

The Perfect Storm: Where the Yen, Private Credit, and the AI Trade Actually Connect

A synthesis, built from three threads already tracked separately in this project

August 10, 2026

The premise

This project has independently developed three lines of analysis: a Japanese carry-trade unwind (the yen report), a private-credit stress cycle with a Bear Stearns precedent (Troy Black's prophetic words and this project's macro analysis), and — new to this document — Michael Burry's AI accounting-valuation thesis. Each has been treated as its own risk. The question worth asking is whether they're actually one risk wearing three faces, and whether the connective tissue is coincidence, correlation, or the same balance sheets. This document adds four further historical precedents, checked precisely rather than invoked

loosely — the Washington Mutual Option ARM analogy, the 1998 LTCM/LTCB dual crisis, the specific asset-class breakdown of the 2007-08 unwind, and the 1997-98 subprime residual-interest failures that preceded both — because each sharpens a different part of the mechanism rather than simply restating it. It also traces where the actual leverage and tail risk sit once the sponsoring firms' own balance sheets are ruled out: overwhelmingly insurance companies and, to a smaller and faster-growing degree, banks — and, following further research, sharpens the distinction between "good PIK" (planned AI infrastructure lending) and "bad PIK" (the genuine distress signal, concentrated in the legacy software companies AI is disrupting rather than the companies building it), adds the Nvidia-OpenAI credit-wrap escalation and the Collateralized Chip Obligation instrument, and notes that any resulting damage would bifurcate sharply between well-protected hard-asset lenders and effectively unsecured chip-only lenders.

The answer, checked rather than assumed: the same balance sheets, in specific and verifiable ways. This document lays out the mechanism precisely, credits every figure to a source, and holds the conclusion to the standard the rest of this project uses — a coherent tail scenario, not a forecast.

Leg one: the carry-trade unwind is a rate shock and a risk-asset shock at once

Covered at length in the companion Japan report; the summary relevant here is narrow. A yen carry-trade unwind means Japanese and yen-funded investors closing positions and repatriating capital. Two things happen simultaneously when that occurs, and they're usually described separately when they should be described together:

It is a global rate shock. Capital leaving dollar and other foreign assets to fund yen positions being closed pushes global bond yields up — this is already visible: the US 10-year pushed past 4.67% in May 2026 partly on Japanese selling, and Japan's own 30-year has moved from ~3.0% to above 3.5% inside a year.

It is a "sell risk, buy safety" event. Classic carry-trade unwinds (August 2024 is the recent proof: Nikkei -12.4% in a day) force closure of the highest-beta, most-crowded trades first, because those are the positions funded most cheaply and levered most aggressively. **What's crowded and highest-beta in 2026 is not evenly distributed across the market. It's concentrated in one trade.**

The person now running the US side of the July 31 joint intervention has a documented history of engineering exactly this kind of currency move from the opposite direction, which is worth having on record. Treasury Secretary Scott Bessent, then 29 and running Soros Fund Management's London office, was a key architect of the 1992 "Black Wednesday" sterling short — the insight being that the UK's near-universal variable-rate mortgages made a Bank of England rate hike to defend the pound politically impossible, regardless of what the central bank said it would do. The trade netted Soros Fund Management over \$1 billion. **In 2013, as CIO of Soros Fund Management, Bessent ran the mirror-image trade against the yen specifically** — shorting it after Abe's Abenomics rollout for a confirmed \$1.2 billion profit as the yen weakened from the high-70s to over 100 per dollar. The same person who profited from that currency's collapse thirteen years ago is now the one publicly engineering its strengthening — and his defining professional insight, demonstrated both times, is that a currency's defense is frequently a **psychological operation rather than a purely financial one**: Black Wednesday was won because traders concluded the Bank of England's stated commitment wasn't credible, not because Soros's fund had enough capital to move the market alone. That history is worth weighing when assessing whether the widely-circulated photograph of Bessent's own handwritten

"to do" note — specific enough to name the ticker JPY, small enough in dollar terms to be market-irrelevant on its own — was an accidental leak or a deliberate signal exploiting exactly the mechanism he is on record having used before.

But the two trades taught opposite lessons, and only one of them prepares a person for the position

Bessent is now in. 1992 was a bet *against* a government's stated commitment — Soros's fund profited from the Bank of England's failure. 2013 was a bet *with* a credible, openly telegraphed government intention — Bessent simply rode Abe's own announced policy. Today, Bessent is neither outside actor: he *is* the government whose intention is being telegraphed, which is a position his own history has never actually tested. Both trades also taught him "I can read this system correctly and act on it before others do." Neither taught him what happens when he is the signal everyone else is now trying to read and front-run first — a fundamentally different game, with no guarantee his old edge transfers to it. Markets that later concluded a stated commitment wasn't credible, or piled onto a telegraphed policy with more force and speed than intended, are the same market Bessent himself helped teach a generation of traders how to read. The largest, most reflexive market in the world does not stay manipulated by someone who has already shown everyone exactly how the trick works —

and every failure examined at length elsewhere in this document, from Bear Stearns' own risk models to WaMu's growth assumptions to the Gaussian copula's author, was a sophisticated actor confident he had correctly modeled a system that turned out to be larger and less controllable than his model allowed. There is no particular reason this instance is exempt from that pattern.

A structural "trilemma" argument, offered as one coherent reading of the policy's underlying logic rather than a confirmed motive, helps explain why this pressure campaign exists at all. The US is simultaneously trying to reshore manufacturing (requires a weaker dollar), maintain price stability (requires a stronger dollar), and keep the bond market stable enough to finance a national debt that crossed \$40 trillion the same week as the joint intervention (also generally easier with a strong dollar and low inflation expectations). Achieving any two tends to work against the third. Under this reading, price stability's strong-dollar requirement is the one being sacrificed — but deniably, since an explicit devaluation announcement would itself trigger the bond-market instability the policy is trying to avoid. The mechanical detail already established in this document's Leg Two-and-a-half section — that the Fed funded its share of the intervention by selling **euros rather than dollars** — fits this reading precisely: the yen strengthens, the

dollar weakens in relative terms, and no US official has to say "devaluation" out loud.

