Edoardo Gaffeo*

When Growth Destroys Collateral: Technological Obsolescence and Macroprudential Risk in the AI Era

https://doi.org/10.1515/ev-2026-0036
Received July 19, 2026; accepted September 11, 2026; published online October 14, 2026

Abstract: Schumpeterian growth creates by destroying. A critical financial consequence arises when capital rendered obsolete by innovation also serves as collateral for debt that markets treat as safe. Using the AI infrastructure boom as a case study, this paper documents the scale of investment, the spread of asset-backed and structured finance, and the emergence of investment-grade claims backed by data centres, GPUs, leases, and customer contracts that are increasingly reaching the portfolios of institutional investors. Systemic risk arises not from technological depreciation itself, but from a common collateral revaluation that transforms widely held, information-insensitive claims into information-sensitive ones. Such a shift can trigger screening, abrupt repricing, deleveraging, and fire sales across institutional portfolios. The mechanism echoes a central lesson of the Global Financial Crisis and carries implications for financial supervision, risk assessment, and complementary industrial policy.

Keywords: artificial intelligence; technological collateral; safe assets; private credit; financial stability

JEL Classification: G01; G23; O33

1 Introduction

In the Schumpeterian model, growth is inherently disruptive. Innovation drives productivity gains by devaluing existing capital structures to make way for new advancements (Aghion and Howitt 1992, 2009). The destruction is usually discussed

I am grateful to Alessio Emanuele Biondo, Artur M. Boscaglia, Saul Desiderio, Andrea Fracasso, Sandro Montresor, Luca Riccetti and two anonymous referees for insightful comments and suggestions on earlier drafts of this manuscript. The usual disclaimer applies.

***Corresponding author: Edoardo Gaffeo**, Department of Economics and Management, University of Trento, Trento, Italy, E-mail: edoardo.gaffeo@unitn.it

Open Access. © 2026 the author(s), published by De Gruyter. This work is licensed under the Creative Commons Attribution 4.0 International License.

on the real side: technologies disappear, rents shrink, firms lose market share, and resources move toward new vintages. Yet productive assets also sit behind credit contracts. When capital serves as collateral, creative destruction can destroy collateral capital as well. The same process that pushes the production frontier outward can weaken the asset base supporting claims written under the previous technological regime.

Technological obsolescence is no novelty, nor is the critical role of finance in driving economic growth. Banks and investors have long financed semiconductor fabrication, telecommunications networks, robotics, and other specialized capital with short economic lives. Crucially, the threat to stability is governed by the nature of the claims issued against that capital. A risky loan held by a specialist lender can absorb a technological loss without threatening the wider system. The stakes change when technology-backed claims acquire investment-grade, safe-like characteristics and migrate into portfolios whose holders have little reason to scrutinize the underlying collateral continuously.

Debt derives its social utility fundamentally from its capacity to remain information-insensitive (Holmstrom 2015; Gorton and Ordoñez 2023). By relying on seniority, guarantees, overcollateralization, long-term leases and high ratings, issuers can transform heterogeneous project risks into claims that circulate beyond specialist lenders. But information insensitivity is not immutable. An adverse technological shock can make previously unimportant information about liquidation values worth acquiring. Investors then start asking questions they had little reason to ask before, and repricing can become abrupt.

Seen from a debt structure's desk, the obvious question is why lenders cannot simply charge a premium for obsolescence risk. At origination, sophisticated lenders can price expected depreciation through spreads, haircuts, covenants, and loan-to-value ratios. As noted above, the systemic issue begins when downstream investors regard structured and rated claims as sufficiently safe. A common shock can then induce many holders to switch simultaneously from reliance on ratings and contractual protections to active scrutiny of collateral, counterparties, and recoveries. Refinancing can tighten and liquidity can disappear even if borrower-level expected losses were initially priced rationally.

The analogy with the Global Financial Crisis arises naturally from this mechanism. The lesson of 2007–08 was not simply that underlying assets could lose value, but that highly rated claims on apparently stable collateral had permeated institutional portfolios. Once doubts surfaced, the scramble to distinguish safe from unsafe claims became destabilizing (Coval et al. 2009; Gorton and Metrick 2012). This paper does not frame technology-backed debt as a new subprime market. Rather, it investigates whether an analogous informational shift can occur when rapidly

obsolescing assets support claims that carry investment-grade ratings and migrate into institutional portfolios.

