WORKING PAPER · MAY 2026

THE CRISIS RATIO

A STRUCTURAL MODEL OF U.S. FISCAL CASCADE RISK

A diagnostic framework for the post-automation fiscal transition

JAY SANDERS

Asheville, North Carolina

A NOTE FROM THE AUTHOR

I'm not a credentialed economist. I'm a musician and entrepreneur who has spent years thinking about fiscal responsibility and sustainability. I have witnessed the destruction of tax base revenue and its effects on our social fabric through my entire lifetime. Now, with the potential for AI displacement on the horizon, the implications are real and significant. I see a pattern that I believe most credentialed people are missing because their training disposes them to look at it through a specific lens. This paper is my attempt to articulate that pattern, to build a framework around what I'm seeing and to put it in a form serious readers can evaluate.

The paper was developed in collaboration with AI tools used as a writing and research assistant. The ideas, framework, and arguments are mine; the prose, calibration against published data, and citations were developed in dialogue over multiple sessions. I think this is increasingly how careful intellectual work outside of academic institutions will be produced, and I'd rather be transparent about the method than pretend otherwise.

I'm publishing this because the question matters and waiting for credentialed permission is not a strategy. The paper is offered as a working paper inviting critique. If parts of it are wrong, I want to know. If parts of it are useful, I hope they circulate. The argument either holds up on its merits or it doesn't, and the only way to find out is to put it in the world.

— Jay Sanders
Asheville, North Carolina · May 2026

ABSTRACT

This paper presents a structural model (the **Crisis Ratio**) that quantifies the relationship between three variables that govern U.S. fiscal sustainability under stress: unemployment, inflation, and the average interest rate on federal debt. The model is intentionally minimal. It compresses the standard sovereign debt sustainability framework (Blanchard, Reinhart-Rogoff, the CBO long-term outlook) into a single ratio that makes the cascade structure legible to non-specialist readers. The argument is diagnostic, not prescriptive, and proceeds in three movements.

The premise. The three variables are not free to move arbitrarily. They are coupled through canonical macroeconomic mechanisms (the Phillips curve, the Federal Reserve's reaction function, fiscal policy feedback) that, under ordinary shocks, move them in offsetting directions. This is why fiscal cascades are rare. Rising unemployment ordinarily produces falling inflation and falling rates, not all three deteriorating together.

The model. The Crisis Ratio expresses federal debt service as a fraction of available federal revenue, where revenue itself contracts as unemployment and inflation rise. The ratio is calibrated against actual FY2024 federal data and the published revenue elasticity literature. It responds to any combination of the three variables, and a dynamic extension carries the static scenarios forward as debt accumulates. The arithmetic is transparent and reproducible.

The alarm bell. AI cognitive displacement is the first plausible shock that overrides the canonical offsetting mechanisms and drives the three variables together, through a chain of identifiable economic links detailed in the paper. The standard policy off-ramps (quantitative easing, reserve currency demand, tax base expansion) do not function as off-ramps in this configuration, because each is endogenous to the cascade, actively shrinking, or politically constrained. Conditional on the override mechanism firing, the ratio rises along trajectories the model makes legible. The mechanism is contested in timing and magnitude. The arithmetic, conditional on the mechanism, is not.

CONTENTS

A Note from the Author **2**

Abstract **3**

1. Introduction **5**

2. The Model **6**

2.1 Specification 6

2.2 Calibration 7

3. Scenarios **9**

3.1 Scenario 0 (Current State (FY2024) 9

3.2 Scenario 1) Mild Cyclical Stress 9

3.3 Scenario 2 (Moderate Stress with Sticky Inflation 10

3.4 Scenario 3) Severe Stress 10

3.5 Scenario 4: Bond Market Cascade 11

3.6 Summary 11

3.7 Compounding within the static model 12

3.8 The Coupling Mechanism: AI Cognitive Displacement 13

A note on the Phillips curve 13

3.9 Dynamic Projection: What Happens When Debt Grows 23

4. Why the Off-Ramps Are Closed **28**

4.1 Quantitative easing does not break the loop; it reroutes it 28

4.2 Reserve currency status is the variable, not the constant 29

4.3 The tax base is contracting, not available for expansion 29

4.4 The fragilities, by contrast, are real and run the other direction 30

4.5 Why not Japan? 31

5. Literature and Limitations **33**

5.1 The Crisis Ratio in context 33

5.2 Coefficient sensitivity 34

5.3 What the model is not 35

5.4 What follows from the diagnosis 36

6. Conclusion **37**

Notes and Sources **39**

1. INTRODUCTION

Discussion of U.S. fiscal sustainability runs along two tracks. The technical track (Congressional Budget Office baselines, primary balance projections, debt-to-GDP ratios) is rigorous, careful, and unread outside the institutions that produce it. The political track (debate over deficits, taxation, entitlement reform) is widely consumed and rarely engages the technical work. The result is a public conversation that orbits the actual fiscal arithmetic rather than entering it.

This paper offers a small bridge. The Crisis Ratio is a closed-form expression that captures the standard sovereign debt sustainability mechanism in three variables. It is not a replacement for the technical literature, and it is not a forecast. It is a teaching framework designed to make the cascade structure legible: to show how unemployment, inflation, and interest rates couple together, why their compression produces a regime change rather than a linear deterioration, and why the conventional safety valves do not function as escape valves in this specific configuration.

The model is calibrated against actual FY2024 federal data: total receipts of $4.92 trillion, debt held by the public of $28.3 trillion, net interest outlays of $949 billion, and a weighted average effective interest rate on outstanding debt of approximately 3.3%.[1] Revenue elasticity coefficients are drawn from the published macroeconometric literature on automatic stabilizers and cyclically-adjusted revenue.[2] The threshold values reflect OECD experience with sovereign fiscal stress episodes.[3]

The argument proceeds in four parts. **Section 2** develops the model and presents its calibration. **Section 3** applies the model to five scenarios spanning current state to bond market cascade. **Section 4** examines why the standard off-ramps fail in this configuration. **Section 5** situates the model within the existing sovereign debt sustainability literature and names its limitations.

The paper concludes without policy recommendation. The intent is to make the diagnostic structure available to readers who may disagree about response but who would benefit from a shared picture of the underlying mechanism.

2. THE MODEL

2.1 Specification

The Crisis Ratio model takes three time-varying inputs and two structural parameters and produces a single scalar measuring fiscal stress.

Inputs:

- u (civilian unemployment rate (%)
- i) headline CPI inflation rate (%)
- r (weighted average effective interest rate on federal debt held by the public (%)

Structural parameters:

- D) federal debt held by the public ($ trillions)
- Y_0: baseline annual federal receipts under target macro conditions ($ trillions)

Revenue function

Federal receipts respond to the business cycle through automatic stabilizers. Higher unemployment depresses payroll, corporate, and activity-based revenues; inflation above target erodes real receipts through bracket creep lag, COLA indexation of outlays, and purchasing power loss on collected revenue. Both effects are bounded: receipts cannot fall below the structural floor that capital income, indexed bases, and non-cyclical sources provide, even under severe shocks.

$$\mathrm{Y(u,\,i)} = \mathrm{Y_0} \times [\mathrm{Y\_floor} + (1 - \mathrm{Y\_floor}) \times \max(0,\, 1 - \alpha{\cdot}\mathrm{u} - \beta{\cdot}\max(0,\, \mathrm{i} - \mathrm{i\_target}))]$$

where α is the unemployment revenue elasticity, β is the inflation revenue elasticity, *i_target* is the Federal Reserve's 2% target, and *Y_floor* is the minimum revenue ratio achievable under extreme cyclical conditions. The Y_floor term prevents the model from producing physically impossible revenue collapses; even at the trough of the Great Depression, U.S. federal receipts retained approximately 55% of baseline in real terms.

Debt service

$$\mathrm{DS} = \mathrm{D} \times \mathrm{r}$$

Debt service is the product of debt held by the public and the weighted average effective rate across the portfolio. The model uses debt held by the public rather than gross debt because intragovernmental holdings (Social Security trust fund, federal employee retirement) do not generate market interest costs and are not credibility-relevant in a bond market stress scenario.

The Crisis Ratio

$$\text{Crisis Ratio} = \text{DS} \div \text{Y}(\text{u}, \text{i})$$

The ratio expresses debt service as a fraction of available federal revenue. It is closely related to the conventional debt service ratio used in IMF Article IV consultations and OECD fiscal monitor reports, with the modification that the denominator is dynamic rather than static, since revenue contracts under the same conditions that expand debt service.

2.2 Calibration

The current calibration uses FY2024 actual fiscal data and revenue elasticity estimates from the published macroeconometric literature.

Parameter	Value	Symbol	Source
Federal debt held by public	**$28.3T**	D	Treasury, FY2024 actual
Baseline annual receipts	**$4.92T**	Y_0	CBO Monthly Budget Review, FY2024
Unemployment elasticity	**0.05**	α	CBO cyclically-adjusted revenue methodology; consistent with Auerbach-Feenberg
Inflation elasticity	**0.04**	β	Real receipts erosion accounting for bracket indexation and COLA pass-through
Target inflation	**2.0%**	i_target	Federal Reserve dual mandate
Revenue floor	**0.55**	Y_floor	Approximate Great Depression peak-to-trough real receipts decline

On α = 0.05. The unemployment revenue elasticity captures the combined effects of payroll tax loss, corporate tax decline, and activity drag on federal receipts per percentage point of unemployment. CBO's cyclically-adjusted revenue methodology implies an elasticity in the 0.04–0.06 range; the 2008–2009 episode saw federal receipts fall approximately 17% against a 4.3 point unemployment increase,

implying ~0.04 net of stimulus-related effects. The Auerbach-Feenberg work on automatic stabilizers supports values in this range. The model uses 0.05 as a midpoint.[4]

On β = 0.04. The inflation revenue elasticity captures real revenue erosion above the 2% target. Modern indexed bracket systems substantially reduce bracket creep effects relative to historical episodes; the binding channels are COLA pass-through to indexed outlays, lag between price increases and tax adjustment, and erosion of purchasing power on revenue collected during a given fiscal year. A coefficient of 0.04 is consistent with the empirical record of post-1985 inflation episodes.