Leg two: private credit's floating-rate structure is a direct, mechanical transmission channel — not a metaphor

This is the piece that makes "a bond problem hitting AI infrastructure" a literal claim rather than a loose analogy, and it's worth stating with the precision the sourcing actually supports.

Direct lending — the core of the private credit market — is overwhelmingly floating-rate by design. Per KKR's own investor materials: "senior-secured, floating-rate credit" is the asset class's defining structure, priced as SOFR plus a spread. Per the Alternative Credit Council, direct lending is roughly 44-47% of total private credit AUM. Median BDC spreads run **SOFR + 400-650 basis points** depending on deal quality (Federal Reserve Bank of Boston research, weeks old at time of writing) — meaning all-in borrower rates already sit at 9-12%.

This is not a theoretical vulnerability — it is already mechanically firing. The same Boston Fed research

(published within the last two weeks) found the share of BDC loans on payment-in-kind — meaning borrowers deferring interest because they cannot service it in cash — rose from roughly 6% in early 2022 to about 10% by Q1 2026, and critically, this rise is **broad-based across nearly every industry sampled**, not concentrated in a few distressed sectors. The researchers' own conclusion: *"floating rate borrowers who could service SOFR plus 5 when SOFR sat near zero face a very different burden with SOFR above 4 percent... [this points] to a mechanical cause rather than an idiosyncratic one."*

Translate that into the scenario at hand: a yen-unwind-driven global rate spike does not need to touch a single private credit borrower's business fundamentals to hurt them. It raises SOFR (or its correlated global rate benchmarks), which mechanically raises debt service costs across a floating-rate book that is, per the same research, already showing broad-based cash-flow strain. This is the identical mechanism already documented elsewhere in this project's macro analysis for KKR and Blackstone's BDC-level stress — a rate shock simply turns the dial further on a mechanism already running.

And technology concentration inside that private credit book is already measured, not speculative. The Boston Fed sample found internet and software

companies average about 20% of the median BDC portfolio, with some lenders running more than a third of their book into the sector — "creating correlated exposure to venture-backed and growth-stage borrowers." Private credit's stress and tech-sector concentration are not two separate facts. They are the same fact from two angles.

A necessary refinement, since "AI and private credit" is not one exposure but two nearly opposite ones, and conflating them would misstate where the actual danger sits. PIK usage inside AI infrastructure lending itself is overwhelmingly what the industry distinguishes as "good PIK" — planned from origination, not a distress signal. Building a hyperscale data center takes 18-36 months and produces zero revenue until the facility goes live and a hyperscaler's lease payments begin; lenders structure interest to accrue rather than pay in cash for precisely that window, then roll the loan to cash-pay once revenue starts. Rating agencies treat this as standard project finance rather than a credit warning, because the eventual payer — an investment-grade hyperscaler under a 10-15 year lease — is known in advance.

The genuine distress signal — "bad PIK," toggled mid-loan by amendment rather than planned from day one — concentrates instead in the companies AI is disrupting, not the companies building AI's infrastructure. CNBC reporting (February 2026) states

plainly that software and services companies account for the largest share of PIK loans specifically, driven by generative AI eroding legacy B2B and enterprise SaaS pricing power at the same time floating rates raise their debt service costs — a genuine double bind, distinct from and roughly the opposite of the infrastructure-lending case. One industry estimate puts BDC concentration in these vulnerable, asset-light tech sectors as high as 30-50% of some portfolios, above the Boston Fed's more conservative ~20% median. The Financial Stability Board's May 2026 report on private credit found PIK toggles specifically (as opposed to PIK-from-origination notes) associated with a **1-2 percentage point increase in the probability a loan turns delinquent the following quarter**, on top of a 3% unconditional base rate — and Fitch's April 2026 data put the sector's true default rate, once selective defaults and liability-management workarounds are counted rather than only headline defaults, at roughly **5%, against a headline figure that has stayed below 2%** for years. **The AI-private credit risk, stated precisely: it runs through two nearly opposite channels at once — well-collateralized, hyperscaler-backed construction lending on one side, and asset-light legacy software credit being actively destroyed by the same technology on the other — and only the second is where the FSB's own delinquency data says the danger currently lives.**

Leg two-and-a-half: two 2008-era precedents that sharpen the private-credit mechanism further

Two historical parallels, checked precisely rather than invoked loosely, both bear directly on how PIK-driven stress in private credit could unfold — and both suggest this cycle could be structurally harder to resolve than 2008 was, not easier.

The Option ARM parallel — and the opacity is worse this time

Washington Mutual's "Pick-a-Payment" mortgages let borrowers pay less than full interest, with the shortfall added to principal — negative amortization, mechanically identical to a PIK loan's capitalized interest. The parallel to today's PIK toggle usage is close to line-for-line:

Mechanism	WaMu Option ARM	PIK toggle loan
The "choice"	Pay less than full interest	Defer interest
The shortfall	Added to principal	Capitalized into principal

Mechanism	WaMu Option ARM	PIK toggle loan
	("negative amortization")	
The cap	Disclosed: 110-125% of original balance	Undisclosed, discretionary, loan-by-loan
The reckoning	"Recast" — forced full reamortization at maturity or cap	Loan maturity or refinancing
The accounting distortion	Full contractual interest booked as income though only the minimum was collected	PIK interest booked as taxable income though nothing was collected in cash

WaMu's own 10-K filings disclosed the cap and, critically, the scheduled recast volume in advance — \$2.2bn in 2008, \$7.1bn in 2009 — giving analysts a real, quantifiable trigger to model. Private credit's PIK exposure has no equivalent industry-wide cap or mandated aggregate disclosure. **This makes today's version of the mechanism more, not less, dangerous from a foresight standpoint:** WaMu's problem was

disclosed and mechanical; private credit's is undisclosed and discretionary, meaning the industry can accurately say "nothing systemic here" right up until enough loans mature simultaneously and it's too late to unwind in an orderly way. WaMu failed in September 2008 – still the largest bank failure in US history – with its own chief risk officer disclosing that roughly three-quarters of the Option ARM book was negatively amortizing well before the recast wave hit.

