FIXED INCOME / MACRO-FINANCE

AI Deflation, Sovereign Risk, and the TIPS Signal

When breakeven inflation can, and cannot, reveal default risk

Sami Miettinen

Working Paper - Combined empirical findings 30 July 2026

ABSTRACT

This paper develops a conditional hypothesis linking artificial intelligence to sovereign credit risk and the relative pricing of nominal Treasury securities and Treasury Inflation-Protected Securities (TIPS). AI may lower task costs, compress labor income, and raise transfer needs, but neither economy-wide deflation nor fiscal distress follows automatically: output, profits, market power, monetary policy, and tax design can offset those channels. A correct yield decomposition is therefore essential. Common sovereign default risk should largely cancel from a maturity-matched nominal-TIPS spread; only differential credit exposure, liquidity, inflation risk, and the TIPS deflation floor can move breakeven inflation away from expected inflation. The paper's testable prediction is that a sufficiently rapid, labor-substituting AI shock can widen direct fiscal-risk measures even as option- and liquidity-adjusted inflation expectations fall. It proposes a two-level empirical design using pre-AI occupational exposure, realized adoption, labor-tax dependence, sovereign credit measures, inflation swaps, and matched nominal-real yield curves. A reproducible public-data diagnostic finds that January 2024 slope-change coefficients have the proposed signs in the full 2019-2026 sample, but both fitted post-break total slopes are positive and the breakeven result reverses under shorter windows. A four-year scan of 44 monthly break dates finds the relative slope-change signs at 35 of 38 adequately powered dates, yet the stronger predicted total-slope divergence at only one date and at no date with joint false-discovery-rate significance. The lowest-BIC break predates ChatGPT. These results validate the code path but neither identify an AI effect nor measure sovereign default. The framework turns a broad post-labor narrative into a set of rejectable propositions while reporting the present evidence as inconclusive.

Keywords: artificial intelligence; TIPS; breakeven inflation; sovereign default risk; labor share; deflation floor; fiscal capacity

Conceptual framework with computational audit. No investment recommendation is made.

1. Introduction

Artificial intelligence creates a plausible but unproven tension for fixed-income markets. On one side, automation can reduce the marginal cost of cognitive services and weaken wage growth in exposed occupations. On the other, a shift of income from labor to capital can erode payroll and personal-income tax receipts, increase transfer spending, and make a highly indebted sovereign more vulnerable. The result could be disinflation in goods and services alongside a higher fiscal-risk premium in government debt.

That tension is economically interesting, but it does not imply that a higher Treasury yield or a wider TIPS breakeven is evidence of AI-driven default risk. A nominal yield combines expected real short rates, expected inflation, real and inflation term premia, liquidity, convenience yield, and potentially credit risk. A TIPS yield combines real-rate compensation with its own liquidity, indexation, deflation-floor, and credit characteristics. The difference between the two is therefore an imperfect measure of expected inflation, but the imperfection must be decomposed rather than labeled.

The paper's central contribution is a disciplined version of the AI-deflation paradox. The strong claim that markets are already confusing sovereign default risk with inflation is replaced by a conditional, falsifiable claim: AI-driven labor substitution can raise sovereign risk while lowering expected inflation, and it can affect breakeven inflation only to the extent that nominal and inflation-linked debt have different state-contingent prices. Recent evidence makes this distinction especially important. Dittmar, Hsu, Roussellet, and Simasek (2026) show that U.S. default risk can matter for the relative pricing of nominal Treasuries and TIPS, but through interactions among default, inflation dynamics, and loss exposure, not through a generic CDS spread simply added to the nominal yield.

CENTRAL IDENTIFICATION CLAIM

If sovereign credit risk affects nominal Treasuries and TIPS equally, it changes both yield levels but drops out of breakeven inflation. Default risk enters the TIPS signal only through differential recovery, state dependence, market segmentation, or interaction with inflation compensation.

The revised framework makes four contributions. First, it separates microeconomic cost reduction from economy-wide deflation and states the conditions required for pass-through. Second, it models fiscal capacity as a portfolio of labor, capital, consumption, and other tax bases rather than treating the wage pool as the sole source of sovereign cash flow. Third, it corrects the sign and content of the breakeven decomposition, including TIPS liquidity and the maturity deflation floor. Fourth, it supplies an empirical strategy with an explicit null hypothesis, observable outcomes, confounders, and falsification tests. Fifth, it reports a reproducible software and public-data audit that separates coefficient recovery from scientific evidence and records where the current design remains non-identifying.