On Y_floor = 0.55. The bounded specification reflects the empirical observation that federal revenue does not collapse to zero even under extreme cyclical stress. Capital income taxation, indexed bases, customs duties, and non-cyclical revenue sources maintain a floor. The 1932 trough of the Great Depression saw federal receipts fall by approximately 45% from 1929; the 0.55 floor is calibrated to that worst peacetime observation.

On the threshold. This paper uses 0.30 as the structurally significant threshold, the level at which debt service consumes the share of revenue that historically supports federal discretionary spending. Above 0.50, debt service exceeds the combined discretionary budget; above 1.0, the government must borrow to pay interest on existing borrowing. These thresholds are not magical; they reflect OECD empirical experience with sovereign fiscal stress episodes and the structural composition of the U.S. federal budget.[5]

3. SCENARIOS

The model is applied to five scenarios spanning current state to severe bond market cascade. The scenarios are not forecasts; they are configurations of the three variables that show how the ratio responds across the relevant policy-relevant range. Each scenario holds debt D constant at $28.3 trillion and baseline revenue Y_0 at $4.92 trillion. Section 3.9 relaxes the constant-debt assumption and shows that the dynamic ratios are substantially worse.

3.1 Scenario 0: Current State (FY2024)

$u = 4.0\%$, $i = 3.0\%$, $r = 3.3\%$

The FY2024 actuals. Unemployment near full employment; inflation above target but moderating; effective rate on debt held by the public reflects the weighted average of legacy issuance at lower rates and recent issuance at higher rates.

$$\text{Y} = \$4.92\text{T} \times [0.55 + 0.45 \times (1 - 0.05 \cdot 4.0 - 0.04 \cdot 1.0)]$$

$$\text{Y} = \$4.92\text{T} \times [0.55 + 0.45 \times 0.76] = \$4.92\text{T} \times 0.892 = \$4.39\text{T}$$

$$\text{DS} = \$28.3\text{T} \times 3.3\% = \$0.93\text{T}$$

$$\text{Crisis Ratio} = 0.93 \div 4.39 = 0.213$$

The current ratio sits at 0.213, approaching but not yet at the 0.30 threshold. Net interest outlays surpassed $1 trillion in FY2025 and now exceed defense spending; CBO projects the ratio to drift upward in coming years even on the assumption of stable employment and inflation.[6] The trajectory is not currently in crisis. It is moving toward the threshold under baseline assumptions.

3.2 Scenario 1: Mild Cyclical Stress

$u = 6.5\%$, $i = 3.5\%$, $r = 4.5\%$

A modest recession combined with persistent above-target inflation. Unemployment rises 2.5 points, which is well within historical postwar range; the Fed maintains restrictive policy, pushing the average rate on debt held by the public up 120 basis points as new and rolled issuance reprices.

$$\text{Y} = \$4.92\text{T} \times [0.55 + 0.45 \times (1 - 0.05 \cdot 6.5 - 0.04 \cdot 1.5)] = \$4.07\text{T}$$

$$\text{DS} = \$28.3\text{T} \times 4.5\% = \$1.27\text{T}$$

$$\text{Crisis Ratio} = 0.313$$

The ratio crosses the 0.30 threshold. Debt service consumes 31% of available revenue, eating into the share of the budget that historically supports federal discretionary spending. This is not yet acute crisis territory, but it represents a regime change in fiscal capacity. The government can still function; what it can no longer do is undertake significant new commitments without offsetting cuts or borrowing.

3.3 Scenario 2: Moderate Stress with Sticky Inflation

u = 8.0%, i = 4.5%, r = 5.5%

Recession-level unemployment combined with inflation that refuses to normalize. The Fed is trapped between fiscal stress and price stability; rates remain elevated even as the labor market weakens.

$$\text{Y} = \$4.92\text{T} \times [0.55 + 0.45 \times (1 - 0.05 \cdot 8.0 - 0.04 \cdot 2.5)] = \$3.81\text{T}$$

$$\text{DS} = \$28.3\text{T} \times 5.5\% = \$1.56\text{T}$$

$$\text{Crisis Ratio} = 0.408$$

Above 0.40. At this level, debt service exceeds combined federal spending on education, transportation, and most non-defense discretionary categories. The fiscal system retains the capacity to function but loses the capacity to respond to additional shocks. Each variable now pulls the others: rising rates compound debt service while elevated unemployment depresses revenue, and the political pressure to ease monetary policy to relieve fiscal stress conflicts with the inflation mandate.

3.4 Scenario 3: Severe Stress

u = 10.0%, i = 5.0%, r = 6.5%

Great Recession-level unemployment with inflation comparable to the early 1980s episode. The combined shock to revenue exceeds the 2008–2009 experience; debt service approaches the combined defense and education budget.

$$\text{Y} = \$4.92\text{T} \times [0.55 + 0.45 \times (1 - 0.05 \cdot 10.0 - 0.04 \cdot 3.0)] = \$3.55\text{T}$$

$$\text{DS} = \$28.3\text{T} \times 6.5\% = \$1.84\text{T}$$

$$\text{Crisis Ratio} = 0.519$$

Above 0.50. Debt service consumes more than half of available federal revenue. This is the regime in which conventional fiscal response (modest tax increases, modest spending cuts) cannot close the gap without major institutional restructuring. The government can continue to meet its obligations only through significantly expanded

borrowing, monetary accommodation, or both. Each of these options carries downstream consequences that feed back into the model.

3.5 Scenario 4: Bond Market Cascade

u = 12.0%, i = 6.0%, r = 8.0%

Every variable pushed beyond any single observation in U.S. postwar experience. Unemployment exceeds the worst of the 1982 recession; inflation at 1979 levels; rates reflecting both Fed tightening and a credibility premium priced into Treasury auctions.

$$Y = \$4.92T \times [0.55 + 0.45 \times (1 - 0.05 \cdot 12.0 - 0.04 \cdot 4.0)] = \$3.24T$$

$$DS = \$28.3T \times 8.0\% = \$2.26T$$

$$\text{Crisis Ratio} = 0.699$$

Approaching 0.70. Debt service consumes seven-tenths of all available federal revenue. The remaining 30% is insufficient to fund mandatory entitlement obligations, much less defense or discretionary categories. The system has exited the regime where standard fiscal policy applies. Operative pathways at this level are restructuring (selective default or maturity extension), monetary accommodation at scale (with consequences discussed in section 4), or institutional reorganization. The cascade has become reflexive: rising yields force debt service higher, which forces additional issuance, which forces yields higher.

3.6 Summary

Scenario	u	i	r	Y ($T)	DS ($T)	Ratio	Regime
0: Current	4.0	3.0	3.3	$4.39	$0.93	**0.213**	Approaching threshold
1: Mild	6.5	3.5	4.5	$4.07	$1.27	**0.313**	Above threshold
2: Moderate	8.0	4.5	5.5	$3.81	$1.56	**0.408**	Discretionary capacity eroded
3: Severe	10.0	5.0	6.5	$3.55	$1.84	**0.519**	Critical
4: Cascade	12.0	6.0	8.0	$3.24	$2.26	**0.699**	Reflexive collapse

3.7 Compounding within the static model

The scenarios in section 3 show that the Crisis Ratio rises more steeply when multiple variables move adversely together than when any single variable moves alone. This is mechanical: the formula contains three positive terms in u, i, and r, and adverse movement in each one contributes additively to the ratio. The compounding is in the arithmetic.

The static model does not, however, endogenize *why* the three variables would move together. The user supplies values of u, i, and r and the formula computes the ratio. Whether the variables in fact co-move under any specific shock is a separate question that the model itself does not answer.

Under most macroeconomic conditions the three variables are coupled by canonical mechanisms that move them in offsetting directions. The Phillips curve relates unemployment and inflation inversely under normal labor market conditions. The Federal Reserve's reaction function adjusts interest rates countercyclically in response to inflation. Fiscal policy responds to cyclical downturns with stabilizers and discretionary stimulus. Under these mechanisms, rising unemployment ordinarily produces falling inflation, falling rates, and stabilizing fiscal response. The cascade configuration the scenarios describe is therefore not produced by ordinary cyclical shocks. It would require something that overrides the canonical mechanisms.

The next section addresses the specific question of when and how the canonical mechanisms could be overridden. The argument is conditional and rests on a named mechanism: large-scale displacement of cognitive labor by artificial intelligence. The Crisis Ratio is the instrument that measures the consequences of that override. Section 3.8 names the cause.

3.8 The Coupling Mechanism: AI Cognitive Displacement

A note on the Phillips curve

The canonical macroeconomic framework treats unemployment and inflation as inversely related through the Phillips curve, and treats interest rates as the Federal Reserve's countercyclical response to inflation. Under those assumptions, the three variables this paper analyzes are not free to move arbitrarily. They are structurally coupled, and the canonical coupling moves them in opposite directions: tight labor markets produce wage pressure, wage pressure produces price pressure, the Fed responds by raising rates, the rates suppress demand, unemployment rises, the cycle resolves. The cascade configuration this paper describes would be ruled out by construction if the canonical coupling held without exception.

The argument of this section rests on three observations that complicate the canonical framework.

First, the Phillips curve has had a contested empirical life since at least 2010. The Federal Reserve itself has spent more than a decade discussing the "flat" or "missing" Phillips curve, in which U.S. unemployment fell to fifty-year lows without inflation moving meaningfully above the 2% target. The 2021-2022 inflation surge occurred without a preceding tightening of the labor market. Olivier Blanchard, whose r-vs-g framework anchors much of this paper's literature context, has been among the most prominent voices observing that the curve is empirically flat under normal conditions and nonlinear only under extreme stress.[22] The orthodox Phillips relationship is therefore not a binding constraint on the analysis at the level of normal cyclical fluctuation.