The interest-rate-escalator comparison sharpens this further. Subprime's other classic structure – the 2/28 hybrid ARM – reset from a fixed teaser rate to a much higher floating rate on a fixed calendar date, independent of the borrower's ability to pay; these loans were widely understood to be "designed to be refinanced" before the reset, and defaulted close to mechanically once home-price appreciation stopped making refinancing available. Private credit has two live analogs, both standard market convention rather than edge cases: **PIK-elected spreads typically price 50-100bp above the cash-pay rate**, and **leverage-based pricing grids** raise a borrower's spread automatically as its leverage ratio worsens. Both are *distress-triggered* rather than *calendar-triggered* – arguably more dangerous than the 2/28, because they fire the moment trouble starts rather than waiting for a scheduled date, and can compound with an ongoing PIK deferral on the same loan simultaneously.

And the PIK mechanism has a second, regulatory distortion channel beyond simply masking stress.

BDCs must distribute 90% of taxable income to keep their RIC tax status — and PIK interest counts as taxable income even though no cash was received.

This can force a fund to draw down cash, sell a genuinely performing asset, or tap its own credit line purely to distribute against income that was never real cash — a mechanism with no analog in the Option ARM story, where the bank itself (not federal tax law) controlled the accounting.

LTCM and LTCB — two crises, one month, one mechanism each, reinforcing rather than replacing each other

The 1998 yen move, precisely measured: on October 7, 1998, USD/JPY moved 7% in a single trading day — the largest one-day yen move ever recorded, and per the Bank of England, the largest two-day move since the yen began floating in 1973. Over the eight weeks around Russia's August 1998 default, USD/JPY fell more than 15%. **This was not incidental to LTCM's collapse — it was the mechanism.** LTCM's own book was partly funded through yen carry positions used for emerging-market and fixed-income arbitrage; when Russia's default forced LTCM's exposures to unwind, the yen carry trade unwound with it. The Fed brokered (but did not fund) a \$3.6bn recapitalization

from 14 banks on September 23, 1998 — the first time the NY Fed organized a private hedge-fund rescue.

Sixteen days after the record yen move, a second and unrelated crisis reached its own breaking point: the Long-Term Credit Bank of Japan was nationalized on October 23, 1998 — the first failure of a major Japanese bank since WWII. LTCB's failure had nothing to do with the carry trade; it was the terminal stage of bad debt accumulated through the 1980s bubble, with disclosed non-performing loans reaching ¥7.6 trillion by 1997. It was the last of a rolling, thirteen-month Japanese banking crisis: Hokkaido Takushoku Bank (Nov 17, 1997, the first postwar city-bank failure) and Yamaichi Securities (closed Nov 24, 1997, ¥30tn in client assets) came first; LTCB's sister institution, Nippon Credit Bank, failed two months after LTCB, in December 1998 — meaning two of Japan's three specialized long-term credit banks collapsed within ten weeks of each other.

The lesson for today, stated precisely: two structurally independent stresses — a long-building domestic Japanese balance-sheet problem, and a global leveraged-carry-trade unwind — landed in the same three-week window in 1998, and almost certainly worsened each other without directly causing each other. That is exactly the shape of the risk this document has been describing as a "perfect storm," except 1998 shows it has already happened

once. Today's version of Japan's domestic fragility runs through JGB yields and life-insurer duration risk rather than 1980s bubble real estate — but a slow-building domestic problem reaching its breaking point at the same moment global leverage unwinds is the same structure, not a new one.

The resolution template is also worth holding in view.

LTCB was nationalized, its shareholders wiped to zero, and it was sold in 2000 for roughly \$1.2bn to a foreign private-equity consortium led by Ripplewood Holdings — the first time a Japanese bank came under foreign control, with the government retaining tail risk via a put-back provision on any loan that fell 20%+ in value within three years. **If a major Japanese financial institution cracked under today's version of this stress, this is a real, already-used playbook** — and the closing irony is structural: the same class of firm now sitting at the center of this document's private-credit leg (Western private equity — the direct successors of Ripplewood in spirit) could end up on both sides of a future crisis, a source of stress in the AI/private-credit book and a buyer of distressed Japanese assets if the sovereign leg actually breaks.

What the 2007-08 unwind actually hit, asset class by asset class

Checked precisely rather than described generally, because the specific pattern bears directly on what

"no safe haven" would mean today.

Emerging-market equities: -67% from the October 2007 peak to the November 2008 trough — the largest single figure in the entire episode. **Commodity and carry-linked EM assets specifically:** Brazil's Bovespa, a favorite carry-trade destination via its commodity exposure (Petrobras, Vale, steel), lost half its market value from its May 2008 all-time high; the Reuters-Jefferies CRB commodity index lost 44% (\$280.6bn) from its July 2008 peak. **Developed-market equities as the second wave**, not the first — the Dow's September 2008 losses (2,900 points, the largest cumulative point loss on record at the time) came after EM and commodities had already cracked, the unwind moving from the periphery inward. **Japan itself, in reverse, six months later** — Japanese exports fell 45% from their 2008 peak to the 2009 trough, more than half of that decline landing in the six months after Lehman specifically, as the violently strengthened yen crushed the export competitiveness of the very country whose currency had caused the unwind. Across the full 2007-09 cycle, USD/JPY moved from roughly 124 to 87 — a 30% yen appreciation.