2. From AI Adoption to Prices and Fiscal Capacity

2.1 Transaction costs, competition, and pass-through

Coase (1937) explains firm boundaries as a comparison between internal coordination costs and the costs of using markets. AI agents may reduce search, contracting, monitoring, and communication costs. They may also lower minimum efficient scale and allow smaller firms to coordinate complex production. These mechanisms support the paper's emphasis on organizational change, but they do not imply that transaction costs become zero or that large hierarchies disappear. Data access, compute, legal liability, trust, integration, and market power remain scarce inputs. OECD evidence likewise treats AI's effect on competition as conditional: AI can lower entry and search costs, while control of models, data, and distribution can strengthen incumbents (Filippucci et al. 2024).

The incidence of a productivity shock depends on market structure. In contestable markets, lower unit costs are more likely to pass through into lower prices. Under concentration, switching costs, or scarce complements, part of the gain may remain in markups, wages, or rents. The empirically relevant object is therefore not an assumed 'zero marginal cost' equilibrium but the pass-through coefficient from AI-induced cost savings to sectoral prices.

2.2 Relative price declines are not automatically aggregate deflation

AI can reduce the relative price of exposed services without producing a sustained fall in the aggregate price level. Higher productivity raises real income and may increase demand elsewhere; scarce energy, housing, physical services, and capital inputs may become more expensive; and monetary policy can offset persistent disinflation. The scale of the productivity effect is also uncertain. Acemoglu (2025), for example, estimates modest aggregate productivity gains under current task-level evidence, while other scenarios are materially larger. The correct empirical prediction is thus sectoral disinflation first, with aggregate deflation only when pass-through is broad, demand offsets are weak, and policy accommodation is limited.

2.3 The fiscal channel is a tax-base reallocation problem

The fiscal mechanism is strongest when AI lowers labor income faster than it raises taxable profits, consumption, or new forms of labor demand. The Congressional Budget Office (2024) emphasizes that AI can raise or lower federal receipts: displaced workers reduce income and payroll taxes, but higher productivity, profits, compliance, wages in complementary jobs, and new output can offset the loss. IMF analysis similarly frames the issue as a need to adapt capital-income taxation and social insurance, not as an automatic collapse of revenue (Brollo et al. 2024).

A compact accounting approximation makes the contingency explicit. Let L denote taxable labor income, K taxable capital income, C consumption, and G transfers and other AI-related spending. For effective tax rates tau_L, tau_K, and tau_C, the first-order change in the primary balance is:

$$\Delta PB = \tau_L \Delta L + \tau_K \Delta K + \tau_C \Delta C - \Delta G \quad (1)$$

The sovereign-risk channel requires the negative terms to dominate after behavioral responses and policy changes. It is amplified by high initial debt, short maturity, weak tax capacity, large automatic stabilizers, and political constraints on shifting the tax base. It is weakened by stronger productivity, taxation of capital rents, new labor demand, lower government operating costs, and monetary-fiscal coordination.

3. What TIPS Breakevens Actually Measure

3.1 A credit-, liquidity-, and option-aware decomposition

Let yN(t,m) be the yield on a nominal Treasury and yR(t,m) the yield on a maturity-matched TIPS. A stylized decomposition separates the common expected real-rate component from inflation compensation and security-specific terms. RTP is the real term premium, IRP the inflation risk premium, LP the liquidity or convenience component, CP the credit component, and Odef the annualized value of the TIPS maturity deflation floor:

$$yN(t,m) = \bar{r}(t,m) + E_t[\pi(m)] + RTP(t,m) + IRP(t,m) + LP_N(t,m) + CP_N(t,m) \quad (2)$$

$$yR(t,m) = \bar{r}(t,m) + RTP(t,m) + LP_R(t,m) + CP_R(t,m) - Odef(t,m) \quad (3)$$

Subtracting the real yield from the nominal yield gives observed breakeven inflation:

$$BEI(t,m) = E_t[\pi(m)] + IRP(t,m) + [LP_N - LP_R] + [CP_N - CP_R] + Odef(t,m) \quad (4)$$

The standard empirical convention treats the TIPS liquidity premium as positive in the TIPS yield. If the nominal Treasury liquidity term is normalized near zero, the contribution of TIPS illiquidity to breakeven is negative, not positive (D'Amico, Kim, and Wei 2018). The floor has the opposite sign: a valuable floor raises the TIPS price, lowers its yield, and increases the observed breakeven. Credit enters only as the difference between the nominal and real security's credit terms.

3.2 The deflation floor is real, but maturity and vintage matter

At maturity, the Treasury pays the greater of the inflation-adjusted principal and the TIPS' original principal. That protection is an embedded put on the cumulative price level. Its moneyness depends on the security's reference CPI and accrued index ratio, so two TIPS with the same remaining maturity can have different floor values. A newly issued security is generally more exposed to the floor than a seasoned TIPS with substantial accrued inflation. The floor protects original par at maturity, not an investor's secondary-market purchase price, and coupon payments can still decline when the inflation-adjusted principal falls.