AI infrastructure provides an unusually clear setting in which to examine that possibility. Frontier AI has generated enormous expenditure on data centers, advanced semiconductors, networking equipment, and power infrastructure, while credit markets have expanded alongside it. This paper documents the scale of that investment, the growth of corporate debt, private credit and securitization, and a more consequential development: selected AI and data-center financings are increasingly gaining investment-grade ratings and reaching broad institutional investor bases. That migration is what gives the AI investment cycle its broader financial significance. Once technological capital supports claims held well beyond specialist lenders, the consequences of obsolescence depend not only on the losses borne by the original financiers but also on how those claims are structured, distributed, and treated across the financial system.

Ultimately, the debt architecture funding the AI buildout exposes a deeper tension between Schumpeterian growth and financial stability. Creative destruction generates growth by rendering old assets obsolete. When those assets double as financial collateral, accelerated technological change threatens both recovery values and the information-insensitivity of the debt built upon them. Through this channel, balance-sheet contagion can travel beyond the directly impacted firms, transforming a routine growth process into a mechanism of systemic fragility.

This distinction also disciplines the policy response. If technological risk is correctly priced and losses remain with those who willingly bear them, there is little macroprudential case for intervention. The rationale emerges when common exposures, refinancing externalities, fire-sale dynamics, or the migration of technology-backed claims into safe-asset portfolios generate effects that bilateral contracts cannot internalize. The task for macroprudential policy is not to restrain creative destruction or ask supervisors to pick technological winners, but to ensure that a financial system increasingly willing to treat technology-backed claims as safe remains resilient when innovation suddenly makes the collateral beneath them worth investigating.

The remainder of the paper is organized as follows. Section 2 documents the scale and financing of the current AI infrastructure buildout. Section 3 develops the concept of technological collateral. Section 4 examines amplification through contracts, refinancing, and asset sales. Section 5 links the mechanism to creative destruction and information-insensitive debt. Section 6 develops macroprudential implications. Section 7 concludes.

2 The Scale and Financing of the AI Infrastructure Buildout

The AI investment cycle is remarkable first for the sheer scale of capital expenditure it has generated. Frontier models rely on advanced semiconductors, servers, networking equipment, data centers, cooling systems, and large electricity connections. Figure 1, based on Federal Reserve data assembled from SEC filings, shows purchases of property and equipment by Amazon, Google, Meta, Microsoft, Oracle, and CoreWeave rising from about US$71 billion in 2019 to US$422 billion in 2025. These totals include substantial non-AI investment, but the sixfold increase documents the acceleration clearly (Soto et al. 2026).

Broader measures tell the same story. U.S. private data-center construction rose from an annualized rate of roughly US$9 billion at the beginning of 2021 to almost US$60 billion by May 2026, while private non-residential investment in computers and peripheral equipment increased from about US$144 billion in 2021Q1 to US$384 billion in 2026Q1. Neither series is AI-specific; taken together with firm-level capital expenditure, however, they document a large physical investment cycle rather than

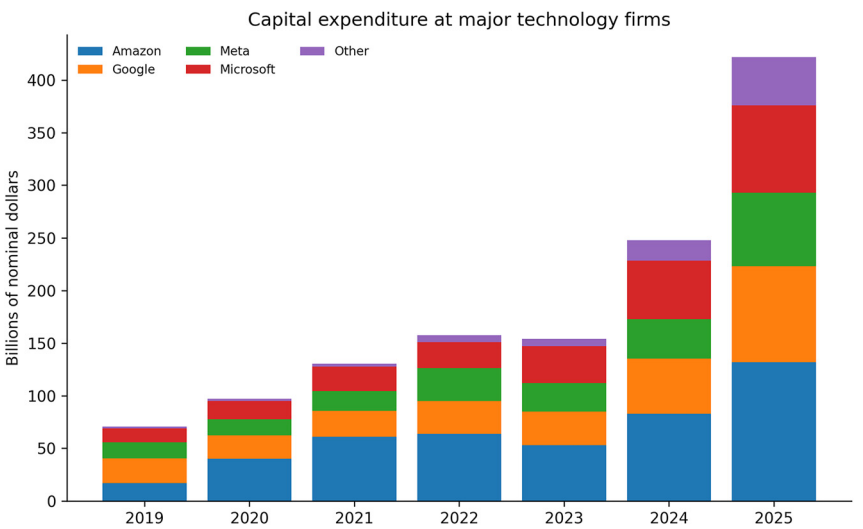


Figure 1: Capital expenditure at major technology firms, 2019–2025. Notes: Capital expenditure is measured as purchases of property and equipment reported in firms' SEC filings. "Other" includes Oracle and CoreWeave. The totals include non-AI investment and should therefore not be interpreted as AI-only capital expenditure. Source: Board of Governors of the Federal Reserve system, based on SEC filings.

one inferred from announcements alone (Van Nieuwerburgh 2026). Investment announcements by AI-native firms provide complementary, if coarser, evidence of the same expansion: Anthropic announced a US$50 billion U.S. computing-infrastructure program in November 2025, while by September 2025 OpenAI had announced more than US$400 billion of planned investment under the Stargate program (Figures 2 and 3).