And Treasuries did what a safe haven is supposed to do in 2008 — yields fell, prices rose, in the classic flight-to-quality pattern. **This is precisely what did not fully happen in the August 2024 episode**, where

Treasury yields barely moved despite the equity carnage — a documented anomaly. One analysis of the pattern frames the underlying mechanism better than a list of asset classes can: *"These are not four separate events. They are four projections of the same event onto four different asset classes... the mark-to-market is imposed simultaneously on every asset that was held in that position... That is why correlations spike in exactly the moments when diversification is supposed to help."*

The escalating pattern across three cycles — 1998, 2007-08, 2024 — is itself the finding. Every major systemic financial crisis of the past three decades has run through a violent yen unwind as part of its mechanism, not as a coincidental backdrop, and the safe-haven function of sovereign debt has weakened with each iteration. **If a 2026-27 unwind originates in Japanese and US sovereign debt themselves — the scenario this document's earlier legs describe — there may be no safe haven left to run to at all, which would be a genuine escalation in kind from both 1998 and 2008, not merely in scale.**

Leg two-and-three-quarters: three generations of the same accounting disease, ten years apart each time

The private-credit PIK mechanism and Burry's AI depreciation thesis are not new inventions. Checked against the historical record, they are the third appearance of a single underlying trick that financial regulators have named, in writing, twice before — and each time, the fix was scoped too narrowly to reach the next decade's version of it.

1997-98: residual interests, the first documented case

A full year before the mortgage-side collapse, the same disease hit subprime auto lending. January 1997 — Mercury Finance, then the largest independent subprime auto lender, disclosed a multi-year earnings overstatement. The American Bankruptcy Institute's own trade journal titled its retrospective "**Subprime Auto Finance: The Year of the Bankruptcies**": Jayhawk Acceptance (February 1997, after an unexpected \$15.5M credit-loss charge), First Merchants Acceptance (July 1997), Western Fidelity Funding (August 1997) — and, in early 1998, four more: Reliance Acceptance Group, Search Financial Services, First Enterprise Financial Group, and Keller Financial Services.

The mortgage-side version followed roughly a year later, and the mechanism is documented precisely in SEC filings. **Cityscape Financial Corp's own filings describe "the interest-only and residual mortgage**

securities that the Company receives upon loan sales through securitizations" — booking today's profit against an *assumption* about future prepayment and default rates on loans already sold. Cityscape deferred an interest payment on its subordinated debt on May 1, 1998, suspended loan origination entirely by November 1998, and was in Chapter 11 by early 1999.

Southern Pacific Funding Corporation, a Lake Oswego, Oregon specialty finance company, filed Chapter 11 on October 1, 1998 — six days before the record 7% single-day yen move, three weeks before LTCB's nationalization. SPF's own bankruptcy filings show it ran on four warehouse lines totaling \$1.2 billion: First Union (\$300M), Lehman Brothers (\$400M), **Morgan Stanley (\$400M, structured as a loan and security agreement rather than a repo)**, and Greenwich Capital (\$100M). Shortly before the petition date, SPF received notices of default from First Union, Lehman, and Morgan simultaneously — the warehouse-lending relationship failing all at once, not gradually.

What actually killed any prospect of a sale illustrates the mechanism better than the accounting theory alone can. There was no strategic buyer, because the entire subprime wholesale-broker business model — buying loans from unaffiliated independent mortgage brokers rather than through a proprietary retail franchise — was purely price-driven, with no durable

competitive asset a strategic acquirer would pay a premium for. There was no financial buyer either, because nobody wanted to hold the residual interest assets: by the time of the failure, they were widely regarded as toxic and worth pennies on the dollar — the market having independently arrived at the same conclusion the accounting assumptions had been obscuring. A last attempt to interest commercial banks in just the mortgage-origination platform, separate from the toxic residual book, drew only "tire kickers" and produced no deal. **And the final blow was operational rather than financial: when the warehouse lenders froze the credit lines, the securitization team had to physically seize the loan documents overnight to preserve the lenders' collateral position — an action that, by removing the operational continuity any acquirer would need, ended the sale process outright, faster than the balance-sheet insolvency question could even be resolved.**

Regulators diagnosed the disease by name within a year of watching it happen. The Federal Reserve's own later testimony records: *"1998 through 2002 — five institutions closed due to problems related to subprime lending, including poor underwriting, fraud, and valuation of securitization and residual interests."* In December 1999, federal banking agencies issued the **Interagency Guidance on Asset Securitization**

Activities, describing proper residual-interest valuation and explicitly flagging situations where such interests "should be assigned no value." The fix, however, only reached federally regulated banks holding residuals on their own books — it never touched the non-bank originators and Wall Street securitization desks that would run the far larger 2000s boom.

2007-08: the same disease, one structural layer more opaque

The 1999 guidance closed the specific loophole it targeted. The underlying trick reopened in a more complex instrument that the guidance never reached: **CDO tranching and Gaussian-copula correlation modeling.**

The model, explained precisely rather than named in passing. A CDO pools hundreds or thousands of individual mortgages, then slices the pool into tranches — senior (paid first, safest), mezzanine, and equity (paid last, absorbing the first losses). Pricing each slice requires knowing not just each loan's individual default probability, but how *correlated* those defaults are with each other — and correlation across thousands of individual mortgages had no clean historical dataset the way default rates did. David X. Li, then at CIBC, published the fix in "**On Default Correlation: A Copula Function Approach**" (Journal of

Fixed Income, March 2000): use credit default swap spreads as a market-price proxy for each loan's survival time, then map them onto a **Gaussian copula** — a mathematical device using the standard bell curve to produce a single joint-default distribution for an enormous, novel pool of assets, in seconds rather than through expensive simulation. Every major CDO desk and rating agency adopted it almost immediately, because it converted an intractable problem into a spreadsheet input.

Two flaws proved fatal, and both fit the same pattern as every other instrument in this document. First, the correlation number was estimated from a short, unrepresentative history — mid-2000s CDS data from a housing boom in which national home prices had never fallen simultaneously in living memory, so the model's key input was borrowed from a period that hadn't yet experienced the event it would later be used to price. Second, and more fundamental: the Gaussian copula has a specific, well-documented mathematical property called **zero tail dependence** — even at a "high" correlation setting, it treats truly extreme joint events (many defaults happening together, all at once) as vanishingly rare by construction. Real mortgage defaults during a national housing crash cluster hard in exactly that tail, and the model had no mechanism to represent the clustering, regardless of what correlation number was

fed into it. **The tranches that broke first in 2007-08 weren't the ones the model flagged as risky. They were the ones it had declared safest** — AAA ratings resting on an assumption (low default correlation across geographically diverse borrowers) that proved catastrophically wrong the moment the housing decline turned out to be national rather than regional.