The floor should not be valued with an unmodified Black-Scholes formula. CPI indexation is lagged and discrete; inflation has a term structure and is not a traded lognormal asset; and the payoff depends on the issue vintage. A no-arbitrage inflation term-structure model, inflation caps and floors, or matched-vintage relative pricing is better suited to the task. Grishchenko, Vanden, and Zhang (2016) find that the option contains information about future inflation, but its historical yield effect was usually small for 10-year TIPS and much larger for some 5-year securities during the 2008-2009 deflation scare. The floor is therefore a state-dependent correction, not a standing explanation for every elevated breakeven.

3.3 Why common default risk usually cancels

Nominal Treasuries and TIPS are obligations of the same sovereign. If a default event applies the same expected loss to both securities, CP_N equals CP_R and the credit term vanishes from equation (4), even though both yields rise. This is the null hypothesis. Differential pricing can arise if recovery, payment delay, legal treatment, inflation at default, or investor segmentation makes expected loss different across the two claims.

U.S. debt-limit episodes demonstrate that Treasuries can contain measurable event risk. Cashin et al. (2017) estimate that yields across Treasury securities rose during the 2011 and 2013 impasses, with especially large effects on bills exposed to delayed payment. More directly, Dittmar et al. (2026) examine nominal Treasuries, TIPS, inflation swaps, and U.S. default risk jointly. Their results support a differential-pricing channel related to the interaction between inflation dynamics and default. This evidence strengthens the paper's topic while narrowing its claim: CDS cannot simply be subtracted from BEI with a constant coefficient to recover 'true' expected inflation.

3.4 A diagnostic map for observed yield movements

Table 1. Competing interpretations of Treasury and TIPS movements

Observed pattern	Inflation reading	Fiscal/default reading	Discriminating evidence
Nominal yield up; TIPS yield little changed; BEI wider	Higher expected inflation or inflation risk premium	Possible only if nominal credit exposure rises relative to TIPS	Compare inflation swaps, surveys, CDS, floor value, and liquidity-adjusted curves
Nominal and TIPS yields rise together; BEI stable	Higher expected real rates or common term premium	Consistent with common sovereign risk that cancels in BEI	Check CDS, OIS or convenience-yield wedges, debt-limit dates, and term-premium models

Observed pattern	Inflation reading	Fiscal/default reading	Discriminating evidence
TIPS yield falls; BEI widens; floor-rich issues outperform	Could look like higher inflation compensation	Not necessary	Compare new and seasoned TIPS with different index ratios; estimate the deflation option
Inflation-swap-hedged BEI moves with CDS	Not explained by expected inflation alone	Evidence for differential credit pricing or credit-inflation interaction	Control for swap counterparty risk, TIPS liquidity, issuance, and risk sentiment

Note: Each row is a diagnostic, not a sufficient identification result. 'Hedged BEI' compares bond-implied inflation compensation with inflation-swap compensation and still requires liquidity and counterparty controls.

4. The AI-Fiscal Risk Hypothesis

The framework predicts a joint pattern, not a single yield movement. A labor-substituting AI shock should first appear in exposed-sector prices, labor income, vacancies, and tax receipts. Sovereign credit measures should react only when the shock is large relative to the government's fiscal buffers and alternative tax bases. TIPS breakevens should reflect the event only after inflation risk, liquidity, the floor, and relative credit exposure are separated.

H1. Labor-tax exposure

For a given AI adoption shock, sovereign credit spreads should widen more in jurisdictions with higher pre-shock dependence on payroll and personal-income taxes, weaker taxation of capital income, larger transfer elasticities, and higher initial debt-service burdens.

H2. Inflation-credit divergence

In the disruptive-substitution scenario, option- and liquidity-adjusted expected inflation should fall or remain anchored while direct fiscal-risk measures, such as sovereign CDS or cash-OIS/convenience wedges, rise. A rise in both expected inflation and credit risk is possible under a fiscal-dominance response, so the sign of inflation must be estimated rather than imposed.

H3. The common-credit null

If nominal Treasuries and TIPS have equal expected loss in default, AI-related fiscal risk should raise their yields similarly and have no direct effect on BEI. Rejecting this null requires evidence that hedged breakevens or matched relative prices move with credit risk after floor and liquidity adjustments.

H4. Deflation-floor heterogeneity

A genuine rise in deflation-tail risk should benefit floor-rich TIPS more than seasoned issues with high index ratios, holding duration and liquidity constant. This cross-sectional vintage prediction is cleaner than inferring the floor from an aggregate breakeven alone.