Second, the Phillips relationship operates through the labor market: tight labor markets produce wage pressure, wage pressure feeds into price pressure, the Fed responds. This mechanism requires the relevant unemployment to be wage-pressure-relevant unemployment. AI cognitive displacement produces a structurally different kind of unemployment. The displaced workers are high-wage cognitive workers replaced by capital-intensive automation; their displacement does not produce the demand-side wage pressure the canonical curve assumes. The 2009-2022 episode, in which productivity gains accrued to capital while real wages stagnated, demonstrated the same dynamic at smaller scale.

Third, the cascade scenario this paper describes is not a configuration in which u and i move together in the way the orthodox Phillips relationship would require to be violated at the wage-CPI channel. It is a configuration in which u rises (labor market

contraction in cognitive occupations), goods inflation may remain muted or fall (because automation reduces unit costs and demand destruction is disinflationary), asset inflation rises (because productivity gains accrue to capital seeking yield), and the Fed faces a policy trap that holds r elevated relative to where countercyclical policy would normally place it. The relevant inflation channel is therefore asset prices and selected goods sensitive to supply-chain or tariff effects, not the wage-driven CPI inflation the canonical Phillips curve was calibrated against.

The honest framing is that the Phillips curve is the right reference framework for the analysis it was designed to describe, which is cyclical fluctuation in a closed economy with full domestic production and labor-driven price dynamics. It is not the right framework for analyzing structural shifts in the relationship between labor share, productivity capture, and asset price dynamics. The cascade this section describes operates in the second territory rather than the first. The paper does not argue that the Phillips curve is wrong. It argues that the curve is silent on the configuration the cascade requires, and that a different coupling mechanism (the chain described below) drives the variables together where the canonical mechanism would have driven them apart.

The conditional argument

The cascade configuration the scenarios describe is not produced by ordinary shocks operating through canonical mechanisms. It would require something that overrides those mechanisms. This section argues that AI cognitive displacement at the pace and scale the labor-exposure literature suggests is a plausible candidate for that override, and walks through the chain by which a single shock produces co-movement in u, i, and r through coherent economic channels rather than through generic stress correlation.

The argument is conditional. The Crisis Ratio is a measurement instrument for what happens *if* the override mechanism fires. The mechanism is contested in timing and magnitude. The arithmetic, conditional on the mechanism, is not.

Why this shock is different

Every previous wave of automation in U.S. history shared a common property: the displaced sector was below the cognitive tier, and the cognitive tier remained intact as an absorption mechanism. Agricultural workers displaced in the early twentieth century were absorbed into manufacturing. Manufacturing workers displaced from the 1970s onward were absorbed into services. Service workers displaced in the 2000s were absorbed into knowledge work. At each step, the workers displaced were

lower-wage, lower-margin, and lower-credentialed than the workers in the absorbing tier. The economic shock was real but contained. Aggregate demand recovered because the absorption tier maintained income for the displaced cohort, even if at lower wages.

AI cognitive automation breaks this pattern in three structurally distinct ways.

First, the cognitive tier *is* the absorption tier. There is no obvious layer of work above it that can absorb displaced cognitive workers in the way that services absorbed displaced manufacturing workers. The Eloundou et al. (2024) and Anthropic (2026) labor exposure studies are consistent in finding that the highest theoretical task exposure to large language models is concentrated in occupations requiring four-year degrees and above, the cohort that has historically been the absorption destination for displacement from below. Displacing this cohort therefore does not produce mobility into a higher tier. It produces mobility into a lower tier, which is wage-deflationary economy-wide rather than wage-neutral.

Second, the displaced cohort is the cohort with the highest marginal propensity to consume and the highest debt service obligations. The professional class carries the bulk of U.S. mortgage debt ($12.6 trillion), student loan debt ($1.62 trillion), and revolving credit ($1.21 trillion). NY Fed Q4 2024 data shows household debt at $18.04 trillion concentrated in the credit-active middle and upper-middle income brackets. Displacement of this cohort therefore does not reduce aggregate demand by the displaced wage alone. It reduces demand by the wage *plus* the debt service that becomes unserviceable when the wage stops. This is the mechanism by which an income shock becomes a balance sheet shock.

Third, AI cognitive automation occurs in parallel with continued physical automation. The standard creative destruction defense assumes displaced workers can move toward new productive uses. China’s manufacturing experience between 2013 and 2025 demonstrates what happens when this assumption fails: manufacturing employment fell by more than 30 million workers (approximately 26% of peak employment) even as manufacturing output reached record levels. Robot density rose from 49 per 10,000 workers in 2015 to over 400 by 2025. Physical and cognitive automation are not happening in sequence with intervening absorption windows. They are happening in parallel, with the absorption tier being automated at the same moment displacement from below accelerates.

The combination of these three properties means AI cognitive displacement is not a faster version of prior automation waves. It is a structurally different kind of shock,

with propagation properties the standard absorption framework was not designed to handle.

The chain

The mechanism by which AI cognitive displacement couples u, i, and r operates through nine identifiable links. Each link rests on documented economic relationships; the chain as a whole is conditional but each step is defensible.

Link 1: AI displaces cognitive workers, not random workers.

This is the strongest link in the chain. Multiple independent research programs converge on the finding that cognitive tasks face 60–90% theoretical automation exposure, concentrated in high-wage knowledge-work occupations. The Anthropic labor index found computer and mathematics occupations showing 94% theoretical task exposure against 33% observed real-world usage, a deployment lag that is closing as inference costs fall. Gartner projects inference costs declining by more than 90% by 2030. The gap between theoretical exposure and observed deployment is therefore a timing question, not an existence question.

Link 2: The displaced cohort carries the bulk of consumer debt and the highest marginal propensity to consume.

NY Fed household debt data, ITEP analysis of household income distribution, and standard MPC estimates from the consumption literature all support this. The professional cohort consumes more per household than the working-class cohort, carries more debt per household, and is more credit-active. Income shocks to this cohort therefore produce disproportionate effects on aggregate demand and on financial system stability through the debt service channel.

Link 3: Aggregate consumer demand contracts as the displaced cohort loses income.

The 2008–2009 episode demonstrated this mechanism in operation. Consumer retrenchment by the middle-income cohort, not poverty-class spending changes, drove the demand collapse. The Garmaise-Levi-Lustig research tracking nearly ten million consumers found that sustained twelve-month unemployment highs trigger cumulative spending drops near 5% with no evidence of subsequent reversal, and that employed people change their spending behavior in anticipation of displacement, not only in response to it. The mechanism is well-documented.

Link 4: Capital captures productivity gains as profits rather than distributing them as wages.

The wage share of GDP has fallen from approximately 63% in 1970 to approximately 58% in 2024. The 2009–2022 period featured sustained productivity gains with stagnant real wages, particularly in the lower-wage segments of the labor market. The Bureau of Labor Statistics productivity data and Federal Reserve Distributional Financial Accounts both document the pattern. The question is not whether capital has captured productivity gains historically (it has) but whether this pattern continues under AI deployment. The answer turns on the relative bargaining power of labor versus capital in a context where labor's outside options are themselves being automated. There is no obvious reason to expect the pattern to reverse.

A note on the composition of revenue contraction. The α coefficient is calibrated to historical episodes in which payroll, individual income, and corporate tax receipts contracted in roughly fixed proportion: the 2008–2009 episode is the model calibration source. The AI cognitive displacement scenario decouples these channels. Payroll and individual income receipts fall (in proportion potentially larger than the historical baseline, because the displaced cohort is higher-wage), but corporate receipts may stay flat or rise as productivity gains accrue to capital rather than to labor. The effective corporate tax rate paid by large multinationals has trended toward 9–13% on book income through legal mechanisms including GILTI deductions, FDII, accelerated depreciation, stock-based compensation deductions, and foreign tax credit utilization. The Inflation Reduction Act's 15% corporate alternative minimum tax has produced less revenue than projected, and the 2025 H.R. 1 provisions expanded several of these mechanisms.[16] Capital gains on asset appreciation are not taxed until realized, and the largest holders can defer indefinitely through step-up basis at death.

The total revenue contraction in an AI scenario therefore arrives at roughly the magnitude implied by historical α, but through a structurally different channel mix: payroll and individual income taxes contracting more sharply than the historical baseline would suggest, while corporate taxes fail to provide the cyclical offset that operated in past recoveries. The formula's magnitude survives the composition shift by coincidence. The interpretation does not. In historical cyclical episodes, α measured how revenue responds to demand collapse with profits and payroll falling together. In the AI scenario, α measures something different, namely the structural inability of the revenue base to capture value when that value flows to capital rather than to labor. The cascade does not require corporate profits to fall. It requires only

that the labor channels contract while the capital channels remain effectively untaxed.

The implication for the off-ramps argument in Section 4 is that monetary stimulus cannot restore fiscal balance through the profits-and-corporate-tax channel that operated in past cyclical recoveries. The standard countercyclical mechanism (Fed cuts rates, demand recovers, profits return, tax receipts return) assumes the profit-tax channel is functional. If profits are being captured without generating proportional revenue, monetary stimulus does not restore the fiscal balance. It accelerates the wealth concentration that produced the cascade in the first place. This is one additional reason, beyond the inflation channel already discussed in Section 4.1, that QE does not function as an off-ramp in this configuration.

Link 5: Wealth concentration drives asset inflation while goods inflation may stay muted.

This is where the AI cascade story diverges most clearly from the standard inflation-recession framework. In a classic recession, both asset prices and consumer prices fall together. In a stagflation episode like the 1970s, both rise together. The AI cascade produces a third configuration: asset prices rise (because capital captures productivity gains and seeks yield in a low-real-rate environment) while consumer prices may fall or remain muted (because automation reduces unit production costs for many goods and services).