The "residual" didn't disappear in this transition — it became the CDO equity tranche, several layers removed from a simple interest-only strip, dressed in a peer-reviewed model rather than a simple accrual assumption. The instrument changed, and the sophistication of the disguise increased. The move — recognize value today based on an assumption about tomorrow, and let reality true it up later — did not.

Today: the same move, in two places at once

Private credit's PIK mechanism books deferred, uncollected interest as current taxable income, which a BDC is then required to distribute against under its RIC tax status — real cash leaving the fund to cover income that was never actually received. **Burry's AI depreciation thesis** extends the assumed useful life of GPUs and servers from Nvidia's actual 2-3 year obsolescence cycle to 5-6 years, deferring real economic cost into 2026-2028 and inflating current reported earnings by an estimated \$176 billion in the process. Different instruments, same move: recognize

today, true up later, on a schedule nobody is required to disclose in aggregate until the schedule arrives.

The pattern across all three generations is now precise enough to state as a rule rather than an analogy: each version of this disease gets fixed only narrowly enough that the next decade's instrument evades the fix, and the interval between generations has run close to ten years each time — 1997-98, 2007-08, and a version now underway across 2025-2028. Whether this generation's regulatory response, if one comes, will be scoped broadly enough to reach whatever instrument carries the disease into the 2030s is, on this record, the wrong question to bet on.

Leg three: the risk asset being sold in an unwind is now overwhelmingly AI/tech

This is where the "sell risk" half of Leg One meets a market structure that didn't exist at anything like this scale in prior carry-trade unwinds.

The "Magnificent Seven" — Apple, Microsoft, Alphabet, Amazon, Nvidia, Meta, Tesla — reached approximately 40% of S&P 500 market capitalization in March 2026, per multiple concurring trackers, with more typical readings through the year in the 32-35%

range. For comparison: the ten largest stocks peaked at **27% during the dot-com bubble**. Today's top ten run **38-39%**. This is, without qualification, the highest concentration in the index's history, and the group is overwhelmingly AI-infrastructure-exposed (Nvidia directly; Microsoft, Amazon, Alphabet, and Meta via massive AI capex programs).

The mechanical consequence for a risk-off, deleveraging event: when carry-funded positions unwind and "sell risk assets" becomes the dominant flow, there is proportionally more forced selling pressure landing on a smaller number of names than at any point in market history. A shock that in 1999 would have hit a 27%-weighted cohort now hits a 35-40%-weighted one. The index-level damage from any given percentage decline in "risk appetite" is mechanically larger today, independent of whether anything is actually wrong with the AI trade's fundamentals.

Leg four: Burry's thesis is the fundamentals crack, and it rhymes with subprime more precisely than a loose comparison

Verified directly rather than summarized secondhand. Burry's argument has two separable parts, both mechanical rather than sentiment-based:

The depreciation-accounting claim. Hyperscalers are depreciating GPUs and servers over 5-6 years while Nvidia's own product cycle obsoletes the hardware in 2-3. Burry's estimate: this understates industry-wide costs by roughly **\$176 billion between 2026 and 2028**, with Oracle's earnings overstated by an estimated 26.9% and Meta's by 20.8% by 2028. This is not a demand call — it's a claim that **current reported earnings are mechanically inflated by an accounting convention**, and that the gap closes automatically as hardware actually needs replacing on its real schedule.

The circular-financing claim. Nvidia-style vendor financing — where a supplier funds its own customer's purchases — distorts the demand signal running back through the chain. UBS estimated the Nvidia-OpenAI arrangement alone at roughly 13% of Nvidia's projected 2026 revenue. Burry's own framing, confirmed directly: not that Nvidia is "Enron," but that it is "Cisco" — a real, useful technology sold into a demand cycle whose scale is partly a financing artifact rather than organic end-user pull.

The circularity escalated substantially in late July 2026, in a way worth stating with precision rather

than headline shorthand. The Wall Street Journal, independently confirmed by CNBC, reported Nvidia in talks for a **\$250 billion guarantee** on the lease and construction debt for OpenAI's planned 10-gigawatt Ohio data center campus — needed because OpenAI, despite a near-\$1 trillion private valuation, carries no investment-grade credit rating, so Nvidia's own balance sheet would stand in for one, letting lenders price the debt against the chipmaker rather than the borrower. Nvidia is reportedly discussing a **separate \$350 billion facility** to finance OpenAI's chip purchases for the same site — together, \$600 billion in discussed exposure, all of it ultimately securing demand for Nvidia's own products, from the company that makes them. Nvidia shares fell nearly 5% on the report.

One qualifier matters enough to state plainly: this is a reported negotiation, not a completed transaction, and the gap between the two is itself informative. Nvidia's own most recent quarterly filing caps its total *disclosed* guarantee exposure across all partner facility leases at **\$3.5 billion** — roughly two orders of magnitude below the \$250 billion figure under discussion. That gap doesn't mean the reporting is wrong; the negotiations are real, WSJ-sourced, and market-moving. It means that if the deal closes anything like as reported, the resulting balance-sheet exposure would appear suddenly, from a base that

current filings show as nearly immaterial — precisely the kind of disclosure gap Moody's has already warned about elsewhere in this document regarding hyperscaler lease guarantees generally.

A specific instrument now exists that turns Burry's depreciation thesis directly into a tradeable structure, worth naming rather than leaving implicit.

