Oil, Hormuz, and the Price of AI Capital: Geopolitical Inflation, Monetary Tightening, and the Crowding Out of Treasury Debt

Mark Uni
Pufferfish Artificial Intelligence Research Group

Corresponding author: Mark Uni. Email: to be supplied by the author at submission.

Word count: 11,948 words, inclusive of tables, references, figure captions, footnotes and endnotes.

> **Companion papers.** Finance: *Debt-Financed AI Expansion and Shareholder Value at Oracle* (Journal of Finance format). Technology: *The Adoption-Impairment Paradox* (Research Policy format). Both cross-cite the paid *Oracle Short Report*; core figures are identical across all four.

Formatting note. Not submitted to, accepted by or published in any journal; no volume, DOI, submission number or acceptance date is claimed.

Declaration of interest (for the editor). Mark Uni is a researcher at Pufferfish Artificial Intelligence Research Group. He has no employment or consulting relationship with Oracle and holds no position in ORCL securities. The related *Oracle Short Report* is a separate paid product prepared by other members of the group using the results of this research; he does not prepare or sell it. This paper contains no investment recommendation, target price, options combination, stop-loss or trading advice.

ABSTRACT

Artificial intelligence infrastructure is being built with borrowed money, on a scale large enough to compete with sovereign issuance. This paper asks how geopolitical oil risk reaches the cost of that capital. It traces, as scenario and mechanism analysis, one transmission chain: escalation or disruption risk around the Strait of Hormuz raises oil prices and inflation expectations; a central bank charged with price stability has no costless alternative to higher or persistently high policy rates, so the risk-free leg of every borrower's financing cost rises mechanically, before any firm-specific spread moves. When governments, AI hyperscalers and other borrowers then draw on the same pool of capital at once, liquidity tightens and yields rise for all; reported commentary describes crowding running in reverse, with large, price-insensitive AI issuance pushing Treasury yields higher. Oracle, the borrower in the companion finance paper, is a case of exposure, not proof: on a claim stack of $457.1 billion, each additional 100 basis points adds about $4.57 billion of annual interest and, at a 34 per cent margin, about $13.4 billion of required annual revenue. Geopolitical risk is not background noise for AI capital. It is a direct input to its break-even rate.

KEYWORDS: geopolitics of finance; oil shocks; Strait of Hormuz; monetary tightening; crowding out; Treasury issuance; artificial intelligence infrastructure; cost of capital

1. Introduction: a geopolitical shock arrives as an interest rate

Debates about the artificial intelligence (AI) investment boom divide into two conversations: technologists and equity investors ask whether the technology will work, how fast it will be adopted, and who will win; macroeconomists ask about oil, inflation and the path of policy rates. For a debt-financed AI builder they are the same conversation, joined by a single price: the interest rate at which the infrastructure is funded.

The reason is timing. AI infrastructure is bought now — land, power connections, buildings and accelerators — and paid for over fifteen to twenty years, the horizon of the leases and bonds that finance it. The revenue meant to justify it arrives later, at a pace set by adoption. Anything that raises the financing rate in the interval raises the break-even revenue the investment must eventually produce, whether or not the technology succeeds. A geopolitical shock need not touch a data centre to damage its economics; it need only move oil, inflation expectations and the policy rate.

The research question is narrow and distinct from those of the companion papers. The finance companion asks whether one firm's debt-financed expansion destroys shareholder value; the technology companion asks why AI hardware can be impaired whether the technology advances or not. This paper asks a prior, external question: *how does geopolitical oil risk transmit through inflation, monetary policy, liquidity and Treasury issuance to raise the cost of AI capital, even for a firm whose own balance sheet is taken as given?*

The argument is built from three predecessor literatures, used as mechanisms rather than as decoration: bargaining theory (Fearon 1995; Powell 2006), oil-shock macroeconomics (Hamilton 1983; Kilian 2009), and the political economy of credit and Treasury supply (Strange 1986; Frieden 1991; Krishnamurthy and Vissing-Jorgensen 2012).

Three clarifications govern everything that follows. First, the argument is scenario and mechanism analysis of an escalation or disruption risk centred on the Strait of Hormuz and the United States–Iran standoff. It does not assert that a United States–Iran war has occurred, and no conclusion depends on treating a scenario as an established fact. Second, the firm-level figures used to size the final step are taken, unchanged, from the companion finance paper's filing-based anchors: a claim stack of $457.1 billion, a five-year credit default swap (CDS) spread of 261 basis points, and fiscal 2026 capital expenditure of $55.7 billion. Third, reported remarks by Treasury Secretary Scott Bessent on AI issuers' willingness to borrow are treated throughout as reported market commentary, not as a verified quotation; no quotation is invented to fill that gap.

Sections 2 to 8 build the chain. Sections 9 to 12 cover siting, sovereign clouds, resilience and policy-rate scenarios; Sections 13 to 18, the foundations; Section 19, the two growth regimes; Sections 20 to 22, implications, limitations and the conclusion.

2. When credit felt unlimited: the zero-rate regime and its memory

Textbook mechanics shaped a generation of borrowers. A bank keeps a fraction of each deposit as reserves and lends the rest; the loan is spent and redeposited, and the process repeats, so an initial $100 can support a multiple in deposits and credit: at a 10 per cent reserve fraction the arithmetic limit is $1,000, at 5 per cent $2,000. The multiplier is finite — constrained by capital, regulation and demand as much as by reserves — but the loop's disciplining price is the interest rate.

$$\text{Total deposits supported} \approx \text{initial deposit} \div \text{reserve fraction}$$

At a policy rate near zero, that disciplining price nearly vanishes. Borrowing to build an asset whose payoff lies a decade away costs almost nothing; refinancing is routine; maturities are rolled rather than repaid. Credit expansion feels unlimited — not because the multiplier is infinite, but because nothing in the price of credit tells the borrower to stop. The AI boom's institutional memory was formed in that regime, which lasted, with interruptions, from the financial crisis to inflation's return in the early 2020s (Blyth 2013; Tooze 2018).

Once inflation returns, the same loop runs under a positive and rising price of credit, rollover becomes a cost rather than a formality, and projects financed on the assumption of cheapness must be re-underwritten at the new rate. Nothing about the project needs to have changed for its break-even point to move.

The AI capital boom is exposed to geopolitics before it is exposed to competition or obsolescence. Much of it was conceived, and its earliest tranches priced, where the financing leg could be treated as a constant.

3. The Hormuz channel: from a strait to an inflation expectation

The Strait of Hormuz is the passage between Iran and the Arabian Peninsula through which roughly one-fifth of global petroleum liquids consumption transits, on the United States Energy Information Administration's chokepoint assessments, together with a substantial share of liquefied natural gas trade. Iran controls the northern shore and the islands commanding the shipping lanes, and its capacity to harass, mine or temporarily close the strait in a confrontation with the United States is a standing feature of energy security analysis (Krasner 1976; Gilpin 1987). The scenario is escalation along that fault line: a confrontation, blockade threat or partial disruption severe enough to remove, or credibly threaten to remove, a material share of transit for weeks to months.

The first-round effect does not require a barrel to be lost. Oil is priced forward, on expectations of scarcity: a credible disruption risk reprices crude immediately, raises refined product and shipping insurance costs, and widens the distribution of outcomes traders must hedge. Hamilton (1983) established the empirical association between oil shocks and macroeconomic contraction in the post-war United States; Kilian (2009) showed the effect depends on the shock's source: a geopolitical supply disruption behaves differently from a demand-driven price rise, and is more likely to be contractionary alongside higher prices. Kilian's third category, a precautionary-demand shock driven by fear of future scarcity, is the closest analogue to a Hormuz threat that has not removed a barrel.

The second-round effect runs through inflation expectations. Energy enters headline inflation directly, production and transport costs indirectly, and the public's and the market's judgment about whether a price-level shock will become an inflation process. A central bank that might look through a one-off relative price change cannot look through a de-anchoring of expectations: the event reaches monetary policy not through the volume of oil lost but through the expectation it creates.

First, the channel is asymmetric: transit can be threatened quickly and restored only slowly, because insurance, escort and inventory behaviour normalise more slowly than headlines. Second, it is indiscriminate: the shock raises the energy and inflation outlook for every economy importing through the strait and, through globally priced oil, for economies importing none of it.