This matters for the Fed's policy reaction function. If headline CPI inflation remains near or below target while asset prices inflate and the real economy contracts, the Fed faces a signal that does not fit its mandate cleanly. Standard policy rules would suggest accommodation. But accommodation feeds the asset inflation that is concentrating wealth and worsening the demand collapse. This is the precondition for the policy trap in Link 6.

The link is less mechanically certain than Links 1–4. It depends on the relative balance of demand destruction (which is disinflationary) against supply chain stress, fiscal stimulus, and asset-driven wealth effects (any of which could push consumer prices up). The honest statement is that the *configuration* of mixed price signals is the central possibility the cascade story rests on, and that configuration is consistent with both economic theory and recent precedent (the 2009–2022 episode produced exactly this pattern in attenuated form).

Link 6: The Federal Reserve faces a policy trap.

If Link 5 produces the configuration described (disinflationary or stable goods prices, inflationary asset prices, contracting real economy, fiscal stress) the Fed has no policy that resolves all three simultaneously. Cutting rates relieves fiscal stress through lower debt service but fuels the asset inflation that is concentrating wealth and worsening the demand collapse. Holding rates or raising them maintains nominal price stability but accelerates the fiscal cascade through the debt service channel and intensifies the real economy contraction. Either choice degrades the configuration in a different direction.

This is the actual coupling mechanism. It is not that u, i, and r are correlated by some statistical regularity. It is that the Fed's policy reaction function, in this specific scenario, produces co-movement by construction: whatever the Fed does, two of the three variables move adversely.

Link 7: Either policy choice creates conditions favorable to a credibility event on Treasuries.

This is the contested link in the chain, and the section should be honest about that.

A credibility event on U.S. Treasuries is a sustained repricing of sovereign risk that causes foreign and domestic holders to demand significantly higher yields. The historical record contains few clean examples. The 2011 debt ceiling episode produced a brief downgrade and a modest yield response but no sustained repricing. The 2013 taper tantrum was a duration shock, not a credibility shock. The 2022 UK gilt crisis was the clearest recent example of a credibility event in a major sovereign, and it was triggered by a specific fiscal proposal rather than a structural cascade.

The chain does not automatically produce a credibility event from the policy trap. It produces *conditions favorable* to one. The trigger may be a specific fiscal proposal that the bond market interprets as unsustainable. It may be a foreign holder reducing holdings at a politically charged moment. It may be a debt ceiling episode that markets price as a real default risk rather than a political performance. The mechanism that connects the policy trap to the credibility event runs through expectations, not through direct economic causation.

There is a fifth path to a credibility event that has become more salient since 2024: the variance of policy itself across administrations. Sovereign creditworthiness rests on a durability premise: that government commitments are honored across time, that institutional rules do not change retroactively, that policy direction is bounded across electoral cycles. A market participant pricing a 30-year Treasury cannot know who will hold executive power for most of those 30 years. The wider the swing

between administrations on tax law, trade policy, regulatory regime, central bank independence, and treaty commitments, the wider the implied variance on every long-dated U.S. commitment. Even when each administration is internally consistent, the *whip-saw between them* damages the durability premise.

This effect is structural rather than partisan in origin. It reflects a broader trend in which executive action has expanded relative to legislative process across multiple administrations, increasing the variance of policy across electoral cycles. The 2017 Tax Cuts and Jobs Act, the 2021 American Rescue Plan, the 2022 Inflation Reduction Act, and the 2025 H.R. 1 package each represented a substantial recasting of fiscal policy with limited bipartisan input and partial reversal of the preceding administration's direction. The same pattern applies to trade policy, immigration enforcement, environmental regulation, and the relationship between the executive branch and the Federal Reserve. The credibility cost is not assigned to any single administration. It accumulates from the variance itself.

The clearest recent illustration is the IEEPA tariff episode of 2025-2026. Tariffs imposed under the International Emergency Economic Powers Act collected approximately $165 billion in revenue between February 2025 and January 2026, comprising approximately 52% of total U.S. customs duties at their peak.[20] On February 20, 2026, the Supreme Court ruled in *Learning Resources, Inc. v. Trump* that IEEPA does not authorize the president to impose tariffs; the tariffs were terminated four days later.[19] The Court of International Trade ordered refunds on March 4, 2026, and the potential refund liability is in the $100-200 billion range. The fiscal arithmetic moved by hundreds of billions of dollars based on a single court decision. The episode demonstrates the mechanism cleanly: an executive action of substantial fiscal magnitude was imposed under contested statutory authority, sustained for thirteen months, struck down on review, and produced a refund liability comparable in magnitude to a year of corporate income tax receipts. A market participant pricing a 30-year Treasury cannot price this kind of variance with confidence. The credibility cost is structural: what is being eroded is the ability to extrapolate any current fiscal configuration into the future.

Moody's downgrade of U.S. sovereign debt from Aaa to Aa1 on May 16, 2025 was the first such downgrade in the agency's 108-year history of rating U.S. debt. The explicit rationale included that "successive U.S. administrations and Congress have failed to agree on measures to reverse the trend of large annual fiscal deficits and growing interest costs." The framing is bipartisan and structural. Moody's also noted "diminished prospects" that the dollar's reserve role and Treasury market depth

would continue to offset widening deficits, language that directly mirrors the off-ramps argument in Section 4 of this paper. The downgrade is itself a partial credibility event: the term premium widened measurably, the 30-year yield traded above 5% in the days that followed, and the third of three major rating agencies stripped the U.S. of its top-tier status. A series of weak Treasury auctions in March 2026 reinforced the pattern: primary dealers absorbed approximately 24% of a 2-year note auction (roughly twice the share normally absorbed by these dealers), and 5- and 7-year securities also saw weak demand.[21] The full cascade Link 7 describes has not fired. The conditions favorable to it have measurably accumulated across multiple independent indicators, and the most conservative of the rating agencies has now acknowledged the accumulation in writing.[17]

The honest framing is that the policy trap *makes a credibility event more likely* by removing the standard reassurance mechanisms (the Fed cannot credibly commit to either supporting markets or maintaining price stability without sacrificing the other) but the trigger and timing are uncertain. This is the link in the chain where reasonable analysts most clearly disagree.

Link 8: The term premium rises, raising the average effective rate on federal debt.

If Link 7 fires, this follows mechanically. The term premium on U.S. Treasuries has been near zero or negative for fifteen years, well below historical averages of 150–200 basis points. Even a partial repricing toward historical norms produces large effects on debt service costs at current debt levels. The duration structure of Treasury debt means the repricing propagates through the average rate (the r variable in the model) over a multi-year window as new and rolled issuance prices at the new term premium.

Link 9: The cascade fires.

At this point the static model takes over. With u elevated by Link 1, the inflation channel contributing through Link 5 (which produces some inflation through asset and import channels even if goods inflation is muted), and r elevated by Link 8, the Crisis Ratio rises along the trajectory the scenarios in section 3 describe. The mechanism is no longer the chain; it is the arithmetic.

What is well-supported and what is contested

Of the nine links, six are robustly supported by existing empirical and theoretical work: Link 1 (cognitive exposure is concentrated in knowledge work), Link 2 (debt

and consumption are concentrated in the professional cohort), Link 3 (consumer demand responds to middle-income unemployment), Link 4 (capital has captured productivity gains historically and there is no clear reason for the pattern to reverse), Link 6 (the policy trap follows mechanically from the mixed price signal), and Link 8 (term premium repricing produces large rate effects).

Two links are well-supported in theory but uncertain in timing and magnitude: Link 5 (the mixed inflation configuration is one of several possible outcomes, depending on the relative balance of demand destruction and other inflationary forces) and Link 9 (which is mechanical given the inputs but depends on the prior links firing).

One link is the central contested step: Link 7, the credibility event. The chain does not generate a credibility event mechanically. It generates conditions favorable to one. The trigger and timing are uncertain and depend on factors outside the model.

The honest summary is that the cascade is conditional on a chain of nine links, six of which are well-supported, two of which are plausible but not certain, and one of which depends on conditions favorable rather than mechanical causation. This is a defensible structural argument. It is not a forecast.

The Crisis Ratio measures consequence, not cause

The mechanism described in this section is not part of the static model. The static model is a measurement instrument that responds to whatever values of u, i, and r are presented to it. The mechanism explains *why* those values might co-move toward stress configurations under a specific scenario.

The two arguments (the static model and the mechanism) are intentionally separable. A reader who accepts the diagnostic framework but rejects the AI mechanism can still use the model to evaluate other scenarios that produce co-movement in u, i, and r through different causes. A reader who accepts the AI mechanism but disputes the calibration of the static model can engage with the chain on its own terms. The contribution of this paper is to offer both arguments and to be explicit about which one rests on what.

The cascade is not predicted. The mechanism that would produce it is named. The arithmetic, conditional on the mechanism, is mechanical. What is contested is whether the mechanism fires, when it fires, and how completely. The Crisis Ratio is the instrument that measures what happens if it does.

3.9 Dynamic Projection: What Happens When Debt Grows

The scenarios presented in section 3 hold debt held by the public constant at $28.3 trillion. This is a pedagogical simplification that isolates the effects of changes in u, i, and r on the ratio. It is also the largest single understatement in the static model. In any real stress scenario, debt does not stay constant. It grows with the primary deficit, which itself widens under unemployment. This section makes the dynamic explicit.

The growth of debt under FY2025 conditions

Federal debt held by the public grew by $1.97 trillion in FY2025, from $28.31 trillion at end-FY2024 to $30.28 trillion at end-FY2025.[11] Of this growth, approximately $1.03 trillion was net interest (debt refinancing itself into more debt) and the remaining $940 billion was the primary deficit. The total federal deficit for FY2025 was $1.83 trillion.[12]

Looking forward under baseline conditions, the CBO projects the structural gap between federal outlays (23.5% of GDP) and revenue (17% of GDP) at approximately 6.5% of GDP.[13] At current nominal GDP of approximately $29 trillion, this implies a structural deficit of roughly $2.0 trillion per year before any cyclical adjustment. The static specification of the Crisis Ratio therefore understates the cascade by treating as constant a variable that grows at roughly $2 trillion per year in the absence of any stress, and substantially more under stress.