H5. Organizational exposure in corporate credit

Within corporate bonds, spreads should widen for firms whose margins rely on automatable coordination, search, or routine cognitive tasks and narrow for firms that capture scarce data, compute, distribution, or physical complements. The relevant exposure is an ex ante task and rent structure, not firm size or headcount by itself.

5. Empirical Design

5.1 Measuring an AI labor-income shock

The original proposal combined the inverse labor share, software spending, and earnings-call language in a principal component. That index risks building the hypothesized outcome directly into the treatment: labor share is itself a result of automation, software spending is endogenous to expected demand, and management language can respond to financing conditions. A cleaner design separates predetermined exposure from realized adoption.

Construct baseline exposure before the generative-AI shock from occupational task content, using an established index such as the ILO-NASK task measure or comparable occupation-level exposure. Aggregate it to a country, region, industry, or firm using pre-period employment weights. Interact that fixed exposure with a common time series of frontier capability or adoption. Validate realized adoption independently with software investment, AI vacancies, firm disclosures, and measured tool use. This shift-share structure does not solve every identification problem, but it avoids treating a contemporaneous fall in labor share as proof of its own cause.

$$AIShock(j,t) = Exposure(j,pre) \times GlobalAI(t) \quad (5)$$

5.2 Sovereign outcomes and pricing controls

For the United States, use maturity-matched zero-coupon nominal and TIPS curves, zero-coupon inflation swaps, TIPS issue-level prices and index ratios, survey inflation expectations, sovereign CDS, OIS or other cash-derivative wedges, and debt-limit event indicators. The Federal Reserve's D'Amico-Kim-Wei decomposition provides a benchmark for expected inflation, inflation risk, and TIPS liquidity. Issue-level data are required to value the floor and separate on-the-run liquidity from vintage effects.

The U.S. aggregate time series is short and AI adoption trends smoothly, so causal claims from a single VAR would be fragile. The principal design should be an international or sub-sovereign panel that combines predetermined AI exposure with fiscal structure. Countries with both nominal and inflation-linked sovereign bonds offer the closest cross-sectional analogue to TIPS; U.S. evidence then serves as a high-quality case study.

5.3 Main panel specification

Let Risk(j,t) be a direct sovereign-risk outcome, such as a CDS spread, a cash-OIS wedge, or an inflation-swap-hedged nominal-real bond wedge. Let LaborTax(j,pre) be the pre-period share of revenue tied to labor income. A difference-in-differences or local-projection specification is:

$$\text{Delta Risk}(j,t+h) = \text{beta\_h}\ [\text{AIShock}(j,t) \times \text{LaborTax}(j,pre)] + \text{gamma\_h}\ X(j,t) + \text{mu\_j} + \text{lambda\_t} + \text{error}(j,t+h) \quad (6)$$

Country fixed effects absorb stable institutions; time effects absorb global rates and common AI news. Controls should include debt and maturity, primary balance, growth surprises, monetary-policy surprises, commodity and supply shocks, inflation surprises, exchange-rate regime, market volatility, bond issuance, and liquidity. The coefficient beta_h is the fiscal-sensitivity test. Equation (6) should be estimated separately for direct credit measures, expected inflation, floor-adjusted BEI, and real yields so that a single price is not forced to identify several latent components.

5.4 U.S. event and issue-level tests

The U.S. design can exploit cross-sectional TIPS variation around common deflation or fiscal-risk shocks. Match new and seasoned issues by duration and liquidity, then test whether securities with more in-the-money floors outperform when AI news lowers exposed-sector inflation expectations. Separately, compare nominal Treasuries, TIPS, inflation swaps, and CDS around debt-limit events and large, pre-specified AI capability or adoption news. The identifying event set must be fixed before inspecting returns to avoid narrative selection.

5.5 Corporate-credit complement

At the firm level, estimate spread changes using pre-period occupational exposure, business-model text, and industry competition. Include firm and time fixed effects, duration, leverage, profitability, rating, issue liquidity, and industry demand. A widening spread for exposed incumbents accompanied by entry, price declines, or margin compression would support the organizational-disruption mechanism. A rally in firms controlling AI infrastructure would be consistent with rent capture rather than universal deflation.

5.6 Falsification and robustness

Pre-trends. High- and low-exposure jurisdictions should not display diverging fiscal-risk trends before the common AI shock.

Placebos. Exposure to tasks that AI cannot plausibly perform should not predict the same spread response; future adoption should not explain past yields.

Alternative inflation measures. Results should survive replacement of raw BEI with survey expectations, inflation swaps, and model-based expected inflation.

Liquidity and issuance. Control for on-the-run status, bid-ask spreads, TIPS supply, dealer constraints, and flight-to-safety episodes.

Policy confounders. Separate AI news from fiscal announcements, debt-limit episodes, monetary surprises, energy prices, tariffs, and recession news.