4. The central bank's constraint: no costless alternative

A central bank with a price-stability mandate, faced with an oil-driven inflation shock, has three options: raise rates, hold rates high for longer, or cut rates and accept the inflation. Only the first two defend the mandate; the 28 October 2026 decision — a 10–2 hold at 3.75 to 4.00 per cent, both dissents preferring an increase — is the second in operation (Board of Governors 2026b). The third is not costless; it lets an energy shock become embedded in wages, contracts and expectations, to be removed at higher cost. The 1970s supply shocks remain the canonical demonstration, and the literature, from Bernanke, Gertler and Watson (1997) to Blanchard and Gali (2010), turns on the same trade-off: policy can cushion the output loss or contain the inflation — a choice over the distribution of pain, not its existence.

For borrowers, the implication is mechanical. A corporate borrowing cost decomposes, approximately, into a risk-free rate for the relevant maturity plus a credit spread:

$$\text{Borrowing cost} \approx \text{risk-free rate} + \text{credit spread}$$

Tightening moves the first term for every borrower, including borrowers whose own creditworthiness has not changed. Spread widening is a second movement: an analysis that watches only a firm's CDS spread understates its exposure, because the CDS measures the second term while the shock is moving the first.

Persistence matters as much as size: AI infrastructure is financed with instruments whose cost is set by the expected path of rates over years. An escalation that persuades markets that rates will stay high for longer reprices the curve at which data centre bonds and leases are written, including issuance planned.

Nor can fiscal policy be assumed to neutralise the move: a government facing an energy shock borrows more — to subsidise energy costs, support defence and fund industrial policy — at the same higher rates, in the same market.

5. One pool of capital: the liquidity squeeze

Interest rates are set at the margin, in a market, by the balance between the supply of savings and the demands on it. The AI boom has introduced a new class of borrower at exceptional scale and speed: a small number of hyperscalers and their infrastructure partners, issuing investment-grade debt in sovereign-scale sizes, on timetables driven by competitive pre-emption rather than by the cost of funds. They borrow alongside, not after, everyone else: governments funding deficits and refinancing maturing stock, utilities building generation and grid capacity, and other corporates refinancing zero-rate-era debt.

When demands of this size arrive together, the result is not rationing but a price. Yields rise until the marginal buyer is induced to hold the supply, term premia rebuild, dealer balance sheets command a higher return, and liquidity — the ease of selling without moving the price — deteriorates when issuance is heaviest. Adrian and Shin (2010) and Brunnermeier and Pedersen (2009) formalise the amplification: funding and market liquidity reinforce each other, so a tightening that begins as a repricing can become, under stress, a spiral of higher haircuts and wider spreads forcing sales that widen spreads.

First, the squeeze is not confined to the weakest borrower: strong borrowers pay more, their advantage a smaller increase, not immunity. Second, it interacts with maturity: AI projects need long money, and long money is what a market repricing inflation risk supplies most reluctantly. Third, it is pro-cyclical with the geopolitical shock: the same escalation that raises the policy rate raises government borrowing needs and investor demand for liquidity, so the demands peak together rather than offsetting.

A profitable, investment-grade borrower can always fund itself; the market will clear. The question is at what price, and what that price does to the revenue an AI asset must earn (Section 7).

6. Reverse crowding: when AI issuance moves the Treasury

The textbook mechanism is crowding out: government borrowing absorbs savings, raises yields, and displaces private investment. The empirical literature, from Engen and Hubbard (2004) to Laubach (2009), estimates a positive effect of federal debt and deficits on interest rates; Krishnamurthy and Vissing-Jorgensen (2012) and Greenwood, Hanson and Stein (2015) explain why the Treasury's particular supply — of safe, liquid, maturity-specific assets — has its own price effects.

Public reporting on the AI debt boom describes a mechanism running, in part, the other way. Treasury Secretary Scott Bessent is reported noting AI companies' willingness to issue debt regardless of borrowing costs, and market commentary argues that large AI hyperscaler issuance competes with Treasury issuance for the same buyers, the adjustment appearing in higher Treasury yields rather than in wider corporate spreads. On that account, a private investment boom of scale and price-insensitivity does not merely suffer the sovereign's cost of funds; it raises it.

> **A note on sources.** The preceding paragraph reports the direction of a mechanism as described in public reporting and market commentary. The exact wording attributed to Secretary Bessent has not been verified against a primary source — a Treasury transcript or the original interview — and no quotation is offered. Readers should treat the attribution as reported commentary pending that source, and the mechanism as a hypothesis about relative supply pressure, assessable.

A borrower that must build, because a rival is building, is a price-insensitive supplier of bonds. Price-insensitive supply in sovereign scale must be absorbed by price-sensitive buyers, at a yield inducement. If the marginal buyer chooses between a Treasury and a hyperscaler bond of similar maturity, the two yields cannot move independently. The sovereign faces a higher cost of issuing and refinancing its stock, which feeds deficits and increases sovereign supply — a fiscal-financial loop in which AI capital and the state are no longer borrower and backdrop, but competitors.

Strange (1986) and Frieden (1991) argued credit creation is a form of power: who may borrow, on what terms, decides what gets built. Where a few private issuers can move the sovereign's yield curve, the build-out's cost of capital is partly endogenous to its scale — and geopolitics (Sections 3 and 4) can tighten it from outside. Meta's October 2026 release describes that competition: a balanced mix of operating cash flow, investment-grade debt, strategic partnerships and complementary equity, competing with sovereign issuance for the same buyers and raising the yield inducement — not evidence that debt markets have closed.

7. From geopolitics to a break-even rate

This section converts the chain into the finance companion's break-even condition. Its claim stack — debt, leases and committed obligations as assembled there — is $457.1 billion. Table 1 asks the annual interest that stack carries at different blended rates, and the incremental revenue required to cover the interest alone at a 34 per cent margin. Principal repayment, capital expenditure and operating costs are excluded; the table is a floor, not a forecast.

Each 100 basis points on the rate adds $4.57 billion of annual interest — 1 per cent of $457.1 billion — and about $13.4 billion of required annual revenue at a 34 per cent margin. Two hundred basis points of risk-free and supply-driven repricing (Sections 3 to 6) adds $9.1 billion of annual interest and roughly $26.9 billion of required revenue, while the stack is outstanding.

The 34 per cent margin is exposed to the competition the technology companion examines: if price competition among equivalent AI capacity compresses the incremental margin to 25 per cent, every required-revenue figure in Table 1 rises by a factor of about 1.36; at 20 per cent, by about 1.70. Geopolitical tightening and competitive margin compression are not alternative stories: they multiply.

7.1. The chain, stated once

Figure 1 restates the sequence of markets through which a shock must pass; each arrow is a claim defended in the section noted, and in a crisis the arrows move together.

The chain does not claim that a Hormuz disruption must happen, that rates must rise by any particular amount, or that any borrower must fail. It claims something weaker: that the borrower's economics contain a priced exposure to a class of events outside its control, large relative to the revenue the assets produce, and invisible in any analysis which holds the financing rate constant at its announcement level. If a geopolitical shock lifts oil-driven inflation expectations and a central bank responds by holding policy rates high or delaying cuts, a borrower financing in tranches meets a higher hurdle at each new tranche: an exposure mechanism, not a prediction that dearer oil compels a rate rise or that any borrower must fail.

8. Oracle as a case: exposure on all three channels at once

Oracle shows that the mechanism has somewhere concrete to bite; the figures are those established in the finance companion from company filings, not new estimates. The firm has assembled a claim stack of $457.1 billion, spent $55.7 billion of capital expenditure in fiscal 2026, and is priced by the credit market, on the companion paper's snapshot, at a five-year CDS spread of 261 basis points.

Through the *geopolitical-rates channel*, a Hormuz-driven tightening raises the risk-free leg under the entire stack and under the larger volume of financing yet to be raised; the firm's own performance cannot offset a movement in the base rate. Through the *liquidity channel*, the firm refinances and issues in the same crowded calendar as the sovereign and its largest customers and competitors, several issuing for the same purpose. Through the *Treasury-competition channel*, the sector's aggregate issuance — to which this borrower contributes — pressures the very yield curve from which its spreads are measured.

The case also clarifies what would falsify the paper's concern for this borrower. If AI revenue grows fast enough, for long enough, at margins at or above those assumed in Table 1, the higher hurdle is cleared and the geopolitical exposure, though real, is affordable. The finance companion examines the firm's revenue, backlog and margins against that test; this paper takes no position on its outcome. What the case establishes is the structure of the bet: a leveraged claim on future AI revenue whose carrying cost is set, in part, by oil markets, a central bank's reaction function and the Treasury's auction calendar.