The growth function

A dynamic projection requires a simple rule for how D grows. The model uses:

$$\Delta D/\text{year} = \$2.0T + \$0.2T \times \max(0,\ u - 4.0)$$

The first term reflects the structural deficit at baseline conditions. The second term reflects the cyclical addition: each percentage point of unemployment above the natural rate (here treated as 4%) widens the primary deficit by approximately $200 billion through automatic stabilizers (unemployment insurance, SNAP, Medicaid expansion) and modest discretionary response. This calibration is consistent with the 2008–2009 experience, when the deficit rose from $458 billion to $1.41 trillion as unemployment increased approximately 5 percentage points, implying a cyclical sensitivity in this range net of stimulus effects.[14]

The rule is intentionally simple. It does not endogenize tax law changes, supplemental appropriations, or fiscal stimulus response. It assumes the primary deficit responds to unemployment alone with a constant elasticity. These

simplifications are conservative: actual deficits in past stress episodes have widened by more than this rule predicts, not less.

The dynamic scenarios

Applying the growth rule to each scenario for one and three years produces the following:

Scenario	u	ΔD/year	D (now)	D (+1yr)	D (+3yr)	Static R	R at +3yr
Current	4.0	$2.0T	$28.3T	$30.3T	$34.3T	**0.213**	**0.258**
Mild	6.5	$2.5T	$28.3T	$30.8T	$35.8T	**0.313**	**0.396**
Moderate	8.0	$2.8T	$28.3T	$31.1T	$36.7T	**0.408**	**0.529**
Severe	10.0	$3.2T	$28.3T	$31.5T	$37.9T	**0.519**	**0.694**
Cascade	12.0	$3.6T	$28.3T	$31.9T	$39.1T	**0.699**	**0.966**

The static model shows the cascade scenario reaching 0.699. The dynamic model shows the same scenario reaching 0.966 within three years if conditions persist. That is the difference between debt service consuming seven-tenths of revenue and debt service consuming nearly all of revenue. The same input configuration produces a qualitatively different conclusion when the static assumption is relaxed.

The mild scenario produces a similar pattern at lower magnitude. A moderate cyclical shock that produces a static ratio of 0.313 produces a three-year dynamic ratio of 0.396, comfortably above the 0.30 threshold and approaching 0.40. The static reading suggests "approaching threshold." The dynamic reading suggests "above threshold and accelerating." Both are accurate descriptions of the same configuration. They differ in what they assume about persistence.

What the dynamic projection adds

The dynamic projection adds three things to the static analysis.

First, it makes the timeline explicit. Static scenarios describe *configurations*, not *trajectories*. A reader looking at the static cascade scenario can reasonably ask: but how fast does this actually happen? The dynamic table answers that question. Under sustained extreme stress, the system moves from "debt service is seven-tenths of revenue" to "debt service is essentially all of revenue" in three years. The cascade has a clock.

Second, it shows that the current trajectory is not benign even without a stress shock. The current state row of the dynamic table shows the ratio rising from 0.213 to 0.258 over three years at u=4%, with debt accumulating at the structural baseline rate. This is the CBO baseline trajectory expressed in the Crisis Ratio framework. Even with full employment, target inflation, and no credibility event, the ratio drifts toward the 0.30 threshold under the structural deficit alone. The cascade scenarios describe what happens when this drift meets a shock.

Third, it corrects what would otherwise be a conservative bias in the static specification. Holding D constant in stress scenarios systematically understates the cascade. The dynamic projection makes that understatement explicit and quantified.

The countervailing dynamic

The dynamic projection above holds baseline revenue Y_0 constant while debt grows. This is also a simplification, and it cuts in the opposite direction. Under favorable conditions, Y_0 itself grows with nominal GDP. Federal receipts grew from $4.92 trillion in FY2024 to $5.24 trillion in FY2025, an increase of $317 billion driven primarily by customs duties (+$118 billion) and individual income tax collections (+ $230 billion).[15] Net interest grew by only $80 billion over the same period. The actual Crisis Ratio at end-FY2025 was therefore 0.198, not 0.213, because revenue growth outpaced interest growth. Through the first seven months of FY2026 (October 2025 to April 2026), receipts continued to outpace interest costs: receipts ran $200 billion higher than the same period of FY2025, while net interest rose $41 billion. Annualized, the FY2026 in-progress Crisis Ratio computes to approximately 0.190.[18] The favorable dynamic has now held for three consecutive fiscal years.

A material caveat applies to the FY2025 customs duty growth, however. On February 20, 2026, the Supreme Court ruled in *Learning Resources, Inc. v. Trump* that the International Emergency Economic Powers Act does not authorize the president to impose tariffs.[19] All IEEPA-based tariffs were terminated on February 24, 2026. Penn Wharton Budget Model estimates that IEEPA tariffs collected approximately $165 billion cumulatively from February 2025 through January 2026 and had reached approximately 52% of total customs duties by the time of the ruling.[20] The Court of International Trade ordered nationwide refunds on March 4, 2026; the government has appealed, and the refund mechanism is unresolved, but the potential refund liability is in the $100-200 billion range. A meaningful portion of the FY2025 customs duty growth that drove the ratio improvement from 0.213 to 0.198 was therefore subsequently ruled to exceed statutory authority. Tariff revenue from Section 301 (China), Section 232 (steel, aluminum, autos), and other statutory

authorities remains in place and continues to support FY2026 receipts, but the dominant IEEPA piece is gone.

The structural argument of this paper is reinforced rather than weakened by this development. The paper argues that the favorable revenue dynamic operating in FY2024 and FY2025 is conditional and fragile. The IEEPA episode is a concrete example of that fragility: $165 billion in revenue, generated through an executive action that exceeded constitutional limits, supporting a ratio decline that the model otherwise treats as evidence of sustainability. When the action was struck down, the revenue stream stopped, the refund liability appeared, and the favorable dynamic became contingent on whatever new revenue sources the political system could substitute. The favorable dynamic continues to hold through FY2026 because nominal growth and remaining statutory tariffs have so far filled the gap. The structural fragility the paper describes is exactly the kind of fragility the IEEPA episode illustrates: favorable revenue conditions sustained by political choices that may not survive judicial review, political turnover, or external shock.

This is the r-vs-g dynamic Blanchard identified.[8] Under favorable conditions, growth in receipts mitigates the cascade. Under unfavorable conditions, the same dynamic reverses: receipts contract while debt accumulates at accelerated rates, and the ratio compounds adversely. The static model captures neither effect. The dynamic projection captures the unfavorable side. A complete picture would capture both.

For the purpose of this paper, the asymmetry is informative rather than disqualifying. The static specification is conservative in benign environments (Y_0 grows, ratio improves) and anti-conservative in stress environments (D grows, Y_0 contracts, ratio compounds). The dynamic projection corrects the anti-conservative bias in the direction that matters for evaluating cascade risk. A reader who wants the symmetric treatment can mentally subtract baseline revenue growth from the dynamic ratios to recover the favorable-conditions trajectory. The cascade scenarios are not improved by this correction; they are at unemployment levels where receipts contract regardless of secular growth.

Limitations of the dynamic projection

The dynamic projection inherits all the limitations of the static model and adds several of its own. It assumes a single growth rule for D that does not respond to monetary policy, fiscal stimulus, or political constraints on debt issuance. It holds the inflation channel and the interest rate channel constant across the projection

horizon, when both would be expected to evolve in a real stress scenario. It does not endogenize Federal Reserve response. It does not model rollover risk, in which debt service can rise faster than the smooth r-update the model assumes.

These limitations are real and a complete dynamic specification would address them. The purpose of this section is narrower: to show that relaxing the most conservative simplification of the static model (holding D constant) produces ratio trajectories that are substantially worse than the static scenarios suggest, on a timeline that is policy-relevant rather than indefinite. The static model understates the cascade. The dynamic model corrects the understatement at a single dimension. A complete dynamic model would compound the correction across additional dimensions, all of which would push the trajectory in the same direction.

4. WHY THE OFF-RAMPS ARE CLOSED

The standard rebuttal to fiscal cascade models is that real economies have safety valves the model leaves out. The list usually runs: the U.S. can print money through Federal Reserve intervention; the dollar's reserve status guarantees a deep buyer base for Treasuries; in a real crisis, governments raise taxes and broaden bases; and so on.

Each of these is real in the abstract. None of them functions as an off-ramp in this specific cascade. The reasons are worth stating directly.

4.1 Quantitative easing does not break the loop; it reroutes it

The Federal Reserve can purchase Treasury debt to suppress yields. This is the standard response to fiscal stress and it has worked in past episodes.

The mechanism by which QE works, however, is inflationary. Expanding the Fed balance sheet to buy debt injects monetary base into the economy. In an environment where inflation is already above target (which is the configuration the cascade assumes) additional monetary expansion accelerates the inflation channel *already eroding revenue in the model's denominator*. The β coefficient is doing this work. QE does not exit the cascade. It trades a faster numerator collapse for a faster denominator collapse, and the net effect on the ratio depends on which channel moves harder.

The historical precedent often cited (the 2009–2022 QE programs) does not refute this. That episode featured QE into a *deflationary* environment, where the inflation channel had headroom to absorb monetary expansion before it bound. The current configuration does not have that headroom. QE during disinflation is a different policy than QE during persistent above-target inflation, and treating them as the same tool is the error.

A more sophisticated version of the monetary objection invokes yield curve control rather than QE: the central bank can simply cap yields at a chosen level, as the Federal Reserve did in the 1940s and the Bank of Japan did from 2016, and hold them there regardless of market demand. If the Fed caps the 10-year yield, the objection goes, then r in the model cannot rise and the cascade cannot run through the interest channel.