"Collateralized Chip Obligations" (CCOs) — GPU-only collateral, no underlying real estate or power infrastructure attached — carry deliberately compressed amortization schedules, typically 24-36 months, explicitly faster than the chip obsolescence cycle Burry's thesis describes; dynamic loan-to-value triggers tied to the secondary market price of the underlying chips, forcing the borrower to post additional collateral the moment a new Nvidia generation crashes the resale value of the old one; and cash-flow sweeps that redirect all revenue from active servers to debt paydown the moment utilization drops for 60-90 days. Lambda Labs' Morgan Stanley-led facility is a named example. **This is the sharpest possible version of the accounting-disease lineage traced earlier in this document — a structure explicitly engineered around the assumption that the collateral will lose most of its value on a known, short timeline, priced and covenanted accordingly, rather than an accounting convention quietly assuming it away.**

The subprime parallel, stated precisely rather than loosely: the 2008 crisis was not fundamentally a story of assets nobody understood — it was a story of **assets whose true economic quality was obscured by a permissive convention** (the Gaussian-copula correlation model detailed above, producing AAA ratings on tranches that didn't deserve them) until a trigger forced recognition. Burry's AI thesis has the identical shape: **a permissive accounting convention (extended useful life) is producing reported earnings that don't reflect true economic depreciation**, and the gap is scheduled to close on its own, mechanically, by 2028 regardless of what markets do in the meantime. Both are stories where the crisis wasn't "the assets were bad" — it was "the valuation convention was wrong, and eventually arithmetic wins."

Where the four legs are not four separate towers — they are one structure

This is the finding this document actually adds to the project, verified rather than inferred.

Blue Owl Capital is simultaneously one of the BDCs already showing rising stress in this project's own macro analysis (the MFS escalation watch) and the

single largest private financier of AI infrastructure in the world. \$27-30 billion for Meta's Hyperion data center in Louisiana (rated A+ by S&P, structured with PIMCO as the debt), and roughly \$13 billion alongside JPMorgan into the SPV that owns Oracle's OpenAI facility in Abilene, Texas. **Blackstone and KKR — the other two names already flagged in this project's macro analysis — both competed for the Meta mandate and lost only on price.**

The Federal Reserve has quantified the broader exposure chain. Up to 25% of bank loans to non-bank financial institutions now go to private credit firms, up from 1% in 2013. Life insurance companies hold nearly **\$1 trillion** in private credit. The Financial Stability Board, in a May 2026 report, explicitly warned that private credit is "concentrated in technology" and "untested through a severe downturn."

And the credit rating on the trophy AI infrastructure debt rests explicitly on an assumption Burry's thesis directly threatens. Coverage of the Meta-Blue Owl deal states plainly that the high ratings depend on "the underlying assumption of sustained AI demand and full occupancy" — precisely the assumption the circular-financing and demand-bullwhip argument calls into question. **This is not two risks that might correlate in a crisis. It is one collateral assumption (durable AI demand) underwriting both the equity**

multiple on Nvidia and the credit rating on the bonds financing the buildings that house its chips.

One structural note on precision, since it changes which private credit segment actually transmits which risk: the largest trophy AI infrastructure bonds (Meta-Blue Owl, Oracle's Abilene SPV) appear to be **fixed-rate, investment-grade paper specifically structured to attract insurance-company buyers seeking duration-matched income** — these carry refinancing and mark-to-market rate risk, not the immediate floating-rate cash-flow risk described in Leg Two. The **immediate, already-firing floating-rate mechanism** runs through the broader middle-market direct-lending book these same firms (Blue Owl, KKR, Blackstone, Ares, Golub) also manage — the book already showing the PIK increase and default stress documented above. Both segments sit on the same platforms and are managed by the same firms; a reputational or liquidity shock at the parent level (per the Bear Stearns precedent already noted in this project's macro analysis) does not respect the boundary between them.

How levered are the sponsors themselves, compared with Bear

Stearns — and does an SPV actually isolate the lender

Bear Stearns, precisely: 33-36:1. By late 2006, \$11.1 billion in tangible equity supported roughly \$350-395 billion in assets — a leverage ratio in that range depending on the exact snapshot. This wasn't an outlier: Morgan Stanley ran 33:1, Merrill Lynch 32:1, Lehman 31:1, Goldman 26:1, all in 2007. It was also regulator-enabled rather than accidental: a 1975 SEC rule had capped broker-dealer leverage at 12:1, but a 2004 rule change let the five largest investment banks opt into self-regulated net capital requirements — precisely what allowed them to reach 30-40:1 within three years.

None of the four firms central to this document run anything close to that, and the comparison is different in kind, not merely degree. BlackRock's own corporate debt-to-equity ratio sits at roughly 0.23-0.34 — essentially unleveraged. Blue Owl, Blackstone, and KKR are structurally fee-and-carry businesses: they manage other parties' capital for a percentage, rather than borrowing against their own balance sheet the way Bear Stearns' trading desks did. **The BDCs these firms sponsor operate under a genuine statutory ceiling, not a voluntary one:** since the 2018 Small Business Credit Availability Act, BDCs are capped at 150% asset coverage, which caps debt-to-equity at a

maximum of 2:1. Blue Owl's own Q1 2026 filing, describing "lower leverage" on \$8.5bn of BDC debt, is consistent with operating well inside that ceiling. **Bear Stearns, at its peak, ran at roughly fifteen to eighteen times the maximum leverage a BDC is legally permitted to carry today.** The post-2008 regulatory architecture genuinely did fix this specific vulnerability.

But this is the wrong place to be looking, and the JPMorgan-Oracle SPV structure shows exactly why.

JPMorgan is one of the lenders in the Blue Owl/JPMorgan special purpose vehicle that owns Oracle's Abilene, Texas OpenAI facility — roughly \$13 billion total, \$10 billion of it debt — and per Financial Times reporting, "lenders typically have recourse to the facility, land, and hardware rather than Oracle's broader balance sheet." That sentence answers two different questions, and conflating them produces a false sense of security.

Does the SPV isolate Oracle from the debt? Yes — that is precisely the structure's purpose, keeping roughly \$120 billion of AI infrastructure debt off hyperscaler balance sheets and credit ratings industry-wide over the past eighteen months. **Does the SPV isolate JPMorgan's own balance sheet from being affected? No — and it was never designed to.**

JPMorgan is not the sponsor hiding exposure from itself; it is the lender who voluntarily took on \$10

billion of secured credit risk. That sits on JPMorgan's own books as a real asset, and if the collateral — a purpose-built AI data center with no obvious alternative tenant — turns out to be worth less than assumed, JPMorgan absorbs the loss directly, exactly as with any secured loan gone bad. The SPV protects Oracle's credit rating; it cannot and does not protect the lender's own decision to extend the credit.