A single borrower's difficulties would not prove the macro chain, any more than its prosperity would refute it: the chain stands or falls on the market-level evidence of Sections 3 to 6.

9. Energy geography and the politics of siting

Power is not a commodity a data centre buys. It is a place, a queue and a permission, and each of the three is political. Grid capacity is abundant in some regions and exhausted in others; connection is approved in sequence — interconnection studies, the utility, state and local permitting, environmental review and equipment supply — with later stages re-run when projects withdraw. A site can be legally available, commercially contracted and still years from being energised.

The scale of the queue is documented in the Lawrence Berkeley National Laboratory interconnection study, *Queued Up: 2026 Edition* (Rand et al. 2026): about 2,061 gigawatts were in United States queues at the end of 2025, only about 13 per cent of capacity applied for between 2000 and 2020 was built by end-2025, about 75 per cent was withdrawn, and projects completed in 2025 had spent a median of more than five years (about 61 months) in queue.

Local opposition is part of the same mechanism. Objections across contested projects concern electricity bills, water use, noise, land and green space, backup-generator testing and feared effects on house prices. The political point does not depend on resolving that dispute: a feared effect, believed by residents and councillors, delays or defeats permits as effectively as a measured one. Bypassing the grid does not bypass this politics; it moves it: on-site gas generation and fuel cells exchange a grid queue for an air permit. Oracle's Project Jupiter in New Mexico in labelled form: a reported design of 2,462 megawatts of Bloom fuel cells, with no delivered megawatts found, an air-permit decision due 23 November 2026 (may slip to 18 December), a public hearing opening 7 December 2026 at Sunland Park, a Green Chile pipeline in-service date of 1 February 2027 and a force majeure notice reported 24 September 2026 — evidence about permitting time, not verified performance.

China: where United States connection is fragmented across operators, utilities and jurisdictions, Chinese grid expansion is state-coordinated at national scale, and no directly comparable verified series is claimed. Competition over AI capacity is in part a competition over energisation speed, which chip export controls do not settle. Horizon: orbital and subsea data centres, examined in the technology companion, remain at demonstration scale and are not a near-term substitute for terrestrial capacity.

10. Sovereign clouds, neoclouds and who controls compute

Sections 5 and 6 assumed a single pool of capital and a recognisable class of borrower; the supply side is changing both. Capacity is now built simultaneously along four lines: hyperscalers building and owning; hyperscalers leasing campuses built by specialist landlords; joint ventures in which an asset owner holds most of a project that a hyperscaler then leases back, of which Meta's Hyperion structure (Blue Owl holding 80 per cent, as reported) is the example used in the companion papers; and neoclouds — specialised GPU cloud operators such as CoreWeave, Nebius and Crusoe — which rent raw accelerator capacity by the hour. In every case, contracted is not energised (Section 9).

For political economy, the significance is control, not valuation. A market in which hyperscalers own the capacity, a second in which specialist firms own the accelerators and rent them to anyone, and a third in which asset managers own the buildings, allocate power over the same scarce inputs — chips, power connections and long-term customer contracts — to different hands, within different states' reach. Neoclouds weaken hyperscaler gatekeeping over compute while deepening dependence on the same chip supplier and grid queues. The price consequences of the simultaneous build — four supply lines arriving against one demand curve, and the margin compression that follows if demand disappoints — are modelled in the finance companion.

Sovereign AI is that fragmentation, stated as policy. States increasingly require data localisation, sponsor sovereign clouds, screen foreign ownership of compute and treat domestic accelerator capacity as a strategic asset — dividing what was becoming a single global pool of capacity into national and bloc pools, duplicating fixed costs, and lowering the utilisation across which those costs can be spread. A geopolitical shock no longer meets one integrated market for compute, but several partially separated markets, each with its own subsidised supply, queue and claim on the same credit markets.

Anthropic illustrates the new allocation. The frontier laboratory's Form S-1, filed on 26 October 2026, is not priced and contains no price range; the widely reported valuation of $1.5 to 2 trillion is press reporting, not a filing figure. Commitments as reported from that filing, summarised in the companion finance paper — of the order of $111 billion to Google, $110 billion to Amazon Web Services and $31 billion to Microsoft — lock that laboratory's compute to three hyperscalers, and not to Oracle. The political-economy point, with no valuation conclusion drawn, is that the newest large buyer of compute has allocated its demand inside one bloc of suppliers — the kind of concentration that makes aggregate AI capital a question of industrial structure and state policy, not only of one firm's backlog.

11. War, resilience and the responsibility of the state

A shock can reach the asset directly. Energy and digital infrastructure have been reported as targets in the conflicts of 2026, and a concentrated campus — one substation, one water source, one fibre entry — is a legible target in a way that a distributed network is not. This paper makes no claim about any specific incident beyond that reported pattern, and no cost figures are invented for it: no audited series for per-megawatt hardening, redundancy or war-insurance cost was found, and none is estimated here.

The division of responsibility can be stated. Redundancy, geographic dispersion, hardening, backup generation and duplicate network entry are company decisions, raising the fixed-cost stack whose break-even arithmetic is set out in Section 7 and Online Appendix A. Air and missile defence is not in that category: it is a function of the state; a private operator does not buy an air-defence network, and describing resilience spending as if it could substitute for one confuses two goods. Insurance does not close the gap: standard property and business-interruption policies typically exclude war, so a residual risk that is neither insurable nor privately defensible sits with the operator and the host state. For a government promoting domestic AI capacity as a strategic asset, that residual is a contingent public liability — and, for a reader of the finance companion, a reason the survival term in Online Appendix A cannot be set to one.

12. Policy-rate scenarios: a cross-reference, not a second model

The interest-rate scenario machinery is built in the finance companion and summarised in Online Appendix A. **As of 29 October 2026.** On 28 October 2026 the Federal Open Market Committee voted 10–2 to maintain the target range for the federal funds rate at 3-3/4 to 4 per cent, continuing its policy of maintaining ample reserves (Board of Governors 2026b). The statement records activity expanding at a strong pace, uncertainty elevated owing in part to geopolitical developments, and inflation remains elevated; the Committee will deliver price stability. Lisa D. Cook and Beth M. Hammack dissented, preferring an increase of 1/4 percentage point. The ten-year Treasury yield closed at 4.99 per cent on 27 October 2026 (Board of Governors 2026a). The Implementation Note held interest on reserve balances at 3.90 per cent from 29 October, standing repo at 4.0 per cent, reverse repo at 3.75 per cent ($160 billion per counterparty) and primary credit at 4.0 per cent. Treasury par yields for 28 October, mirror-reported from the Treasury's daily curve (the Treasury page was not live-read here), were 5.05 per cent at ten years, 4.67 per cent at two and 5.37 per cent at thirty, labelled reported, not primary-verified. Interpretation, not the statement's language: this is a pause, not a pivot, and the dissent is the position of this paper's December +25 basis point scenario, which remains hypothetical, not a decision or a forecast. The next meeting, on 8–9 December 2026, carries a Summary of Economic Projections; its outcome is not predicted here. Cumulative +50 and +100 basis point scenarios are treated likewise.

Their arithmetic, at the 6.5 per cent centre rate, is set out in Appendix A.3: a rate path alone does not break the case. What breaks it is the combination: the same geopolitical shock that moves the policy rate also widens spreads, tightens the issuance calendar (Sections 5 and 6), delays energisation (Section 9) and, through competition (Section 10), compresses the margin.

13. The rivalry underneath: Saudi Arabia, Iran and the structure of Gulf insecurity

The Hormuz scenario of Section 3 sits inside a regional rivalry, not a bilateral contingency. The Council on Foreign Relations describes the Saudi–Iranian contest as a 'Middle East Cold War': two states that rarely fight directly but compete for regional primacy through proxies and partners — Hezbollah in Lebanon, militias and parties in Iraq, the Assad government and its opponents in Syria, the Houthis and their adversaries in Yemen, and the opposition and government in Bahrain. Riyadh's reading is encirclement by Iranian proxy depth on its borders; Tehran's mirror reading is a Saudi Arabia armed and underwritten by the United States. Both are stated as their holders state them; the transmission runs on whichever prevails in a crisis.