Yield curve control does not escape the cascade; it relocates it. Capping a yield below the level the market demands requires the central bank to buy whatever quantity of debt the market will not absorb at the capped rate. In a credibility event, that quantity is unbounded, and the purchases are made with newly created money in an already-inflationary environment, which accelerates the same denominator collapse described above. More fundamentally, for the reserve currency issuer, a yield cap that holds nominal rates below market-clearing levels drives capital out of the currency. The cascade then runs through the exchange rate and the reserve status rather than through the nominal yield. Japan could sustain yield curve control in part because the yen is not the global reserve asset and because Japanese debt is domestically financed (Section 4.5). For the dollar, yield curve control trades a Treasury-market cascade for a currency-and-reserve-status cascade, which is a different failure mode and arguably a worse one. The tool does not remove the trap. It changes which wall the system hits.

4.2 Reserve currency status is the variable, not the constant

The model is sometimes criticized for not accounting for the dollar's reserve currency status, which generates structural demand for Treasuries from foreign central banks and provides a buffer no other sovereign enjoys.

This critique misunderstands what the cascade describes. The Crisis Ratio is a model of the conditions under which reserve currency status comes into question. Foreign central banks hold dollars and Treasuries because the dollar is the most credible store of value among large reserve assets. Credibility is not a fixed property of the dollar; it is a function of perceived U.S. fiscal sustainability. When debt service consumes a large and rising share of revenue, the credibility premium that supports reserve status erodes.

Treating reserve status as a buffer the model omits is therefore circular. It assumes the conclusion the cascade contradicts. Reserve currency status is among the things that fail when the ratio climbs far enough, not a buffer that holds against it.

4.3 The tax base is contracting, not available for expansion

The standard fiscal recovery move is to raise taxes, broaden the base, and impose windfall levies. The model holds baseline revenue Y_0 constant; this is sometimes flagged as a conservative assumption, on the theory that a government in crisis would expand Y_0 through new revenue.

The empirical record points the other direction.

The trajectory that produced $28.3 trillion in publicly held federal debt is not a story of inadequate political will to spend. It is in significant part a story of revenue-side decisions: the 2001 and 2003 Bush cuts, the 2017 Tax Cuts and Jobs Act, and the 2025 extensions of TCJA provisions. The political system has spent two generations narrowing the lever that fiscal recovery would require. Calling it an unused buffer treats as available a tool the political system has actively constrained.

Windfall levy proposals fail the same test. The 2022 oil price episode produced record corporate margins and bipartisan calls for windfall taxation. None passed. If the political system cannot impose a windfall tax on the most visible windfall in a generation, the model should not assume the tool is operationally available.

The active direction is the other way. Mass deportation, currently underway as policy, directly reduces Y_0. Undocumented workers paid $96.7 billion in federal, state, and local taxes in 2022, of which $59.4 billion went to the federal government. [7] Removal of these workers is a revenue contraction *on top of* the unemployment channel the model already prices. The "buffer not included" framing has the sign wrong.

4.4 The fragilities, by contrast, are real and run the other direction

The model excludes several mechanisms that would make the cascade worse, faster, than it shows:

- Rollover risk operates on a faster timeline than smooth changes in r. Roughly one-third of U.S. Treasury debt matures and rolls each year. A buyer strike at auction can force yields up in days, not quarters.
- Foreign holder behavior is endogenous to the very crisis the model describes. Major foreign holders will sell when others sell, not when the model gradually predicts they would.
- Credit spread blowouts widen the gap between policy rates and average portfolio yields. The model uses a single r; in an actual cascade, long-duration debt service can detach from short-end policy.
- Political constraints (debt ceiling fights, statutory caps, technical default risk) are not modeled. These are not buffers; they are additional points of fracture.

The pattern across the four conventional escape valves is the same. Each is presented in standard fiscal discourse as something the model omits to its disservice.

In fact, each is either endogenous to the cascade (QE, reserve status), actively contracting in the opposite direction from its supposed buffer function (tax base), or politically constrained (revenue expansion). The genuine fragilities the model omits (rollover, foreign holders, credit spreads, political process) all push the cascade in the worse direction.

The honest statement is therefore: the Crisis Ratio is conservative. The real system is more locked in than the model shows, not less.

4.5 Why not Japan?

The most common single objection to any U.S. fiscal alarm is Japan. Japan has run gross debt above 250% of GDP for years, sustained near-zero interest rates for decades, and has not experienced the cascade this paper describes. If high debt mechanically produced the cascade, the objection goes, Japan would have collapsed long ago. The Japanese counterexample deserves a direct answer, because answering it clarifies what the cascade actually requires.

Japan does not cascade because the specific conditions that drive the U.S. cascade are largely absent there. Four differences matter.

First, Japanese government debt is overwhelmingly held domestically. The Bank of Japan and Japanese institutions and households own the large majority of it, with foreign holders a small minority. The Link 7 credibility mechanism in this paper runs substantially through foreign holder behavior and the term premium foreign capital demands. A sovereign whose debt is held by its own central bank and its own savers, who buy for reasons that are not purely return-driven, is insulated from the foreign-capital credibility dynamics that drive the U.S. cascade. The U.S. relies on foreign holders for a large share of its marketable debt, which is precisely the exposure Japan lacks.

Second, Japan ran its high debt alongside a persistently high domestic savings rate and sustained current account surpluses. The debt was financed by a domestic saving pool large enough to absorb it at low yields. The U.S. has a low domestic savings rate and runs persistent current account deficits, which means it depends on foreign capital inflows to finance its debt. The structural capacity to self-finance that protected Japan is not present in the U.S. case.

Third, Japan is not the global reserve currency issuer. This cuts in a direction that surprises people. Reserve status is often cited as a U.S. advantage, and in normal times it is, but it also means the consequences of a U.S. credibility event propagate

globally and reflexively in a way Japan’s never could. Japan could run an idiosyncratic high-debt experiment precisely because the yen is not the asset the entire global financial system treats as risk-free. The dollar’s reserve role is a larger buffer in calm conditions and a larger fracture surface in a crisis.

Fourth, and most important for this paper, Japan never faced the specific shock the cascade requires. Japan’s high-debt era coincided with deflation or near-zero inflation, an aging and shrinking workforce that suppressed wage pressure through demographics rather than displacement, and no episode of simultaneous cognitive-tier automation against a high-debt, low-savings, foreign-financed balance sheet. The cascade this paper describes is not "high debt causes collapse." It is "a specific shock drives unemployment, inflation, and rates together in a balance sheet that lacks the domestic-financing and savings buffers that would otherwise absorb it." Japan had the buffers and lacked the shock. The U.S. lacks the buffers and faces the candidate shock.

Japan is therefore not a counterexample to the cascade. It is a description of the conditions under which high debt is survivable: domestic financing, high savings, no coupling shock. The U.S. configuration is the photographic negative of the Japanese one. The lesson of Japan is not that high debt is safe. It is that high debt is safe under conditions the U.S. does not meet.

5. LITERATURE AND LIMITATIONS

5.1 The Crisis Ratio in context

The model presented here is a teaching simplification of the established sovereign debt sustainability framework. The technical literature is extensive and the major contributions deserve acknowledgment.

Blanchard's r vs g framework. Olivier Blanchard's 2019 AEA presidential address argued that when the real interest rate on government debt (r) is below the real growth rate of the economy (g), public debt may be sustainable even at elevated levels because the debt-to-GDP ratio falls mechanically over time.[8] The Crisis Ratio model does not contradict this framework; it describes the conditions under which r-vs-g favorability fails. When inflation forces nominal rates higher and unemployment depresses growth, the favorable r-vs-g configuration that supported the 2010s low-rate environment reverses.

Reinhart and Rogoff. The Reinhart-Rogoff threshold work (debt-to-GDP above 90% correlated with slower growth) has been contested in detail but the underlying observation is robust: high debt levels constrain fiscal flexibility and increase vulnerability to interest rate shocks.[9] The Crisis Ratio expresses the same mechanism through the debt-service-to-revenue channel rather than the debt-to-GDP channel, which is more directly observable in policy-relevant timeframes.

The CBO long-term outlook. CBO projects net interest reaching 4.1% of GDP by 2035, consuming approximately one-sixth of all federal spending under current-law baseline assumptions.[10] The Crisis Ratio model does not extend CBO's projections; it makes them legible by expressing the trajectory as a single ratio whose evolution depends on three intuitive variables.

IMF and BIS debt sustainability frameworks. The IMF's Debt Sustainability Analysis (DSA) framework and the BIS work on sovereign debt vulnerabilities are the technical standards in this space. They use richer multi-variable specifications, scenario stress testing, and country-specific parameter calibration. The Crisis Ratio model trades that richness for accessibility; readers seeking the technical version should consult those frameworks directly.

5.2 Coefficient sensitivity

The revenue function uses two sensitivity coefficients: α (the fractional reduction in receipts per percentage point of unemployment) set to 0.05, and β (the fractional reduction per percentage point of excess inflation) set to 0.04. These are stylized values rather than structural estimates, and a reasonable objection is that the model's output depends on coefficient choices the paper asserts rather than derives. This section addresses that objection directly.

An illustrative regression on historical annual data (federal receipts, unemployment, CPI inflation, and nominal GDP from 1990 to 2024) provides an empirical reference range. Regressing the receipts-to-GDP ratio on unemployment and excess inflation yields an unemployment coefficient implying α of approximately 0.029, statistically significant at well below the 1% level. A log-receipts specification controlling for log GDP yields an unemployment semi-elasticity of approximately 0.034. Both historical estimates are therefore in the 0.03 range, somewhat below the 0.05 used in the model.[23] The GDP elasticity of receipts in the same regression is 0.94, consistent with the standard CBO finding of revenue elasticity near unity, which serves as a sanity check on the method.