State dependence. Estimate effects by initial debt, tax structure, monetary regime, and labor-market institutions rather than assuming one global coefficient.

Break-date and window sensitivity. Report all defensible AI event dates, sample starts, and HAC lag choices. A selected January 2024 break is not an intervention unless it is tied to a pre-specified, economically relevant adoption or capability shock.

Persistence and inference. Test for residual serial correlation, stochastic trends, and structural instability. Newey-West standard errors address a class of covariance problems but do not prevent spurious level regressions or repair an endogenous break date.

5.7 Estimand discipline and minimum evidence

In a segmented trend model, the coefficient on post_trend is a change relative to the pre-event slope. It is not the post-event slope itself. The latter equals the coefficient on time_trend plus the coefficient on post_trend and must be reported with its covariance-based standard error. The post_shock coefficient is an immediate level change only when post_trend is centered at zero in the event month. The supplied prototype instead multiplied the uncentered global trend by the post indicator; the accompanying implementation corrects this parameterization.

A result will count as evidence for H2 only if a pre-specified AI shock changes a direct sovereign-risk measure and an adjusted inflation-expectations measure in the predicted joint direction, survives window and event-date sensitivity, and is not explained by fiscal, monetary, recession, liquidity, or supply news. A corporate spread or a broad financial-stress index can debug the pipeline, but neither outcome is a sovereign default rate or probability.

6. Scenarios and Policy Implications

6.1 Augmentation and broad productivity growth

If AI complements workers, creates new tasks, and raises wages and output, the tax base can expand even as some service prices fall. Real yields may rise with expected growth, credit risk may decline, and BEI may remain near target. This scenario rejects the paper's disruptive fiscal channel without rejecting AI's economic importance.

6.2 Labor substitution with fiscal adaptation

If labor income falls but governments tax capital rents, consumption, land, or excess profits and use AI to improve compliance and service delivery, fiscal capacity can be preserved. The distributional conflict may be large, but default is not mechanically implied. This scenario predicts changes in tax composition more strongly than changes in sovereign spreads.

6.3 Rapid substitution with delayed policy response

The paradox is strongest when substitution is fast, pass-through is broad, labor-tax dependence is high, transfers rise, and tax reform lags. Disinflation increases the real burden of nominal debt while weaker primary balances raise refinancing needs. Sovereign-risk measures can then widen even as adjusted expected inflation falls. For a monetary sovereign, however, distress may be resolved through inflation, financial repression, maturity extension, taxation, or delayed payment rather than outright default. The pricing consequence depends on which response investors expect.

6.4 Universal basic income is a policy branch, not an inevitability

A large universal basic income would raise gross financing needs if introduced without offsetting taxes or spending reductions. Yet its fiscal effect depends on design, replacement of existing transfers, behavioral responses, and the tax base used to fund it. The useful research question is not whether UBI creates a mathematically inevitable 'doom loop,' but which transfer-tax combinations stabilize aggregate demand without causing debt service to outrun revenue capacity. That policy branch can be simulated within equation (1).

7. Limitations and Conditions for Rejection

The framework is deliberately conditional. AI exposure is not realized displacement; adoption may augment jobs; task measures are noisy; tax systems adapt; and the central bank influences aggregate inflation. Sovereign CDS is also an imperfect credit measure because it is thinly traded, affected by contract definitions and counterparty risk, and sensitive to technical default as well as insolvency. Inflation swaps introduce their own liquidity and counterparty components. No single series is a clean sufficient statistic.

The current public-data exercise has additional limitations. January 2024 is imposed rather than derived from a pre-registered shock; the 2019-2026 window includes the pandemic, reopening inflation, aggressive monetary tightening, debt-limit episodes, and later fiscal and trade news; raw five-year breakeven inflation is not option- or liquidity-adjusted expected inflation; and the Baa-Treasury spread is corporate rather than sovereign credit. The short post period, persistence of yield series, and inspection of many possible breaks make nominal HAC p-values insufficient evidence of causality.

The hypothesis should be rejected or narrowed if high-exposure, labor-tax-dependent jurisdictions do not show relative deterioration in fiscal cash flows or credit measures after adoption; if the effect disappears after ordinary macroeconomic and liquidity controls; if nominal and real sovereign bonds respond identically; or if measured adoption raises wages, output, and revenue enough to offset displacement. A result that AI lowers sectoral prices but leaves fiscal risk unchanged would support technological disinflation without the sovereign default channel.