First, sect is the vocabulary of the rivalry more often than its cause: Iran, a Shia theocracy, has backed the Sunni Hamas and a secular Alawite-led government in Syria — difficult to explain on sectarian grounds, straightforward on strategic ones. Second, the rivalry has an oil-market expression: in the 2016 OPEC episode, Saudi market-share defence against shale and a returning post-sanctions Iran helped produce a price collapse that damaged both treasuries.

The rivalry has an off-ramp, and its limits matter for the model of Section 17. On 10 March 2023, in Beijing, Iran and Saudi Arabia agreed to restore diplomatic relations, reopen embassies and reactivate earlier security and economic agreements. The text is normalisation, not alliance: it contains no Chinese mutual-defence guarantee, and it leaves Saudi red lines — on Yemen, on Iranian missiles, and on an Iranian nuclear weapon — unresolved. The rapprochement lowers the baseline probability of confrontation without removing the tail.

14. Israel, survival security and the domestic-politics question

Israel's part in the escalation chain follows from a doctrine, not only from a government. The Begin Doctrine, named for the prime minister who ordered the 1981 Osirak strike and restated before the 2007 al-Kibar strike, holds that Israel will not permit a regional adversary to acquire nuclear weapons and will act preventively if diplomacy and deterrence fail. On 13 June 2025, opening Operation Rising Lion, Prime Minister Benjamin Netanyahu framed the strikes, describing Iran's programme as a threat to Israel's 'very survival'; the Israel Defense Forces presented the operation as pre-emption against an enrichment and weaponisation effort approaching a point of no return. Tehran's framing is the opposite: the strikes were unprovoked aggression against a state whose programme it maintains is peaceful, and its retaliation a lawful act of self-defence. Third-party legal assessments remain contested on both the anticipatory self-defence claim and proportionality. Three framings are recorded because the bargaining model of Section 17 must enter each side's reservation values as that side states them.

A further claim is that the timing is explained by Netanyahu's domestic politics — coalition fragility, corruption proceedings, the rally effect of crisis. This paper labels it a hypothesis: no primary evidence in the record assembled here establishes domestic causation for the June 2025 decision, which is fully consistent with the stated doctrine and with the International Atomic Energy Agency (IAEA) finding summarised in Section 16. Domestic incentive enters the model of Section 17 as a variable to be specified and tested, not as an explanation already proven.

14.1. The religious-political apex, and why decapitation is a different shock

Iran is not a state with a religious establishment beside it; under *velayat-e faqih* (guardianship of the jurist), the Supreme Leader is the apex of both. He appoints the head of the judiciary and of state broadcasting, six of the twelve Guardian Council members, and holds supreme command of the armed forces including the Islamic Revolutionary Guard Corps (IRGC), inside a vetting loop in which the Guardian Council screens candidates for the Assembly of Experts that nominally selects him. A strike on that apex is not a strike on a general staff: it attacks the node where religious legitimacy, candidate selection and military command meet, in a political culture where martyrdom is an organising narrative. Decapitation carries the largest coefficient in the model of Section 17, and the record of who was, and was not, killed must be stated exactly, as the next section does.

15. A necessary correction of the record: who was killed, and when

> **Fact-check, date-stamped.** Three episodes are routinely conflated, including in early drafts of the argument this paper formalises. They must be kept apart. (i) On 3 January 2020 a United States strike in Baghdad killed Qasem Soleimani, commander of the IRGC Quds Force. Soleimani was a senior military commander, not Iran's religious leader; Supreme Leader Ali Khamenei was not the target and was not killed. (ii) On 13 June 2025 Israel's Operation Rising Lion killed, among others, IRGC commander Hossein Salami, Armed Forces chief of staff Mohammad Bagheri, Khatam al-Anbiya commander Gholam Ali Rashid, IRGC aerospace chief Amir Ali Hajizadeh, and nuclear scientists Fereydoon Abbasi and Mohammad Mehdi Tehranchi. Khamenei survived: he issued statements during the exchange and appointed successors to the officers killed. A reported Israeli plan to kill Khamenei in June 2025, said by unnamed United States officials to have been vetoed by President Trump, is not officially confirmed; Defence Minister Israel Katz later said Israel would have struck him had it been able to locate him. (iii) Multiple secondary sources in 2026 report that Khamenei *was* killed on 28 February 2026 in joint United States–Israeli strikes, with Iranian confirmation reported on 1 March 2026. This paper cites that episode only as widely reported with reported confirmation: the primary Iranian or United States confirmation has not been retrieved into this paper's record, and it is listed as unverified in Section 21. Casualty totals for all three episodes conflict across sources, and no single total is relied on anywhere in this paper.

If the decapitation variable of Section 17 is coded from a conflated record — a 2020 military killing misread as a religious-leader assassination, or a June 2025 survival misread as a death — it measures journalism, not events. The February 2026 reports, if confirmed by primary sources, would be the apex shock Section 14.1 describes: succession inside a theocratic vetting loop, under wartime conditions, with martyrdom mobilisation raising any successor's cost of de-escalation. The mechanism is specified so it can be applied once the primary record is retrieved; it is not applied yet.

Strike details are cited as reported; the source note in Section 18 lists what is primary-verified.

The scenario of Sections 3 and 18 is a *sustained* disruption, not the June 2025 exchange, which markets absorbed within days, and not any single reported strike: escalation and disruption risk, not an accomplished closure, is what the evidence supports modelling.

16. The nuclear file and the Iranian domestic economy

The United States withdrew from the Joint Comprehensive Plan of Action on 8 May 2018; Iran thereafter expanded enrichment beyond the agreement's limits. On 12 June 2025 the IAEA Board of Governors (resolution GOV/2025/38, adopted 19 to 3, with 11 abstentions and 2 members not participating) found Iran in non-compliance with its safeguards obligations. The Agency's reporting around the June strikes (GOV/2025/65) put Iran's total enriched uranium stockpile at 9,874.9 kg as of 13 June 2025, of which 440.9 kg was enriched to 60 per cent U-235. Sixty per cent is not weapons-grade, which is conventionally around 90 per cent; and breakout — producing enough fissile material for one device — is not weaponisation, which additionally requires a warhead and delivery integration. After the June 2025 strikes the Agency reported a loss of continuity of knowledge over parts of the programme, which compresses diplomatic warning time even where physical capacity was damaged.

Inside Iran, the same structure frames household welfare. The World Bank describes a lost decade in which poverty rose from 20 per cent in 2011 to 28.1 per cent in 2020, with roughly ten million more Iranians falling into poverty. Oil's measured weight is smaller than the debate suggests — about 8.6 per cent of GDP over April–December 2023 — and three shares must not be conflated: oil's share of GDP, of exports and of government revenue are different quantities, of which the government-revenue share is unverified in this record and is flagged as a gap, not estimated. Documented are prices and the currency: inflation above 40 per cent for several years, around 42–43 per cent in 2025, a rial near 1.42 million per dollar in December 2025, and Tehran bazaar protests that month beginning among currency and gold traders. This is the domestic constraint on Tehran's bargaining position: a leadership that can absorb external pressure longer than its households can absorb the price level, but not indefinitely.

17. Framing the unquantifiable: a war-probability model with stated limits

Geopolitics is quantifiable as a *model* — outputs only as good as assumptions — not as a measurement. The formal literature supplies the frame. In Fearon's (1995) bargaining model, war is ex post inefficient, so a bargaining range exists: agreements in the interval $(p - c_A,\ p + c_B)$, where p is the expected outcome of fighting and c is each side's cost of war; war occurs when private information, commitment problems or indivisible stakes collapse that range. Bueno de Mesquita et al. (2003) and Powell (2006) on commitment problems make the same point: war probability rises as the bargaining range narrows, and the variables of Sections 13 to 16 — enrichment proximity shrinking the time over which commitments can hold, decapitation raising audience costs, alliance signalling, sanctions raising the cost of the status quo — are range-narrowing forces. Event-data systems such as the Integrated Conflict Early Warning System (ICEWS) and the Global Database of Events, Language, and Tone (GDELT) attempt such estimates from news flows, with modest precision and recall, media bias and reflexivity; the formalism structures scenarios, not forecasts.

$$P(\text{war within horizon}) = 1 / (1 + e^{-z}), \quad z = -2.20 + 0.90\cdot\text{oil} + 1.10\cdot\text{enrichment} + 1.40\cdot\text{decapitation} + 0.70\cdot\text{alliance} + 0.50\cdot\text{sanctions}$$

Each coefficient's sign, if not its size, is fixed by theory before scoring: every range-narrowing force enters positively. The model is falsifiable in one way only: pre-register codings and threshold before the horizon opens, score out of sample against the base rate (Brier score, precision–recall), and prohibit retroactive recoding. A specification that cannot beat the base rate out of sample is discarded. Table 2 does one job: the mechanisms compound multiplicatively, so a world where several are active is a different regime.

