0377-2217(00)00071-0.

Appendix A. Simulation Parameters

The state-loss weights are 1.5 for income shortfall, 1.5 for unemployment, 1.1 for inflation deviation, 1.2 for inequality, 0.9 for fiscal stress, 1.4 for housing-and-crime distress, and 1.0 for education deficit. Implementation cost is $0.055 \| a_t \|^2$; adjustment cost is $0.10 \| a_t - a_{t-1} \|^2$; and predictive uncertainty cost is 0.18 times weighted local forecast variance. Transition noise has standard deviation 0.055. The fixed vectors are market-leaning $(0.15, 0.15, 0.20, 0.10, 0.20)$, redistributive $(0.65, 0.80, 0.45, 0.75, 0.25)$, developmental $(0.45, 0.50, 0.80, 0.55, 0.80)$, and balanced $(0.45, 0.45, 0.45, 0.45, 0.45)$. The main seed is 20260807; the robustness seed is 20260808.

Appendix B. Proposed Decision Record

Each recommendation should publish: decision identifier; date and jurisdiction; objective-function version; data sources and vintages; missingness and revisions; model names, owners, versions, and validation status; retrieved historical analogues; identification assumptions; candidate policies; rejected rights or feasibility screens; predicted aggregate and distributional effects; uncertainty and OOD status; stress tests; expert and model dissents; LLM prompts and verified sources where legally publishable; congressional committee record; individual votes and recusals; enacted text; implementation owner; monitoring indicators; sunset and reversal triggers; realized outcomes; appeals; incidents; and post-implementation audit.

Appendix C. Funding, Conflicts, and Data Availability

Funding. The author received no external funding for this research.

Conflicts of interest. The author declares no conflict of interest.

Data and code availability. The computational evidence is entirely synthetic. Source code, parameter files, and summary outputs are included in the accompanying reproducibility supplement. No human-subject or confidential data were used.