8. Computational Validation and Preliminary Diagnostic

8.1 What the tests establish

The accompanying Python package implements a centered monthly interrupted time series with Newey-West/HAC covariance and a corporate PanelOLS design with CUSIP and month fixed effects. Fourteen automated tests pass. The fixtures contain 50 CUSIPs observed for 60 months (3,000 rows), verify that event month -1 is omitted, verify every exposure interaction exactly, run both estimators without rank failure, recover the imposed coefficient directions, and produce explicit errors for missing data. This establishes internal software consistency. It does not support H1-H5 because the synthetic outcomes were generated from the hypothesized coefficients.

8.2 Public FRED interrupted-time-series diagnostic

As a transparent stress test, monthly means of the FRED Baa-Treasury spread (BAA10Y), five-year breakeven inflation (T5YIE), and St. Louis Fed Financial Stress Index (STLFSI4) were fit from January 2019 through July 2026 with an imposed January 2024 break and four Newey-West lags. Table 2 distinguishes the immediate level change, the slope change relative to the pre-period, and the total fitted post-break slope.

Table 2. January 2024 public-data ITS diagnostic

Outcome	Level change	Slope change	Total post slope	Scientific reading
Baa-Treasury spread	-0.313 ($p<0.001$)	+0.0164 / month ($p<0.001$)	+0.00568 / month ($p=0.034$)	Corporate risk proxy fell at the break, then drifted upward; not sovereign default.
5-year breakeven	-0.516 ($p=0.037$)	-0.0145 / month ($p=0.032$)	+0.00742 / month ($p=0.027$)	Trend slowed relative to its pre-trend but did not become negative; raw BEI is unadjusted.

Outcome	Level change	Slope change	Total post slope	Scientific reading
Financial stress index	-0.196 (p=0.222)	+0.0162 / month (p=0.053)	+0.00641 / month (p=0.142)	No conventionally significant level or total-slope result.

Notes: Level change is beta_post_shock. Slope change is beta_post_trend. Total post slope is beta_time_trend + beta_post_trend, tested using the full HAC covariance matrix. Spread and breakeven coefficients are percentage points; STLFSI4 is in index units. Values are descriptive and unadjusted for multiple break-date inspection.

8.3 Four-year break-date scan and evidentiary verdict

To test whether the January 2024 result is an artifact of choosing one date, the same centered segmented-trend model was re-estimated at every monthly break from July 2022 through February 2026. The scan contains 44 candidate breaks; 38 have at least 12 post-break months and enter the multiple-testing analysis. Four Newey-West lags are used throughout. Benjamini-Hochberg false-discovery-rate q-values are calculated across the 38 adequately powered dates separately for the credit and breakeven outcomes. Named frontier-model releases are annotations, not proof that AGI occurred or that the release was an exogenous macroeconomic intervention.

The weak, relative-sign criterion is common: the Baa-Treasury slope change is positive while the breakeven slope change is negative at 35 of 38 adequately powered breaks. That frequency is not compelling evidence by itself because it can be generated by the same smooth, pre-existing trends under many adjacent break dates. The stronger economic criterion - a positive fitted post-break credit slope and a negative fitted post-break breakeven slope - appears at only 1 of 38 dates, and not significantly. Both relative slope changes survive the 5 percent FDR threshold at only three adjacent dates (September-November 2023); no date has both total post-break slopes significant after FDR adjustment. The lowest joint Bayesian information criterion occurs in September 2022, before the public ChatGPT launch, which is more consistent with a pre-existing inflation-credit regime than with a frontier-model event.

FOUR-YEAR SPECIFICATION-SEARCH RESULT

Relative signs: 35/38. Strong total-slope divergence: 1/38. Joint FDR-significant relative changes: 3/38, clustered in September-November 2023. Joint FDR-significant total slopes: 0/38. Best joint BIC: September 2022.