Scale matters because external finance has become more important at the margin. The largest hyperscalers retain formidable internal resources, but the BIS reports that Alphabet, Amazon, Meta, Microsoft, and Oracle alone issued more than US$100 billion in gross corporate bonds in 2025. This figure covers only these five companies and therefore captures just one segment of the financing associated with the broader infrastructure buildout. Private-credit funds, banks, insurers, asset managers, and securitization markets are also financing data centers and computing infrastructure (Bank for International Settlements 2026).

These financing arrangements differ in economically important ways. Corporate bonds rely on the sponsor's broader balance sheet, whereas project and non-recourse financing relies more narrowly on the assets and contractual cash flows associated with the infrastructure being financed. Private-credit facilities are often secured by GPUs, servers, subsidiary equity, customer contracts, or combinations of

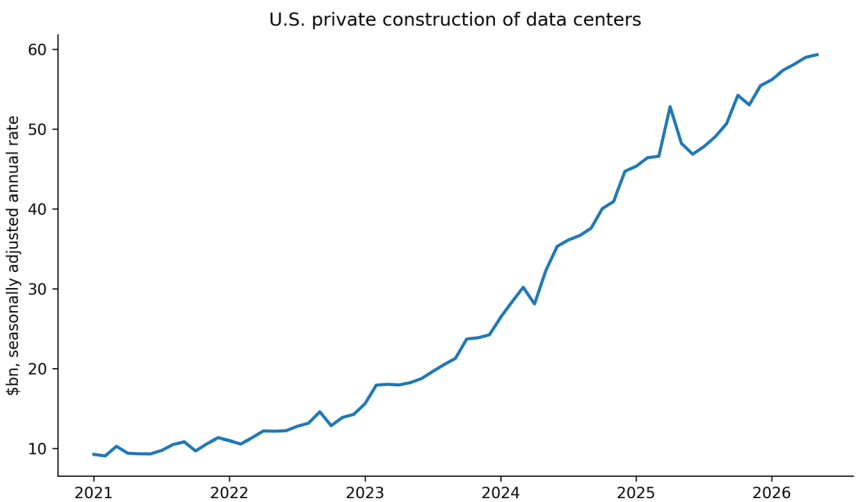


Figure 2: U.S. private construction of data centers. Notes: Monthly value put in place, seasonally adjusted annual rate, billions of nominal U.S. dollars. The series measures construction structures and does not capture the value of servers and other computing equipment installed in data centers. Source: U.S. Census Bureau, as compiled by the Board of Governors of the Federal Reserve system.

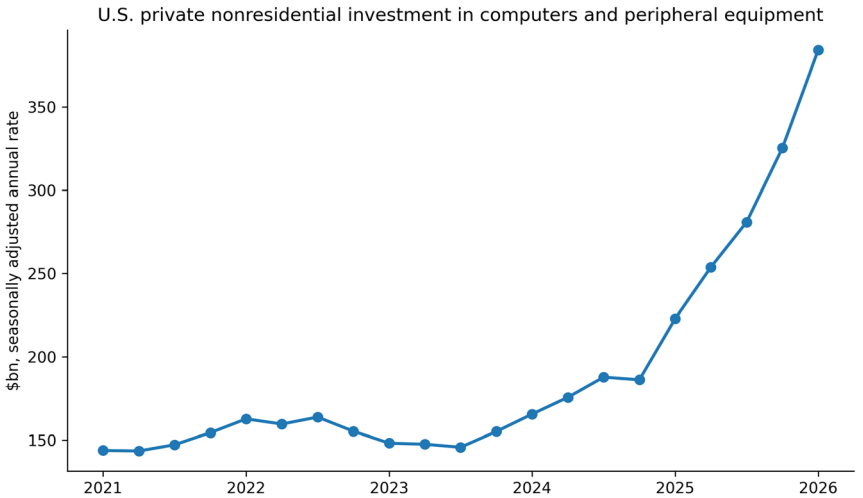


Figure 3: U.S. private nonresidential investment in computers and peripheral equipment. Notes: Quarterly seasonally adjusted annual rate, billions of nominal U.S. dollars. The category includes substantial non-AI activity. Source: U.S. Bureau of economic analysis, as compiled by the Board of Governors of the Federal Reserve system.