Two honest observations follow. First, the model's α of 0.05 is above the historical estimate. This is deliberate and consistent with the paper's central thesis: the AI displacement scenario is argued to produce larger revenue contraction than historical cyclical episodes, because the displaced cohort is high-wage and the corporate-tax channel does not provide its usual offset (Section 3.8, Link 4). An α at or above the historical high end is the appropriate parameterization for a shock the paper argues is worse than its historical analogues, not a claim that 0.05 is the normal cyclical sensitivity. Second, the inflation coefficient β is the weakest-identified parameter in the model. In the historical regression, excess inflation does not reliably reduce receipts and may slightly increase nominal receipts through bracket creep before indexing adjusts. The inflation channel is retained because the AI scenario's asset-inflation and supply-side dynamics differ from the historical sample, but its empirical support is thin and the paper does not lean on it.

The decisive question is whether the cascade conclusion survives across the plausible coefficient range. It does. The following table computes each scenario's static ratio under four coefficient sets: the paper's values, the conservative historical estimates (α = 0.03, β = 0), a low-inflation-sensitivity variant, and a high-sensitivity variant.

Scenario	Paper a.05 b.04	Hist. a.03 b0	Low a.03 b.02	High a.07 b.05
Current	0.213	0.201	0.203	0.223
Mild	0.313	0.284	0.288	0.340
Moderate	0.408	0.355	0.364	0.457
Severe	0.519	0.432	0.446	0.605
Cascade	0.699	0.549	0.574	0.837

The threshold is 0.30. Under the paper's coefficients, the mild scenario and everything more severe breach it. Under the most conservative historical coefficients, the moderate scenario and everything more severe still breach it, and the cascade scenario still reaches 0.55, meaning debt service consuming more than half of available revenue. The coefficient choice affects how quickly the threshold is breached, not whether. The cascade conclusion is robust to the full plausible range of revenue sensitivities. This is the property a diagnostic of this kind most needs: the qualitative result does not depend on a knife-edge calibration.

5.3 What the model is not

The Crisis Ratio is a simplified diagnostic. It is not a forecast. It is not a substitute for the cyclically-adjusted primary balance, the present value of future deficits, or any of the standard technical measures. Its purpose is to make the cascade structure visible to readers who would not otherwise engage with the technical literature. Several limitations should be named explicitly.

- The model is linear in u and i (with bounded specification), but real revenue elasticities likely vary with the level of stress. Behavior near the floor differs from behavior at moderate stress.
- The model uses a single weighted average r, which abstracts over a heterogeneous maturity structure. Real debt service evolves with the rollover schedule, not smoothly with the policy rate.
- The static specification holds the debt stock D constant within scenarios. Section 3.9 relaxes this assumption with a simple growth rule for D as a function of unemployment, but does not endogenize the primary deficit structure, monetary policy response, or political constraints on debt issuance.

- The model does not endogenize Federal Reserve policy. In practice, monetary response to fiscal stress is itself a strategic variable, with the choices and constraints described in section 4.
- The model does not incorporate growth (g) directly. A complete sustainability analysis requires the r-vs-g comparison; the Crisis Ratio is an indirect measure that implicitly depends on growth assumptions embedded in Y_0.

Despite these limitations, the model has a property that more sophisticated frameworks often lack: it is legible. A serious reader can verify the calibration, change a variable, and see the response. The relationship between the three inputs and the output is transparent. For the purpose of making fiscal cascade risk available to readers outside the specialist community, this transparency is the model's primary value.

5.4 What follows from the diagnosis

This paper concludes without policy recommendation. The diagnostic (that the cascade is structurally locked in under plausible shock paths, that the standard off-ramps do not function as off-ramps in this configuration, and that the timeline is compressed by variable coupling) leaves open a range of possible responses. The categories include monetary regime change, structural reform of healthcare and retirement security (which currently bind labor mobility to specific employment relationships), industrial policy aimed at maintaining the consumer base, sovereign wealth-style citizen dividends, debt restructuring, and combinations of the above. Each of these is contestable. Each has serious proponents and serious critics in the economics literature.

The argument for any specific response is outside the scope of this paper. The intent here is to make the diagnostic shared. A serious response to fiscal cascade risk requires agreement on the structure of the problem before debate over the structure of the response. That agreement has been elusive in part because the technical literature is inaccessible to most policy-relevant readers and the political literature does not engage the technical case. The Crisis Ratio is offered as a bridge between them.

6. CONCLUSION

The Crisis Ratio is a simple tool. Federal debt service divided by available federal revenue, with revenue calculated as a function of unemployment, inflation, and a target inflation rate. Three intuitive inputs, two structural parameters, one scalar output.

Calibrated against actual FY2024 data, the current static ratio sits at 0.213, approaching but not yet at the 0.30 threshold where debt service begins to consume the federal discretionary budget. The actual end-FY2025 ratio is 0.198, lower than the FY2024 figure, because favorable revenue growth from tariffs and individual income tax outpaced the rise in interest costs. CBO projections suggest the ratio drifts upward under baseline assumptions even with full employment. A mild cyclical shock combined with persistent above-target inflation pushes the static ratio above the 0.30 threshold. The dynamic projection in Section 3.9, which relaxes the constant-debt assumption, shows that the same configurations produce substantially worse ratios when carried forward three years: the mild scenario reaches 0.396, the moderate scenario reaches 0.529, and the cascade scenario approaches 0.97. The three variables are ordinarily coupled by canonical macroeconomic mechanisms that move them in offsetting directions; the chain described in Section 3.8 argues that AI cognitive displacement is a plausible mechanism that overrides the canonical coupling and drives the variables together. Conditional on the override mechanism firing and persisting, the ratio rises along trajectories the dynamic projection makes legible.

The standard policy off-ramps (quantitative easing, reserve currency demand, tax base expansion) do not function as escape valves in this configuration. Each is either endogenous to the cascade, actively shrinking in the opposite direction from its supposed buffer function, or politically constrained. The genuine model omissions (rollover risk, foreign holder behavior, credit spread dynamics, political process) all push the cascade in the worse direction.

The model does not forecast. It describes a structural relationship. What it shows is that the conditions under which fiscal cascade becomes operative are closer than separate consideration of any single variable would suggest, and that the off-ramps assumed in standard fiscal discourse are not available in this configuration. The implications for response are left to readers, analysts, and policymakers who engage with the diagnostic.

The arithmetic is the model. No other assumptions are introduced. The conclusions follow.

NOTES AND SOURCES

[1] FY2024 data. Total federal receipts of $4.92 trillion: Congressional Budget Office, *Monthly Budget Review: Summary for Fiscal Year 2024* (November 2024). Net interest outlays of $949 billion (FY2024); surpassed $1 trillion in FY2025: CBO, *Monthly Budget Review: Summary for Fiscal Year 2025* (November 2025). Federal debt held by the public of approximately $28.3 trillion at end of FY2024: U.S. Department of the Treasury, Monthly Treasury Statement. Weighted average effective interest rate on debt held by the public of approximately 3.3% derived from these figures.

[2] Revenue elasticity literature. The unemployment revenue elasticity is consistent with CBO's cyclically-adjusted revenue methodology and with the empirical work of Auerbach and Feenberg, "The Significance of Federal Taxes as Automatic Stabilizers," *Journal of Economic Perspectives* (2000). The inflation elasticity on real receipts is calibrated to be consistent with post-1985 indexed-bracket experience.

[3] OECD threshold experience. The 0.30 threshold reflects observed fiscal stress episodes in OECD economies where debt service began to crowd out discretionary spending. See OECD Fiscal Monitor and IMF Debt Sustainability Analysis country reports for empirical reference points.

[4] On α calibration. The 2008–2009 episode saw federal receipts fall from $2.52T (FY2008) to $2.10T (FY2009), approximately 17%, against a peak unemployment increase of 4.3 percentage points. Net of stimulus-related effects, this implies an underlying cyclical revenue elasticity in the 0.04–0.05 range. The model uses 0.05 as a midpoint.

[5] Threshold interpretation. At 0.30, debt service approximately equals the federal share of education, transportation, and most non-defense discretionary spending. At 0.50, debt service exceeds combined defense and education. At 1.0, debt service equals total federal revenue and the government must borrow simply to pay interest on existing borrowing. These are structural thresholds, not regulatory ones; they reflect the composition of the U.S. federal budget rather than any specific statutory limit.

[6] CBO projections. CBO, *The Budget and Economic Outlook: 2025 to 2035* (January 2025), projects net interest reaching 4.1% of GDP by 2035 and accounting for about one-sixth of all federal spending.

[7] Undocumented worker tax contributions. Institute on Taxation and Economic Policy, "Tax Payments by Undocumented Immigrants" (July 2024). The $96.7 billion figure represents total federal, state, and local tax contributions in 2022; the $59.4 billion figure represents the federal portion.

[8] Blanchard r-vs-g framework. Olivier J. Blanchard, "Public Debt and Low Interest Rates," *American Economic Review* 109(4): 1197–1229 (2019). The presidential address argued that when r < g, debt may be sustainable even at elevated levels because the debt-to-GDP ratio falls over time absent new borrowing.

[9] Reinhart and Rogoff. Carmen M. Reinhart and Kenneth S. Rogoff, "Growth in a Time of Debt," *American Economic Review* 100(2): 573–578 (2010). The specific 90% threshold has been contested (notably by Herndon, Ash, and Pollin) but the underlying observation that elevated debt constrains fiscal flexibility is robust across the broader literature.

[10] CBO long-term outlook. CBO, *The Budget and Economic Outlook: 2025 to 2035* (January 2025); CBO, *The 2024 Long-Term Budget Outlook* (March 2024).

[11] FY2025 debt growth. Joint Economic Committee, U.S. Senate, "FY2025 Debt Increased by $2.2 Trillion, Stands at Over $37.6 Trillion" (October 2025). Debt held by the public grew $1.97 trillion in FY2025, from $28.31 trillion to $30.28 trillion. Total debt growth (public plus intragovernmental) was $2.17 trillion. The model uses debt held by the public throughout because intragovernmental holdings do not generate market interest costs and are not credibility-relevant.