18. Transmission in a sustained-disruption scenario

The June 2025 exchange is the calibration point. About 20 million barrels per day move through Hormuz — roughly 20 per cent of global petroleum-liquids consumption and 25 per cent of seaborne oil trade; the two percentages have different denominators and are not interchangeable (U.S. Energy Information Administration 2024). Bypass capacity is only about 3.5–5.5 million barrels per day. As reported, in June 2025 Brent rose from about $64 to about $81 and fell back below $70 after the 24 June ceasefire: Iran's parliament voted to close Hormuz and Iran never did. Analyst scenarios for a *sustained* closure — $100–130, and $130–150 or more for four to eight weeks — are analyst scenarios, not an EIA consensus.

A wider war would also target the newest layer of Gulf infrastructure. Gulf data-centre exposure, as reported, is material — Stargate UAE at 1 GW initial within a 5 GW planned project, Khazna above 1 GW by 2030, Saudi plans of about 1.5 GW by 2030 — and fragile: temperatures above 45°C, seawater or liquid-cooling dependence, concentrated grid connections. This is exposure, not damage: no verified damage to Gulf data centres in June 2025 is in this record, and none is claimed. The cable channel has a documented precedent: in February 2024 the AAE-1, SEACOM and EIG systems were cut in the Red Sea, as reported, most plausibly by the anchor of the sinking *Rubymar*; AAE-1 was restored only in July 2024, about five months later. Repair time, not the cut, is the binding constraint.

A correction to a popular version of the chain. War does bid for gold — on 13 June 2025 gold reportedly rose 1.30 per cent to $3,428.10 an ounce, near its April record around $3,500 — but the Federal Reserve does not raise rates to suppress gold, and gold is not in its mandate or its reaction function. The corrected chain is: war raises oil; oil raises inflation expectations; if expectations threaten to de-anchor, the Federal Reserve holds rates high or delays cuts — it hikes only in that de-anchoring case; real and discount rates stay elevated; and the financing cost of AI capital rises. Gold can rise *alongside* high rates in a stagflationary configuration: the safe-haven bid and the real-rate drag operate. The chain's falsifiers are observable: a new nuclear agreement, restored IAEA access, normal Hormuz transit volumes, a falling oil risk premium and narrowing spreads would each code the variables of Section 17 down, and the break-even exposure of Section 7 with them.

Source note. Primary-verified are the IAEA resolutions and stockpile reporting (IAEA 2025a, 2025b), the Federal Reserve statements and H.15 (Board of Governors 2026a, 2026b), EIA chokepoint volumes (U.S. Energy Information Administration 2024), World Bank welfare figures (World Bank 2025) and LBNL queue data (Rand et al. 2026). Cited as reported only are the strike details of Sections 14–15, the February 2026 killing, cable causation, Brent and gold prices, analyst scenarios and Gulf capacity plans. Also reported only, via Reuters pending verification against the Bureau of Economic Analysis release itself: the figures released on 29 October 2026: advance third-quarter GDP growth of 3.4 per cent, September PCE inflation of 3.6 per cent year on year, core PCE of 3.1 per cent and initial claims of 194,000; no new Oracle filing, IAEA or LBNL release was found for that date.

19. Two regimes of growth: liquidity, population, and the support ratio

The transmission chain of Sections 3 to 12 moves the interest rate. This section asks the prior question: whether the real side can grow into the claims at any rate. Two regimes of growth must be distinguished, because the debate conflates them. In the first, liquidity, the summary of the monetary literature prepared for the Congressional Research Service finds faster money and credit growth largely dissipated in inflation, with little lasting effect on real output, the output-gap effect disappearing after about five years. Bank for International Settlements Working Paper No. 569 locates the constraint on the other side of the cycle: a debt-service gap above equilibrium depresses subsequent output growth, and policy transmits through the financial cycle and the collateral loop, not through money quantities alone. Liquidity can finance a build-out and inflate asset prices; it cannot by itself create the real income that must service the claims.

In the second regime, long-run real growth, the arithmetic is demographic. In the Solow steady state, output per worker grows only at the rate of technical progress g; population growth n lowers steady-state output per worker through capital dilution, and a higher saving rate s raises its level without raising its long-run growth rate. The support-ratio identities state the same point in accounting form: $Y/N = (L/N) \times (Y/L)$, and $C/N = (L/N) \times [(1 - s)(Y/L)]$, where L/N is the share of workers in the population. East and Southeast Asia, 1970–2000, illustrate the magnitudes: income per effective consumer grew at 4.32 per cent a year, the first demographic dividend contributing 0.59 percentage points and the second 1.31. Headcount alone is not growth per person; what matters is the support ratio L/N and productivity Y/L.

The demographic facts, corrected where the popular record is wrong, are these. The United Nations *World Population Prospects 2024* puts world population at 8.2 billion in 2024, peaking at about 10.3 billion in the mid-2080s and standing at 10.2 billion in 2100, with an 80 per cent probability that the peak occurs within this century; 63 countries, holding 28 per cent of world population, have already peaked. The post-war expansion is misremembered: total fertility was 5.0 in 1950, and world population rose from 2.5 billion in 1950 to 6.2 billion in 2000, driven by mortality decline with sustained high fertility, not mainly by the Western baby boom (Bloom 2025). The claim that the post-1945 world stopped fighting is overstated: there has been no Second-World-War-scale war, but many wars continued, and the expansion happened alongside them, not because war disappeared.

Country corrections matter more. China: 1.41 billion in 2024, total fertility about 1.0, population falling since 2022, projected at 1.21 billion in 2054 and 633 million in 2100; National Bureau of Statistics figures carried in secondary reporting put end-2025 population at 1,404.89 million, down 3.39 million (births 7.92 million, deaths 11.31 million). South Korea: just under 52 million in 2024, total fertility from 0.72 in 2023 to 0.75 in 2024 and 0.80 in 2025 on preliminary figures, with natural decrease since 2020. India runs the other way: 1.45 billion in 2024 and total fertility of about 1.96, below replacement, but a population still growing on momentum to 1.69 billion in 2054, peaking at about 1.7 billion in the early 2060s and standing at about 1.5 billion in 2100. Any claim that India's population is falling is false. China and India together are 2.86 billion of the world's 8.2 billion — 34.9 per cent, about one-third, not one-half.

Ukraine's population is about 10 million lower than in 2014 and about 8 million lower than in February 2022, with total fertility about 1.0 and about 6.7 million refugees (UNFPA, October 2024); decline predated the war, and war-time dips combine tempo and migration effects, not automatically permanent. A national war dip is not evidence of incipient global decline.

The *State of Food Security and Nutrition in the World* (2025) counts 673 million people — 8.2 per cent — facing hunger in 2024, down from 8.5 per cent in 2023 and 8.7 per cent in 2022, Africa at 307 million and 20.2 per cent. The *Global Report on Food Crises* (2026) counts about 266 million people in IPC Phase 3 or above across 47 countries in 2025 — 22.9 per cent of the analysed population — with conflict at about 147 million and climate extremes at about 87.5 million, and famine (Phase 5) in Gaza and Sudan within a Phase 5 total of 1.4 million. These are acute, concentrated crises, not a global physical shortage: world cereal supply was a record, about 2,961 million tonnes on the September 2025 forecast; the FAO Food Price Index stood at 124.3 in December 2025, 22.4 per cent below its March 2022 peak; and the cereal index averaged 107.9 in 2025, down 4.9 per cent and the lowest since 2020. Heat is a supply risk: yields decline above about 30°C, and the IPCC-summarised estimates of Zhao et al. put the loss per degree Celsius at 6.0 per cent for wheat, 3.2 per cent for rice, 7.4 per cent for maize and 3.1 per cent for soybean, without carbon-dioxide fertilisation or adaptation, within a projected range of 4 to 10 per cent for maize and wheat — but futures prices are expectations and risk premia, not realised output.