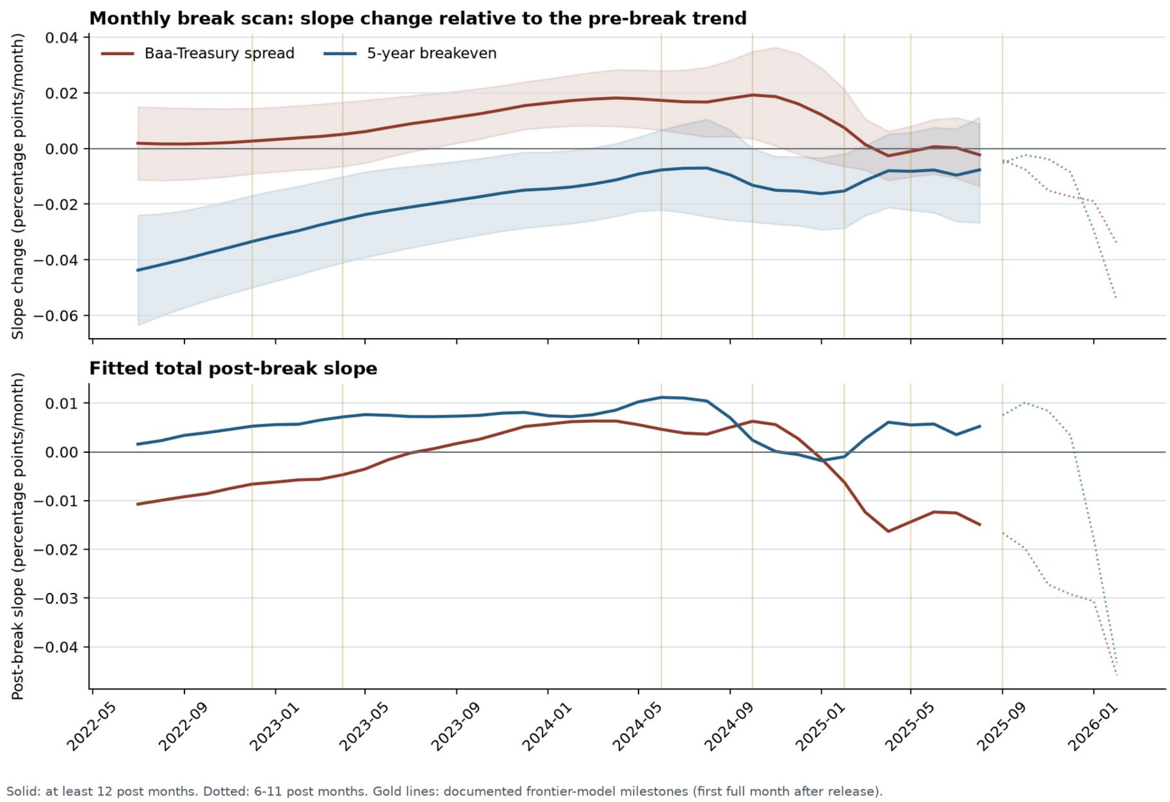

Figure 1. Monthly break-date scan, July 2022-February 2026

Notes: Top panels show slope changes relative to the pre-break trend; bottom panels show fitted total post-break slopes. Shading shows 95 percent HAC confidence intervals for the relative slope changes. Solid segments have at least 12 post months; dotted segments have 6-11. Vertical gold lines mark the first full calendar month after public model releases. BAA10Y is corporate credit and T5YIE is raw breakeven inflation.

8.4 Named frontier-model milestones

No named milestone supplies robust joint evidence. Table 3 reports scan-adjusted q-values for each relative slope change. A milestone would still fail the paper's substantive H2 test if its fitted total slopes do not diverge in the predicted direction, even when one relative slope change is statistically detectable.

Table 3. Named frontier-model release diagnostics

Public milestone	Break	Post months	Slope-change q credit / BEI	Joint reading
ChatGPT launch	Dec 2022	44	.870 / .0005	BEI only; fitted total directions oppose H2.
GPT-4	Apr 2023	40	.638 / .0046	BEI only; fitted total directions oppose H2.
GPT-4o	Jun 2024	26	.008 / .331	Credit only; fitted BEI slope remains positive.
OpenAI o1-preview	Oct 2024	22	.047 / .076	Closest named fit, but not joint at 5%; both total slopes positive.
DeepSeek-R1	Feb 2025	18	.502 / .047	BEI only; both fitted total slopes negative.

Public milestone	Break	Post months	Slope-change q credit / BEI	Joint reading
OpenAI o3/o4-mini	May 2025	15	.870 / .299	No relative evidence; fitted total directions oppose H2.
GPT-5	Sep 2025	11	Not adjusted	Underpowered by the pre-set 12-month rule.

Notes: Break is the first full calendar month after each public release. q-values are Benjamini-Hochberg adjusted across the 38 adequately powered monthly breaks, separately by outcome. Credit is BAA10Y; BEI is T5YIE. Release dates are from official OpenAI and DeepSeek materials.

CURRENT EVIDENTIARY VERDICT

Not proven. No named milestone produces jointly robust credit-breakeven evidence, the best-fitting break predates ChatGPT, the public outcomes do not measure sovereign default, and raw BEI is unadjusted. The scan also does not disprove the conditional theory: these proxies and an endogenous break search are inadequate to adjudicate the required sovereign-risk and adjusted-inflation prediction.

9. Conclusion

Artificial intelligence can plausibly create a world in which selected prices fall while fiscal risk rises, but TIPS do not provide a one-line proof of that world. Breakeven inflation is a relative price. It contains expected inflation, inflation risk, a negative contribution from TIPS illiquidity, a positive contribution from the deflation floor, and any difference in credit exposure between nominal and inflation-linked claims. Common sovereign default risk raises yield levels and normally cancels from the spread.