[12] FY2025 deficit. Bipartisan Policy Center, "Deficit Tracker" (November 2025), drawing on CBO and Treasury Department data. FY2025 federal deficit of $1.83 trillion comprises approximately $1.03 trillion in net interest and $0.80 trillion in primary deficit. The growth of debt held by the public in FY2025 ($1.97 trillion) exceeds the total deficit slightly due to Treasury cash management and operating cash balance changes.

[13] CBO baseline gap. Congressional Research Service, "Federal Debt and the Debt Limit in 2025" (September 2025). CBO current-law baseline projects FY2025 outlays at 23.5% of GDP and revenues at 17.0% of GDP, implying a 6.5% of GDP structural gap. At nominal GDP of approximately $29 trillion, this corresponds to a structural deficit of approximately $1.9 trillion per year before cyclical adjustment. The model uses $2.0 trillion as the structural baseline for the growth rule.

[14] Cyclical deficit sensitivity. The 2008–2009 episode saw the federal deficit rise from $458 billion (FY2008) to $1.413 trillion (FY2009) as unemployment increased from 5.0% (FY2008 average) to 9.3% (FY2009 average), a peak rise of approximately 5 percentage points. The deficit increase of $955 billion against a 5 point unemployment increase implies a gross cyclical sensitivity of approximately $190 billion per percentage point of unemployment, net of stimulus-related effects. The model uses $200 billion as the cyclical addition per percentage point of unemployment above 4%.

[15] FY2024 to FY2025 revenue growth. Bipartisan Policy Center, "Deficit Tracker" (November 2025). Federal receipts rose from $4.92 trillion (FY2024) to $5.24 trillion (FY2025), an increase of $317 billion (6%). Individual income tax collections rose by $230 billion (9%), customs duties rose by $118 billion (153%) following increased tariffs from February 2025 onward, payroll taxes rose by $39 billion (2%), and corporate income tax fell by $78 billion (-15%) due to allowances for accelerated investment deductions under H.R. 1. The actual end-FY2025 Crisis Ratio computed against these receipts and a $1.03 trillion interest charge is 0.198.

[16] Effective corporate tax rates. The federal statutory corporate tax rate is 21% (post-TCJA). The effective rate paid by large U.S. multinationals on global book income has run substantially lower, with most large-cap technology companies reporting effective rates in the 9–13% range and some reporting near-zero rates in specific years through legal mechanisms including the Foreign-Derived Intangible Income (FDII) deduction, the Global Intangible Low-Taxed Income (GILTI) regime, accelerated depreciation, stock-based compensation deductions, R&D credits, and foreign tax credit utilization. The Inflation Reduction Act of 2022 introduced a 15% corporate alternative minimum tax (CAMT) on adjusted financial statement income of corporations with average annual book income exceeding $1 billion; subsequent Treasury and CBO scoring revised CAMT revenue projections downward as guidance and exemptions narrowed the effective base. H.R. 1 (2025) expanded several accelerated-deduction mechanisms, contributing to the $78 billion FY2025 decline in corporate income tax receipts noted in [15]. Sources: Joint Committee on Taxation reports on CAMT; Institute on Taxation and Economic Policy, "Corporate Tax Avoidance in the First Five Years of the Trump Tax Law" (2024); Tax Foundation analyses of effective corporate rates.

[17] Moody’s downgrade. Moody’s Ratings, "Moody’s downgrades the United States of America’s ratings to Aa1 from Aaa; changes outlook to stable" (May 16, 2025). The downgrade was the first by Moody’s in its 108-year history of rating U.S.

sovereign debt. The agency cited fiscal trajectory, projected interest payments rising from 9% of federal revenue to 30% by 2035, and "diminished prospects" that reserve currency status would continue to offset widening deficits. The 30-year Treasury yield rose above 5% and the 10-year above 4.5% in the days following. Standard & Poor's downgraded U.S. sovereign debt from AAA to AA+ in August 2011; Fitch downgraded from AAA to AA+ in August 2023; the Moody's 2025 action made the U.S. downgrade unanimous across the three major rating agencies. Market reaction analysis from Western Asset Management, "End of an Era: Moody's Downgrades US to Aa1" (May 19, 2025); macro context from CSIS, "Moody's Downgrade Signals Deeper Risk" (May 20, 2025).

[18] FY2026 in-progress data. Congressional Budget Office, *Monthly Budget Review for April 2026* (May 8, 2026); Bipartisan Policy Center, "Deficit Tracker" (May 2026); Peter G. Peterson Foundation, "The Current Federal Deficit and Debt" (May 2026). Through the first seven months of FY2026 (October 2025 through April 2026), federal receipts totaled approximately $3.3 trillion (up $200 billion from the same period in FY2025) and net interest outlays totaled $628 billion (up $41 billion, or 7%, over the same period). Annualized to a full fiscal year: receipts ~$5.66 trillion, net interest ~$1.08 trillion, implied Crisis Ratio approximately 0.190. The FY2026 deficit is projected at $1.9 trillion (5.8% of GDP) per CBO's January 2026 baseline. Federal debt held by the public reached $31.26 trillion in April 2026, with total public debt outstanding (debt held by public plus intragovernmental holdings) at $38.92 trillion on pace to cross $39 trillion in May 2026.

[19] Supreme Court ruling on IEEPA tariffs. Supreme Court of the United States, *Learning Resources, Inc. v. Trump*, No. 24-1287 (February 20, 2026). In a 6-3 decision by Chief Justice Roberts (joined in full by Justices Gorsuch and Barrett, and in part by Justices Sotomayor, Kagan, and Jackson), the Court held that the International Emergency Economic Powers Act does not authorize the president to impose tariffs. The Court emphasized that the Constitution assigns to Congress the authority to "lay and collect Taxes, Duties, Imposts and Excises" and applied separation-of-powers principles in determining that tariff authority of significant economic and political magnitude requires explicit statutory delegation. Justices Thomas, Alito, and Kavanaugh dissented. All IEEPA-based tariffs were terminated on February 24, 2026. The Court did not order refunds; the U.S. Court of International Trade subsequently issued a nationwide refund order on March 4, 2026, which the government has appealed. The Supreme Court's decision became final on March 17, 2026. See also Holland & Knight, "Supreme Court Strikes Down IEEPA Tariffs" (February 20, 2026); Congressional Research Service Legal Sidebar LSB11398

(February 23, 2026); Holland & Knight, "Court of International Trade Orders Nationwide Tariff Refunds" (March 6, 2026).

[20] IEEPA tariff revenue and refund estimates. Penn Wharton Budget Model, "Supreme Court Tariff Ruling: IEEPA Revenue and Potential Refunds" (February 20, 2026). IEEPA tariffs collected approximately $164.7 billion cumulatively between February 2025 and January 2026, with monthly collections growing from $0.81 billion in February 2025 to $20.8 billion by January 2026. As of January 2026, IEEPA tariffs represented approximately 52% of total U.S. customs duties. The estimated refund liability is in the $100-200 billion range, depending on how unjust enrichment defenses (importers having passed costs to consumers) are resolved, the speed of CBP administrative processing, and the outcome of the government's appeal of the CIT nationwide refund order. Tariff revenue from Section 301 (China), Section 232 (steel and aluminum, autos), and other statutory authorities remains in place. Justice Kavanaugh observed in his dissent that "[t]he United States may be required to refund billions of dollars to importers who paid the IEEPA tariffs, even though some importers may have already passed on costs to consumers or others."

[21] March 2026 Treasury auction stress. Bipartisan Policy Center, "Deficit Tracker" (May 2026). A series of weak Treasury auctions in March 2026 saw primary dealers absorbing approximately 24% of a 2-year note auction in late March (roughly twice the share normally absorbed by these dealers), with 5- and 7-year securities also showing weak demand. Primary dealer takedown is an indicator of softening demand from end investors; when dealers are forced to absorb a larger share, they typically hedge by selling other Treasury positions, contributing to yield rises across the curve. The auction stress is not itself a credibility event but is consistent with the term-premium widening pattern Link 8 describes.

[23] Coefficient sensitivity regression. Illustrative OLS regression on annual data, 1990-2024, constructed from public Bureau of Economic Analysis (nominal GDP), Bureau of Labor Statistics (unemployment, CPI), and Treasury / OMB (federal receipts) figures. Two specifications: (1) receipts-to-GDP ratio regressed on unemployment and excess inflation (inflation above the 2% target), yielding an unemployment coefficient of -0.0050 ratio points per percentage point ($t = -4.7$, $p < 0.001$), implying a fractional receipts sensitivity α of approximately 0.029 against the mean receipts-to-GDP ratio of 0.170; (2) log receipts regressed on log GDP, unemployment, and excess inflation, yielding a GDP elasticity of 0.94 ($t = 47$), an unemployment semi-elasticity of -0.034 ($t = -6.0$), and an excess-inflation coefficient that is not statistically significant ($p > 0.30$). These are illustrative estimates, not

cyclically-adjusted structural elasticities; they do not control for the discrete tax-law changes across the period (notably 2001, 2003, 2017, 2025), do not address simultaneity between receipts and the business cycle, and should be read as establishing an empirical reference range rather than a definitive structural parameter. The CBO's published cyclically-adjusted revenue methodology implies a revenue elasticity to the output gap near unity, which converts via Okun's law (roughly two points of output gap per point of unemployment) to a range consistent with these estimates. The model's $\alpha = 0.05$ sits above this historical range, consistent with the paper's argument (Section 3.8, Link 4) that the AI displacement scenario produces larger revenue contraction than historical cyclical episodes.

This is a working paper. The arguments are subject to revision as the empirical record develops. The mathematical model is offered as a teaching framework, not as an empirical forecast. Comments and corrections are welcome.