Whether AI relaxes or tightens the support ratio is open, and the evidence must be stated at the strength it has. Exposure is not displacement: the IMF puts about 40 per cent of global jobs exposed to AI, and about 60 per cent in advanced economies (Cazzaniga et al. 2024); the ILO (2025) finds about 24 per cent of employment with some generative-AI exposure and only 3.3 per cent in the highest band, with clerical work most exposed and transformation more likely than replacement; the OECD (2023) puts 27 per cent of jobs at high automation risk. The World Economic Forum (2025) expects 92 million roles displaced and 170 million created by 2030 — a net of 78 million, an expectation not an outcome — with 59 of every 100 workers needing reskilling, 11 of 100 unlikely to receive it. Early displacement evidence is narrow and generational: Brynjolfsson, Chandar and Chen (2025) find US employment at ages 22 to 25 in the most exposed occupations down 13 per cent in relative terms, junior developers down about 20 per cent. Augmentation evidence is real: call-centre productivity up 14 per cent, 34 per cent for novices (Brynjolfsson, Li and Raymond 2023); writing time down 40 per cent, quality up 18 per cent (Noy and Zhang 2023). Acemoglu's aggregate estimate — GDP higher by 1.1 to 1.6 per cent over ten years (Acemoglu 2024) — sits far below optimistic claims of about 7 per cent productivity gains. Adoption is generational: the OECD (2025) finds use declining steeply with age — under 20 per cent of over-55s in formal AI training against 46 per cent of 26- to 35-year-olds. There is no established evidence that AI raises population fertility. And the proposition that AI-driven mass unemployment lowers fertility is a hypothesis: no AI-specific causal chain has been established, and no global dollar cost of retraining was found in the sources assembled here.

Precision-agriculture gains are modelled NENA scenarios by 2050 — fertiliser up 18 per cent, planting 13 per cent, spraying 4 per cent, irrigation 10 per cent, in efficiency terms — not a globally observed figure in tonnes per hectare; yields are stated in t/ha (1 kg/mu = 0.015 t/ha). Against them stand the build-out's resource claims (International Energy Agency 2025): data centres used about 415 TWh in 2024 (about 1.5 per cent of world electricity), 485 TWh in 2025, about 945 to 950 TWh in 2030 (base case); water was about 560 billion litres in 2023 — 373 indirect, 140 direct, 47 manufacturing — projected at about 1,200 billion litres in 2030; a hyperscale AI site above 100 MW loads like about 100,000 households. A net saving exists only where farm-attributable savings exceed AI-attributable energy and water for that service in that place, without rebound. One correction is absolute: no source here establishes that data centres lower crop yields directly, and this paper does not claim it; the framing is local competition and opportunity cost.

The implication is a bound, not a new channel: Sections 3 to 12 show how a shock moves the financing rate; this section shows what it must be serviced from — a support ratio and a productivity path liquidity cannot manufacture. If L/N falls while claims are fixed, Online Appendix A's break-even conditions tighten.

20. Implications

20.1. Corporate governance: the financing assumption is a board decision

If the break-even rate on a multi-decade build-out can be moved by geopolitics, the financing-rate assumption in an investment case is not a treasury detail; it is one of the decision's central judgments, and boards should treat it as such. That means stress-testing committed capital expenditure against financing rates several hundred basis points above the rate at signing, disclosing the stack of obligations in a form that lets outsiders run the same test, and distinguishing between contracted revenue and revenue that must exist, at an assumed margin, to cover the carrying cost. The governance failure implied is rarely dishonesty; it is approving a twenty-year obligation on a point estimate of a variable with a geopolitical distribution.

20.2. Energy security: a chokepoint inside the technology budget

Energy security analysis has counted its stakes in fuel costs, competitiveness and the trade balance. The AI build-out adds a new line to that ledger: the capital cost of the decade's defining infrastructure programme is partially indexed, through the inflation and policy channel, to one strait's security. Strategic reserves, naval escort capacity, diversification of supply routes and diplomatic de-escalation have a return that does not appear in energy accounts.

20.3. Monetary policy: tightening into a build-out

An oil shock argues for tightening; the largest private investment programme in the economy is dependent on cheap long-term finance. Tightening does its intended work on inflation expectations partly by raising the cost of exactly such programmes. A central bank cannot, and on its mandate should not, exempt AI capital from that transmission. Policymakers and commentators should describe it honestly: in a supply shock, monetary policy does not merely cool demand, it reprices the nation's strategic investments.

20.4. Fiscal issuance: the sovereign as competitor

Treasury debt management has long assumed deep, elastic demand for the sovereign's paper; Sections 5 and 6 qualify that assumption. When a few private issuers place sovereign-scale supply on overlapping calendars, with a stated indifference to price, debt management becomes partly a competition problem: auction tails, maturity choice and issuance timing must reckon with corporate supply as a state variable. Higher yields raise debt-service costs and widen the deficit to be funded.

21. Limitations

Four limitations follow. First, the Hormuz element is a scenario, chosen because the chokepoint is real, the standoff is live and the transmission is well studied; other geopolitical shocks would enter the chain at different points and with different strengths. Second, the quantitative section is break-even arithmetic on one borrower's established stack, not an econometric estimate of the chain's elasticities; estimating them is an empirical project. Third, the reverse-crowding discussion relies in part on reported commentary whose primary sourcing remains incomplete, as flagged in Section 6; the mechanism is stated so that it can be tested against issuance and yield data whether or not the attribution survives verification. Fourth, several load-bearing claims remain unverified in this paper's record and are used only as labelled reports, never as premises: the primary Iranian or United States confirmation of the reported February 2026 killing of the Supreme Leader; the reported Trump veto of June 2025; any damage to Gulf data centres; all casualty totals; and any single estimate of nuclear breakout time in days.

22. Conclusion

The zero-rate era taught investors to treat the cost of capital as a constant and technology risk as the only risk. This paper has argued the opposite for debt-financed AI infrastructure: layered over technology risk is an older risk, the price of money itself, set in markets where oil, a strait in the Gulf, a central bank's mandate and a treasury's auction calendar are first-order forces. A shock that begins with geopolitics passes, link by link, into inflation expectations, the policy rate and a crowded issuance calendar, and emerges as a higher break-even revenue.

On the established figures, that exposure is not marginal: about $4.57 billion of additional annual interest, and about $13.4 billion of additional required annual revenue, for every 100 basis points the geopolitical and supply channels add to a $457.1 billion stack. Geopolitical risk is not the background against which the AI investment story is told; it is one of the variables that decides whether the story's arithmetic closes. Analysts who hold the financing rate fixed are not neutral about geopolitics: they are forecasting it — forecasting that nothing happens — without saying so.

Acknowledgements

Acknowledgements that could identify the author are withheld in this blinded manuscript for double-blind review.

Declaration of interest

The author is a researcher at an independent research group, has no employment or consulting relationship with Oracle and holds no position in ORCL securities. A related paid short report is prepared by other members of the group from this research; the author does not prepare or sell it. This paper contains no investment recommendation, target price, options combination, stop-loss or trading advice.

References

Acemoglu, D. (2024) NBER Working Paper No. 32487. Cambridge, MA: National Bureau of Economic Research.

Adrian, T. and Shin, H. S. (2010) ‘Liquidity and Leverage’, *Journal of Financial Intermediation* 19(3): 418–437.

Anthropic PBC (2026) *Form S-1 Registration Statement*. Filing with the United States Securities and Exchange Commission, 26 October 2026. [Commitments and filing status as reported.]

Bernanke, B. S., Gertler, M. and Watson, M. (1997) ‘Systematic Monetary Policy and the Effects of Oil Price Shocks’, *Brookings Papers on Economic Activity* 1997(1): 91–157.

Blanchard, O. and Gali, J. (2010) ‘The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s?’, in J. Gali and M. J. Gertler (eds) *International Dimensions of Monetary Policy*. Chicago: University of Chicago Press.

Bloom, D. (2025) ‘The Debate over Falling Fertility’, *Finance & Development*. Washington, DC: International Monetary Fund.

Blyth, M. (2013) *Austerity: The History of a Dangerous Idea*. New York: Oxford University Press.

Board of Governors of the Federal Reserve System (2026a) *Federal Open Market Committee Statement, 16 September 2026; Statistical Release H.15*. Washington, DC.

Board of Governors of the Federal Reserve System (2026b) *Federal Open Market Committee Statement, 28 October 2026*. Washington, DC.

Brunnermeier, M. K. and Pedersen, L. H. (2009) ‘Market Liquidity and Funding Liquidity’, *Review of Financial Studies* 22(6): 2201–2238.