Where this does genuinely rhyme with 2008 is one layer beneath the headline structure, and it remains an open question rather than a resolved one. The SIV vehicles that damaged Citigroup in 2007 were legally off-balance-sheet too — what mattered was that Citi had extended implicit and explicit liquidity backstops that pulled the exposure back onto its own balance sheet the moment the SIVs could no longer roll their funding. There is no public evidence of a comparable backstop specific to the Abilene SPV. But Moody's has explicitly warned that "disclosures may not show the full picture" of hyperscaler lease and guarantee commitments generally — Meta's own \$28 billion residual value guarantee on its Hyperion data center financing appears only in footnotes to its annual report, not on the balance sheet. **The honest position: this specific deal's isolation of JPMorgan appears real, on the terms disclosed. Whether that holds across the sector cannot currently be confirmed from outside, which is precisely the condition that made**

2008's off-balance-sheet exposure far larger than any pre-crisis model had captured.

Where the leverage actually landed, once the sponsors' own balance sheets are ruled out

A natural question follows from everything above: if Blue Owl, Blackstone, KKR, and BlackRock aren't themselves highly leveraged at the corporate level — and they aren't, running debt-to-equity ratios nowhere near Bear Stearns' 33-36:1 — where does the actual risk sit? The answer, traced through, is that it has been distributed outward to three destinations, and the first is more concentrated and more self-referential than most outside observers appreciate.

Insurance companies: the largest concentration, and it is substantially circular

The scale is not modest. Apollo's captive insurer, Athene, carries roughly **\$30 billion of exposure to below-investment-grade loans**, including **\$25 billion in CLOs** — of which only 10% are rated AAA and 62% are rated below AA. KKR's Global Atlantic grew from \$72 billion to **\$158 billion in assets under management** during KKR's ownership (2020 to full

acquisition in January 2024) — more than doubling specifically because KKR could now direct that captive balance sheet into its own credit strategies. Blackstone paid \$2.2 billion in 2021 for a 9.9% stake in AIG's life and annuity unit, gaining effective control over its investment decisions, and separately acquired an Allstate life insurance unit the same year; Blackstone's insurance AUM reached **\$150 billion by the end of 2021 — a third of the firm's total assets under management** at the time. Across the sector, Moody's found roughly **ten large insurers concentrate \$352 billion** of private and illiquid investments between them — this risk is not diversified across the industry; it is clustered in a small number of balance sheets, nearly all affiliated with the largest alternative managers.

And a meaningful share of it is not arm's-length lending — it is the sponsor investing its own captive insurer's premiums into its own funds. AM Best — a mainstream insurance rating agency, not an activist critic — found that **roughly a fifth of the investments held by Athene's US Life Group and by KKR's Global Atlantic are loans to funds affiliated with their own private equity parent.** An Apollo executive has publicly characterized the firm's exposure to this kind of instrument as "very small"; an independent analysis of Athene's own 9,000-page statutory filing found \$25bn in CLOs alone and disputed that

characterization directly. The structural conflict is straightforward to state: a private equity firm that owns an insurance company can direct that insurer's permanent, continuously-replenished premium capital into the PE firm's own credit vehicles, with the same firm managing both the asset (the fund) and the counterparty (the insurer) sides of the transaction.

A fifth generation of the same accounting disease traced through this entire document is hiding inside this specific channel. These affiliated and illiquid holdings are increasingly booked as **Level 3** assets — the accounting classification for instruments with no observable market price, valued by internal model rather than actual transaction data. Level 3 assets reached 18% of the entire US insurance industry's \$3.8 trillion in fixed income holdings, and **a full third of Athene's and Global Atlantic's total assets** by the third quarter of 2025. This is mark-to-model valuation, performed substantially by the same firm that originated the asset — structurally identical to 1997-98's residual interest valuations and 2007-08's CDO correlation models, relocated once again into a new instrument and a new class of balance sheet.

Why insurance specifically, and why the pace has accelerated: fixed and fixed-indexed annuities generated roughly **\$280 billion in new US deposits in 2024 alone**. Unlike a private equity fund, which must periodically re-raise capital and remains exposed to

shifting investor sentiment, an insurer collects premiums continuously — one industry analysis described PE fundraising as "episodic" against insurance capital as "permanent." That structural difference is the entire reason Apollo, KKR, Blackstone, and Brookfield have each spent the last five years acquiring or taking investment control of insurers: it converts an unstable, redeemable capital base into a captive, continuously-replenishing one, largely outside the redemption pressure that hit the BDCs directly in March 2026.

Banks: the second channel, smaller but growing fastest

Federal Reserve research already cited in this document quantifies this precisely: **up to 25% of bank loans to non-bank financial institutions now go to private credit firms**, up from 1% in 2013, and bank credit lines extended to private credit vehicles grew from **\$8 billion in 2013 to \$95 billion by late 2024** — the fastest-growing lending category the researchers tracked. Banks are not generally holding the underlying risky middle-market loans directly; they are financing the funds that do, through subscription lines and NAV-based facilities. This is precisely the kind of intermediated exposure that made 2008's interbank contagion difficult to trace in real time — the risk is

real and growing, but it sits one layer removed from where a simple balance-sheet review would find it.

The distribution pattern, stated plainly

Put the two channels together, and the shape is the same one that defined the run-up to 2008: **the entity making the credit decision does not bear its full consequence.** The sponsors — Blue Owl, Blackstone, KKR, Apollo — collect management fees and carried interest regardless of how the underlying loans perform, while the tail risk has been distributed to parties several steps removed from the original underwriting: annuity holders and retirees, via captive insurers directing policyholder premiums into their own parent's funds; banks, via credit lines extended to the funds themselves; and, through the retail-BDC and "evergreen fund" products increasingly marketed over the past two years, individual wealth-channel investors buying instruments originally built for institutional balance sheets. **This is "originate to distribute" from the 2000s mortgage market, reappearing with insurance products marketed as safe, guaranteed retirement income as the primary distribution channel this time** — arguably a more consequential channel than institutional mortgage-backed securities ever was, because the buyers on the other end are ordinary retirees who were sold certainty.