these assets and claims. Data-center securitizations go one step further, transforming properties, leases, and associated cash flows into tradable securities with different degrees of seniority. The IMF reports about US$46 billion of data-center securitizations since 2018, with issuance reaching a record level in 2025.[1]

Table 1 shows how quickly the market has broadened. Investment-grade data-center ABS appeared alongside large project-style and GPU-backed financings. Hyperion's October 2025 financing produced about US$27.3 billion of A + -rated debt supported by Meta leases and a residual-value guarantee; PIMCO disclosures subsequently showed the notes in fixed-income portfolios. On 31 March 2026, CoreWeave closed its US$8.5 billion DDTL 4.0, rated A3 by Moody's and A(low) by DBRS and backed by HPC infrastructure, a customer contract, and data-center leases. Blackstone Credit & Insurance anchored the transaction, joined by financial institutions, asset managers, and insurers.[2]

The migration of technology-linked debt from specialist lenders toward highly rated claims held by mainstream institutional investors is key, since it changes the

1 International Monetary Fund (2026), box on data-center financing and securitization.

2 S&P Global Ratings, "Beignet Investor LLC's $27.3 Billion Senior Secured Debt Assigned 'A+' Rating; Outlook Stable," 23 October 2025; CoreWeave, "CoreWeave Closes Landmark $8.5 Billion Financing Facility, Achieving First Investment-Grade Rated GPU-backed Financing," 31 March 2026.

Table 1: Selected AI and data-center debt: ratings, collateral and investor channel.

Transaction	Size/rating	Seniority, security and backing	Maturity	Disclosed investor channel
Compass III 2025-3 class A-2 (17 Jul. 2025)	US$415m; S&P A (sf)	ABS backed by four operating data centers, fixtures, leases, reserves, and equity interests in asset entities	Jul. 2030 anticipated/Jul. 2050 legal	Structured-finance institutional market
QTS ABS II 2025-1 class A-2 (30 Sep. 2025)	US$482m; S&P A-(sf)	ABS backed by data-center real estate, fixtures, tenant leases, reserves, transaction accounts, and equity interests in asset entities	Sep. 2030 anticipated/Sep. 2055 legal	Structured-finance institutional market; subsequent registered-fund holdings are publicly disclosed
Meta–Blue Owl/Beignet investor (Oct. 2025)	US$27.3bn; S&P A+	Senior project-style debt at HoldCo; Meta leases and a residual-value guarantee support cash flows; economic exposure to the Hyperion data-center JV	May 2049	Institutional bond market; PIMCO-managed funds took a large share of the offering
CoreWeave DDTL 4.0 (31 Mar. 2026)	US$8.5bn; Moody's A3/DBRS A(low)	Non-recourse; first-priority security over borrower assets, including cloud infrastructure equipment, customer contract, and data-center leases	Mar. 2032	Anchored by Blackstone credit & insurance; global financial institutions, asset managers, and insurers
CoreWeave DDTL 5.0 (15 May 2026)	US$3.1bn; Moody's Ba2/Fitch BB+	First-priority security over substantially all borrower assets; financing tied to GPU infrastructure deployed under two customer contracts	15 Nov. 2031	First publicly syndicated HPC infrastructure-backed facility; broader institutional lender base and secondary-market trading

Transactions are illustrative, not a census. Ratings refer to the cited issue or tranche and do not imply that the underlying technological assets themselves are safe. "Investor channel" records disclosed placement, market structure, or observed institutional holdings, rather than the complete beneficial-owner base. DDTL denotes delayed-draw term loan; ABS denotes asset-backed security; HPC denotes high-performance computing. The CoreWeave DDTL 5.0 facility illustrates the widening institutional distribution of HPC-backed credit through public syndication and secondary-market trading. Sources: S&P Global Ratings; CoreWeave and SEC filings; Meta and Blue Owl disclosures; PIMCO portfolio disclosures.

nature of the exposure. A specialist lender has both the incentive and the expertise to monitor technological obsolescence continuously and price it into the loan. As similar risks enter conventional investment-grade portfolios, however, their apparent safety can reduce the need for such scrutiny. The issue then extends beyond expected credit losses to the informational properties of the claims themselves: debt that circulates because investors have little reason to investigate its collateral may behave very differently once technological change gives them a reason to do so.