Brynjolfsson, E., Chandar, B. and Chen (2025) Working paper on generative AI and early-career employment. Stanford University, August 2025.

Brynjolfsson, E., Li, D. and Raymond, L. (2023) Working paper on generative AI in customer support.

Bueno de Mesquita, B., Smith, A., Siverson, R. M. and Morrow, J. D. (2003) *The Logic of Political Survival*. Cambridge, MA: MIT Press.

Cazzaniga, M. et al. (2024) Staff Discussion Note on generative AI and employment exposure. Washington, DC: International Monetary Fund.

Congressional Research Service (n.d.) Report RL30354. Washington, DC: Library of Congress.

Council on Foreign Relations (2025) ‘The Middle East Cold War: Saudi Arabia and Iran’, *Global Conflict Tracker and Backgrounder series*. New York: Council on Foreign Relations.

Drehmann, M. et al. (n.d.) BIS Working Paper No. 569. Basel: Bank for International Settlements. [Year and title not stated in the record consulted.]

Engen, E. M. and Hubbard, R. G. (2004) ‘Federal Government Debt and Interest Rates’, *NBER Macroeconomics Annual* 19: 83–138.

FAO, IFAD, UNICEF, WFP and WHO (2025) *The State of Food Security and Nutrition in the World 2025*. Rome: FAO.

Fearon, J. D. (1995) ‘Rationalist Explanations for War’, *International Organization* 49(3): 379–414.

Frieden, J. A. (1991) ‘Invested Interests: The Politics of National Economic Policies in a World of Global Finance’, *International Organization* 45(4): 425–451.

Gilpin, R. (1987) *The Political Economy of International Relations*. Princeton: Princeton University Press.

Global Network Against Food Crises (2026) *Global Report on Food Crises 2026*.

References (continued)

Greenwood, R., Hanson, S. G. and Stein, J. C. (2015) 'A Comparative-Advantage Approach to Government Debt Maturity', *Brookings Papers on Economic Activity* 2015: 169–217.

Hamilton, J. D. (1983) 'Oil and the Macroeconomy since World War II', *Journal of Political Economy* 91(2): 228–248.

International Atomic Energy Agency (2025a) *Resolution GOV/2025/38: NPT Safeguards Agreement with the Islamic Republic of Iran*. Vienna: IAEA Board of Governors, 12 June 2025 (adopted 19–3–11, with 2 not participating).

International Atomic Energy Agency (2025b) *Report GOV/2025/65: Verification and Monitoring in the Islamic Republic of Iran*. Vienna: IAEA.

International Energy Agency (2025) *Energy and AI*. Paris: IEA, April 2025.

International Labour Organization (2025) Working Paper No. 140, May 2025 update. Geneva: ILO.

Islamic Republic of Iran, Kingdom of Saudi Arabia and People's Republic of China (2023) *Joint Trilateral Statement*. Beijing, 10 March 2023. [Contains no mutual-defence provision.]

Kilian, L. (2009) 'Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market', *American Economic Review* 99(3): 1053–1069.

Krasner, S. D. (1976) 'State Power and the Structure of International Trade', *World Politics* 28(3): 317–347.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2012) 'The Aggregate Demand for Treasury Debt', *Journal of Political Economy* 120(2): 233–267.

Laubach, T. (2009) 'New Evidence on the Interest Rate Effects of Budget Deficits and Debt', *Journal of the European Economic Association* 7(4): 858–885.

Noy, S. and Zhang, W. (2023) Writing-task experiment with generative AI, *Science*.

OECD (2023) *OECD Employment Outlook 2023*. Paris: OECD.

OECD (2025) Study of generative-AI experience and training by age. Paris: OECD.

Oracle Corporation (2026) *Form 10-K for Fiscal 2026; Form 10-Q for the First Quarter of Fiscal 2027*. Filings with the United States Securities and Exchange Commission.

Powell, R. (2006) 'War as a Commitment Problem', *International Organization* 60(1): 169–203.

Rand, J. et al. (2026) *Queued Up: 2026 Edition*. Berkeley, CA: Lawrence Berkeley National Laboratory, June 2026. DOI: 10.2172/3376147.

Solow, R. M. (1956) 'A Contribution to the Theory of Economic Growth', *Quarterly Journal of Economics* 70(1): 65–94.

Strange, S. (1986) *Casino Capitalism*. Oxford: Basil Blackwell.

Tooze, A. (2018) *Crashed: How a Decade of Financial Crises Changed the World*. New York: Viking.

UNFPA (2024) Ukraine population briefing, October 2024.

Author (2026a) *Debt-Financed AI Expansion and Shareholder Value at Oracle*. Manuscript prepared in the format of The Journal of Finance. [Companion finance paper.]

Author (2026b) *The Adoption-Impairment Paradox*. Manuscript prepared in the format of Research Policy. [Companion technology paper.]

Author (2026c) *Oracle Short Report*. [Paid investor report; cross-cites the three companion papers.]

United Nations (2024) *World Population Prospects 2024*.

United States Energy Information Administration (2024) *World Oil Transit Chokepoints*. Washington, DC: United States Department of Energy.

References (continued)

World Bank (2025) *Iran Economic Monitor* and *World Development Indicators: Iran*. Washington, DC: World Bank. [Oil's share of Iranian government revenue not verified in this record.]

World Economic Forum (2025) *Future of Jobs Report 2025*. Geneva: WEF.

Zhao, C. et al. (2017) Global crop-yield responses to temperature increase, *Proceedings of the National Academy of Sciences*, as summarised in IPCC assessment reporting.

Tables

Table 1. Financing hurdle on a $457.1 billion claim stack, by blended financing rate

Blended financing rate	Annual interest ($B)	Revenue needed to cover interest at 34% margin ($B)
5%	22.9	67.2
6%	27.4	80.7
7%	32.0	94.1
8%	36.6	107.6
9%	41.1	121.0

Notes: Author's calculations from the companion finance paper's established claim stack of $457.1 billion. Interest = stack × rate. Required revenue = interest ÷ 0.34. Figures are rounded; the 7 per cent row reproduces the companion documents' anchor of about $32.0 billion of annual interest and about $94.1 billion of required annual revenue.

Table 2. Illustrative logistic scoring: checked sensitivities

Scenario (variables at 1 unless stated)	z	P
Baseline (all variables 0)	−2.20	0.10
Oil shock only	−1.30	0.21
Enrichment proximity only	−1.10	0.25
Decapitation shock only	−0.80	0.31
Alliance signalling only	−1.50	0.18
Sanctions pressure only	−1.70	0.15
All five at 0.5	0.10	0.53
All five at 1	2.40	0.92

Author's illustrative scoring: coefficients are scenario weights chosen to rank the mechanisms of Sections 13 to 16; they are **not** calibrated or estimated probabilities, and the outputs are model outputs, not facts.

Table A1. Annual profit hurdle and required revenue, claim stack $457.1 billion, centre rate 6.5 per cent ($B per year)

Horizon T	Profit hurdle	Revenue at 34% (illustrative)	Revenue at 25%	Revenue at 20%
2 years	251.1	738.4	1004.3	1255.3
5 years	110.0	323.5	440.0	550.0
10 years	63.6	187.0	254.3	317.9
20 years	41.5	122.0	165.9	207.4

Notes: Author's calculations. Total revenue in fiscal 2026 was $67.4 billion and OCI revenue $18.1 billion, per the filings cited in the finance companion, which is why the two- and five-year rows are presented as arithmetic impossibilities rather than targets. The 6.5 per cent centre rate is an illustrative centre, broadly a risk-free rate of 4 to 5 per cent plus a credit spread in the reported CDS region; it is not a forecast, nor Oracle's weighted average cost of debt.