The compounding mechanism, stated plainly

Put the four legs in sequence rather than in parallel, because the danger is in the sequence:

A yen-driven global rate shock raises the cost of capital everywhere at once, hitting an already-strained floating-rate private credit book mechanically — no new bad news required, just higher SOFR. **The same shock is a "sell risk assets" event**, landing disproportionately on the most historically concentrated equity cohort ever assembled, because that cohort *is* the risk asset the market has been running the carry trade through. **If Burry's timeline is anywhere close to right, the earnings and collateral assumptions underneath both the equity multiples and the private credit ratings are independently scheduled to crack in the same 2026-2028 window** — meaning the fundamentals leg doesn't need the liquidity leg to happen; it's already on its own clock.

The genuinely dangerous case is not either leg alone. It's the two arriving close enough together that forced selling (liquidity) lands on collateral whose underlying value assumption is simultaneously being revised downward (fundamentals) — with no safe-haven release valve, because the shock originates in

sovereign bonds, which are normally where capital flees to, not from. That last point is the precise echo of 2008: the crisis wasn't dangerous because risky assets fell. It was dangerous because the assets presumed safe (AAA tranches) turned out to be the epicenter. A yen/JGB-originated shock puts sovereign debt back in that role.

And the compounding, if it happens, would not land evenly across the private credit and AI-financing complex — it would bifurcate sharply, which matters for anyone trying to assess where the real damage concentrates. Lenders holding the physical infrastructure — land, power grid interconnections, the building itself — are comparatively protected even in a severe AI-demand disappointment, because that infrastructure can be stripped of AI-specific hardware and repurposed for ordinary cloud computing or enterprise storage, and because 10-15 year hyperscaler leases are typically structured as take-or-pay obligations the tenant owes regardless of actual utilization; plausible recovery in a bad scenario runs 80-100%. **Lenders holding only the chips — the Collateralized Chip Obligations described above — have no comparable fallback.** A specialized AI accelerator has essentially no alternative use once a newer generation exists, and a demand disappointment would flood the secondary market with hardware simultaneously, collapsing resale

values toward zero; plausible recovery in the same scenario is 10-30 cents on the dollar. **This means "AI credit risk" is not one number — it is a sharp bimodal split between well-protected hard-asset lending and effectively unsecured chip-only lending dressed in secured-loan documentation,** and which side of that split a given fund sits on matters more than its aggregate AI exposure headline.

Held honestly

This is a coherent, heavily-sourced tail scenario, not a prediction, and every figure above is independently checkable. The four legs do not require the same probability to compound — Leg Two (floating-rate mechanical transmission) is close to certain if Leg One occurs at all; Leg Three (concentration amplifying any risk-off event) is a structural fact already true today, not a future contingency; Leg Four (Burry's thesis) runs on its own accounting clock through 2028 regardless of what Japan does. What makes this a "perfect storm" scenario rather than four separate risk memos is that **the same three or four firms — Blue Owl foremost, with Blackstone and KKR close behind — sit at the literal intersection of all four legs,** which is a fact about market structure in 2026 that didn't exist in anything like this form even three years ago.

What would most cleanly falsify or confirm the compounding thesis, watched forward: whether the next BDC earnings cycle shows PIK and nonaccrual increases specifically concentrated in tech-sector loans (confirming the transmission channel is live rather than diffuse); whether the September 17-18 BOJ meeting (see the Japan report) produces a hike sharp enough to move global floating-rate benchmarks meaningfully; and whether Burry's promised further disclosures, or a hyperscaler's own capex guidance, move first on the depreciation question before any liquidity event forces the issue.

This document synthesizes previously separate threads in this project: the Japan/yen carry-trade material, Troy Black's prophetic words and this project's macro analysis on private credit and their Bear Stearns precedent, and new research on Michael Burry's AI accounting-valuation thesis and its structural links to private credit via Blue Owl, KKR, and Blackstone. Also draws on four additional historical precedents checked directly against primary sources: the Washington Mutual Option ARM analogy (WaMu 10-K filings, 2005-2007), the 1998 LTCM/LTCB dual crisis (Federal Reserve History, Bank of England, Japan Times, GAO), the specific asset-class breakdown of the 2007-08 carry-trade unwind, and the 1997-98 subprime residual-interest failures (SEC filings for Southern Pacific

Funding Corp. and Cityscape Financial Corp., American Bankruptcy Institute, Federal Reserve testimony, HUD research, the December 1999 Interagency Guidance on Asset Securitization Activities). The leverage comparison and insurance-channel analysis draw on Bear Stearns' own historical leverage ratios (Financial Crisis Inquiry Commission, Capital IQ), the JPMorgan-Oracle Abilene SPV structure (Quinn Emanuel client alert, Financial Times reporting), and the private equity-insurance nexus (Moody's, AM Best, CEPR, ABF Journal). The Gaussian copula explanation draws on David X. Li's original 2000 paper (Journal of Fixed Income) and subsequent technical and popular analysis of its role in the 2007-08 crisis. Bessent's 1992 and 2013 currency-trade history and the trilemma/euro-funding framework draw on independent commentary (Andrei Jikh) and independent verification (Newsweek, Fortune, Wikipedia, multiple biographical sources). The good-PIK/bad-PIK distinction, the Nvidia-OpenAI credit-wrap reporting, and the Collateralized Chip Obligation instrument draw on Financial Stability Board (May 2026), Fitch Ratings (April 2026), CNBC, and Wall Street Journal reporting independently confirmed by CNBC and cross-checked against Nvidia's own most recent quarterly filing (via Tom's Hardware) for disclosed guarantee exposure. All figures independently sourced (Federal Reserve Bank of Boston, KKR investor materials, Alternative Credit Council, Financial Stability Board, S&P Global Ratings coverage, TheStreet, CNBC,

Forbes, SEC filings, multiple market-concentration trackers). Not financial or investment advice – this is scenario analysis, not a forecast with an assigned probability, and should be read the way the rest of this project treats compound, multi-leg claims: real and worth tracking, not a single number.