The AI-deflation paradox is therefore a joint empirical prediction: labor-substituting adoption should reduce labor income and exposed-sector prices; the fiscal effect should be strongest where the tax system relies on labor and adapts slowly; direct credit-risk measures should widen relative to adjusted inflation expectations; and any TIPS effect should appear through vintage-specific floor values or differential default pricing. The computational audit shows that the estimators can recover such a pattern when it is built into synthetic data. The four-year break scan shows why the public proxy exercise does not establish it in the economy: weak relative signs recur across many adjacent dates, the stronger total-slope divergence is essentially absent, and the best-fitting date predates the public generative-AI shock. The next evidentiary step is a pre-specified, exposure-based panel using direct sovereign-risk and adjusted inflation measures. Until that test is completed, the mechanism should be presented as coherent and falsifiable, not as a demonstrated effect.

References

Acemoglu, D. (2025). The simple macroeconomics of AI. Economic Policy, 40(121), 13-58. doi:10.1093/epolic/eiae042

Arellano, C. (2008). Default risk and income fluctuations in emerging economies. American Economic Review, 98(3), 690-712. doi:10.1257/aer.98.3.690

Brollo, F., Dabla-Norris, E., de Mooij, R., Garcia-Macia, D., Hanappi, T., Liu, L., and Nguyen, A. D. M. (2024). Broadening the gains from generative AI: The role of fiscal policies. IMF Staff Discussion Note 2024/002. doi:10.5089/9798400277177.006

Cashin, D., Ferris, E. S., Klee, B., and Stevens, C. (2017). Take it to the limit: The debt ceiling and Treasury yields. Finance and Economics Discussion Series 2017-052, Board of Governors of the Federal Reserve System. doi:10.17016/FEDS.2017.052

Coase, R. H. (1937). The nature of the firm. Economica, 4(16), 386-405. doi:10.1111/j.1468-0335.1937.tb00002.x

Congressional Budget Office. (2024). Artificial intelligence and its potential effects on the economy and the federal budget. Source

D'Amico, S., Kim, D. H., and Wei, M. (2018). Tips from TIPS: The informational content of Treasury Inflation-Protected Security prices. Journal of Financial and Quantitative Analysis, 53(1), 395-436. doi:10.1017/S0022109017000916

Dittmar, R. F., Hsu, A., Roussellet, G., and Simasek, P. (2026). Default risk and the pricing of U.S. sovereign bonds. Journal of Finance, 81(2), 829-869. doi:10.1111/jofi.70014

Eaton, J., and Gersovitz, M. (1981). Debt with potential repudiation: Theoretical and empirical analysis. Review of Economic Studies, 48(2), 289-309. doi:10.2307/2296886

Filippucci, F., Gal, P., Jona-Lasinio, C., Leandro, A., and Nicoletti, G. (2024). The impact of artificial intelligence on productivity, distribution and growth. OECD Artificial Intelligence Papers, No. 15. doi:10.1787/8d900037-en

Fleckenstein, M., Longstaff, F. A., and Lustig, H. (2014). The TIPS-Treasury bond puzzle. Journal of Finance, 69(5), 2151-2197. doi:10.1111/jofi.12032

Grishchenko, O. V., Vanden, J. M., and Zhang, J. (2016). The informational content of the embedded deflation option in TIPS. Journal of Banking and Finance, 65, 1-26. doi:10.1016/j.jbankfin.2015.12.004

Gmyrek, P., Berg, J., Kaminski, K., Konopczynski, F., Ladna, A., Nafradi, B., Roslaniec, K., and Troszynski, M. (2025). Generative AI and jobs: A refined global index of occupational exposure. ILO Working Paper 140. doi:10.54394/HETP0387

Federal Reserve Bank of St. Louis. (2026). Five-year breakeven inflation rate [T5YIE]; Moody's seasoned Baa corporate bond yield relative to the 10-year Treasury [BAA10Y]; and St. Louis Fed Financial Stress Index [STLFSI4]. Retrieved 30 July 2026. Source

OpenAI. (2023). GPT-4 technical report and research release, 14 March 2023. Source

OpenAI Economic Research. (2025). How people use ChatGPT. Research paper; documents the public launch on 30 November 2022. Source

OpenAI. (2024a). Hello GPT-4o, 13 May 2024. Source

OpenAI. (2024b). Introducing OpenAI o1-preview, 12 September 2024. Source

DeepSeek. (2025). DeepSeek-R1 release, 20 January 2025. Source

OpenAI. (2025a). Introducing o3 and o4-mini, 16 April 2025. Source

OpenAI. (2025b). Introducing GPT-5, 7 August 2025. Source

TreasuryDirect. (2026). Treasury Inflation-Protected Securities (TIPS). U.S. Department of the Treasury, Bureau of the Fiscal Service. Source