Figures

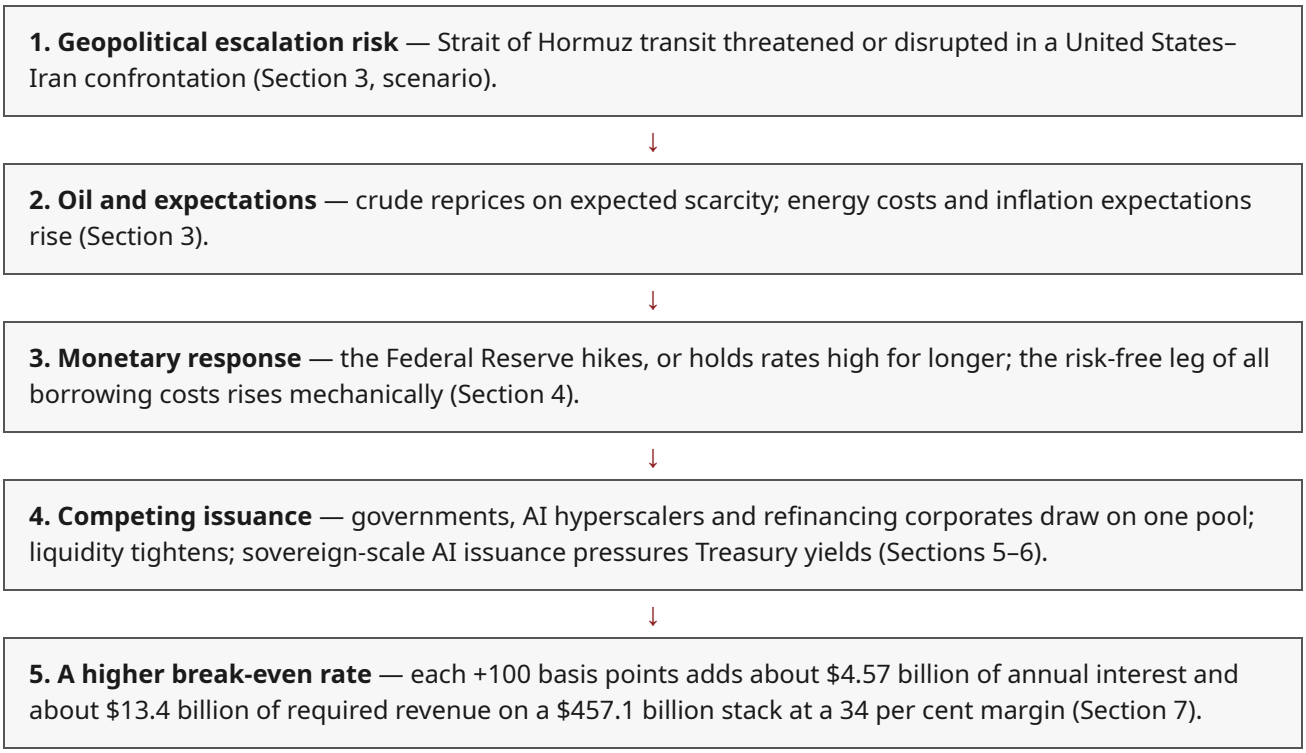

Figure 1. The transmission chain from geopolitical oil risk to the break-even rate on debt-financed AI capital.

Figure captions

Figure 1. The transmission chain from geopolitical oil risk to the break-even rate on debt-financed AI capital.

Online Appendix A. The scenario model in summary: identification, expectation and falsification

This appendix summarises the model the finance companion estimates: the geopolitical chain of Sections 3 to 6 must terminate in a condition that can be wrong, on stated assumptions, or it is rhetoric. Four separations govern every number in all four companion documents. Figures taken from Oracle's filings are labelled filing-verified. Figures taken from reporting or a named assessor are labelled reported estimates. The 34 per cent margin is labelled author-illustrative throughout: Oracle does not disclose a separate Oracle Cloud Infrastructure (OCI) segment gross margin, and no table in any companion document presents 34 per cent as a disclosed figure. Scenario weights are labelled author-assigned; they are not estimated probabilities.

A.1. Formulas

$$A(r, T) = P \times r \div (1 - (1 + r)^{-T}); \quad \text{required revenue} = A \div m$$

$$E[V] = \Sigma\, p_s\, V_s, \quad \Sigma\, p_s = 1; \quad V_s = -P + \Sigma\, CF_t(s) \div (1 + r)^t; \quad S(t) = \exp(-\int \lambda(u)\, du)$$

P is the $457.1 billion claim stack ($123.0 billion debt, $46.0 billion of leases in effect and $288.0 billion of leases not yet effective, as assembled in the finance companion from the fiscal 2026 Form 10-K), r the financing rate, T the horizon, m the gross margin, p_s the weight on state s, $CF_t(s)$ the cash flow in that state, and S(t) the probability that the asset is neither impaired nor destroyed by time t, with λ the combined hazard from obsolescence, substitution, physical attack and stranding. Expected cash flow in year t is $S(t) \times CF_t$. Each proposition has a sign that can fail: required revenue rises with r, with delay and with the RPO haircut, and falls with m and with T; delay is convex in its present-value damage, so a three-year energisation delay destroys more value than a one-year delay three times over. For the credit hazard, $\lambda \approx s \div LGD$ (loss given default) is an approximation, with s read as a reported range — about 141 basis points in January 2026, and reported prints of 288, 295 and 308 basis points on 15, 19 and 21 October, with a weekly summary reported up to 343 basis points — and not as a live quote. The 261 basis point figure in Section 8 is the companion papers' earlier snapshot, retained as such.

A.2. Assumptions, in one place

Remaining performance obligations (RPO) are $455.0 billion at 31 August 2025, $638.0 billion at fiscal year-end 31 May 2026 and $664.0 billion at 31 August 2026, per the Form 10-K and Form 10-Q; the latest filing's maturity buckets are 13, 37, 34 and 16 per cent. OpenAI's approximately $300 billion, five-year contract from 2027 is a reported estimate (Wall Street Journal reporting of September 2025, relayed in a July 2026 S&P review), equal to about 45.2 per cent of $664.0 billion; the straight-line $60 billion a year is illustrative and is not a contractual payment schedule. Circular financing — Nvidia's commitment, revised from up to $100 billion to a $30 billion equity participation, alongside OpenAI's cloud commitments and cloud providers' purchases of Nvidia systems — is alleged by critics and denied as circular by Nvidia; neither characterisation is treated as proven.

A.3. Horizons: what the horizon does to the hurdle

Three sensitivities, stated in full in the finance companion, follow. *Rates, ten-year horizon:* required revenue at 34 per cent is $184.8 billion at −25 basis points, $187.0 billion at the centre, $189.2 billion at +25, $191.4 billion at +50, $195.9 billion at +100 and $200.4 billion at +150 basis points. The full grid is not reproduced. *RPO haircut:* a 10, 20 or 30 per cent haircut applied only to the OpenAI share lowers total RPO realisation to 95.48, 90.96 or 86.44 per cent and raises the ten-year, 34 per cent required revenue to $195.9, $205.6 or $216.3 billion. *Customer growth:* OpenAI's annualised revenue of about $50 billion, as reported to investors in September 2026, must reach $60 billion if wholly paid to Oracle, $120 billion at a 50 per cent share of revenue, $200 billion at 30 per cent and $300 billion at 20 per cent; from $50 billion that takes 4.8, 7.6 and 9.8 years respectively at 20 per cent growth, or 2.2, 3.4 and 4.4 years at 50 per cent growth, before its other reported commitments, of the order of $1.4 trillion across providers, are paid. Combined states multiply rather than add: at +100 basis points, a 20 per cent margin and a 30 per cent haircut, required revenue is about $385 billion a year, roughly 2.1 times the clean ten-year baseline.

A.4. Expectation, falsification and robustness

Under one illustrative set of author-assigned weights — 20 per cent clean, 35 per cent base, 30 per cent adverse, 15 per cent stress — the expected present value in the finance companion's worked example is $111.2 billion. Under friendlier weights of 40, 30, 20 and 10 per cent it is $129.5 billion; under harsher weights of 10, 25, 35 and 30 per cent it is $96.3 billion. The level is weight-sensitive; the direction is the result: under every weighting examined, delay, haircut and margin compression each lower expected value, and no plausible weighting restores the two- or five-year rows of Table A1 to feasibility.

The model is falsifiable, and its falsifiers are observable: RPO converting materially faster than the disclosed 13 per cent first bucket without new signing; OpenAI revenue sustained above the $60 billion a year its Oracle contract implies, while servicing its other commitments; the financing rate falling back below about 4.5 per cent; a disclosed margin at or above the illustrative 34 per cent for two quarters; energisation delivered on the shell timetable rather than the grid timetable of Section 9; or contracted supply being cancelled rather than built. Robustness re-runs at 5- and 20-year horizons and at 4.5 and 8.0 per cent rates change magnitudes only, not signs. What the model cannot do is assign estimated probabilities to geopolitical states; that limit is why this paper's claims are stated as exposures and break-even conditions, not forecasts.