Taken together, these developments reveal a financial architecture increasingly built around large-scale technological investment, expanding credit finance, asset-backed structures, and a broader institutional distribution of technology-linked claims. What matters next is the nature of the collateral supporting those claims: how its economic value responds to technological change, how readily it can be redeployed, and what creditors can expect to recover when innovation renders it obsolete.

3 Technological Collateral and Financial Amplification

Technical change is often embodied in new capital. Advances in productivity arrive not simply as better ideas, but through new generations of equipment that make earlier vintages economically inferior. This is especially visible in technology-intensive sectors, where semiconductor equipment, telecommunications networks, robotics, and other specialized assets can remain physically durable while having remarkably short economic lives. A GPU may remain perfectly functional yet lose much of its value because a new architecture delivers more computing power per unit of energy, space, or cost. When such capital also secures debt, it becomes what I call *technological collateral*: productive capital whose recovery value depends materially on its position relative to a moving technological frontier. Creative destruction then operates through relative productivity. New capital need not make old capital useless; it need only make it sufficiently inferior that users are no longer willing to pay prices consistent with earlier valuations.

For creditors, the relevant object is recovery value rather than physical durability. Recovery depends on alternative uses, secondary-market depth, redeployment costs, and the technological state prevailing at liquidation. Redeployability provides a buffer: older GPUs may serve inference or less demanding workloads, data centers can sometimes be adapted, and contracts may support debt service even when spot prices fall. Obsolescence should not therefore be equated with immediate collapse.

The concern is that a common technological or demand shock can make many lenders revise cash flows and liquidation values at once.

That simultaneity separates ordinary credit risk from a macro-financial problem. At the loan level, expected obsolescence is priced through spreads, haircuts, covenants, and LTV ratios. As historical semiconductor and telecommunications cycles show, rapid obsolescence can be absorbed without systemic disruption – so long as exposure remains localized. The AI boom risks becoming systemically relevant if exposures grow large, leverage or refinancing needs are substantial, and a shared technological frontier causes correlated revaluations.

Under these conditions, diversification fails: a single shift in computing architecture or AI profitability simultaneously impairs underlying assets, long-term leases, and counterparty cash flows. When such assets back specialist loans, collateral devaluation remains a credit-loss problem. But when they back investment-grade claims held by asset managers, insurers, and other institutional investors, the nature of the risk changes. Obsolescence reduces expected recoveries, while uncertainty over those recoveries risks turning previously safe-like claims into information-sensitive ones. Below, we first model this loss-amplification channel before examining the resulting informational regime shift.

4 From Collateral Losses to Financial Amplification

A decline in technological collateral does not become systemic merely because it produces losses. Credit markets absorb losses continuously. The macro-financial issue arises when deteriorating collateral values change behavior elsewhere in the system and magnify the original shock. While technological risk originates in the economics of the asset, financial amplification depends on the contracts and balance sheets built around it.

Secured lending is designed to react when pledged assets lose value, as lower collateral values can reduce borrowing capacity, raise haircuts, tighten loan-to-value constraints, or make financial covenants binding. These responses are individually rational, but a common technological shock can make many of them occur at once. In turn, refinancing reinforces the effect: AI infrastructure requires large up-front investment and may have a long physical life, while the liabilities financing it can mature much sooner. A borrower may therefore return to the market with assets whose technological position has changed substantially since origination. Lower expected recoveries can lead lenders to reduce advance rates, shorten maturities, raise spreads, or demand additional equity. Crucially, the project need not be

insolvent for this mechanism to operate. A widening financing gap may be enough to force firms to cut investment, draw down liquidity, or sell assets. Financial stress can therefore spill into the real economy well before default occurs.

AI infrastructure finance adds complexity because collateral, contractual cash flows, and counterparties are often tightly intertwined. Some transactions are non-recourse, leaving creditors dependent primarily on the assets and cash flows of a dedicated vehicle rather than on the sponsor's wider balance sheet. Large cloud and computing contracts add another layer. When customers commit in advance to purchase computing capacity over several years, the supplier records the contracted but not yet recognized revenue as remaining performance obligations (RPOs). These commitments give lenders greater visibility over future cash flows and can therefore support the financing of the infrastructure needed to fulfil the contracts.[3] These relationships can overlap in what is sometimes described as an "AI money loop": hyperscalers invest in or contract with AI firms, those firms purchase computing capacity from the same or related hyperscalers, and the resulting long-term commitments support further infrastructure investment. This is not a literal circulation of the same funds, but a web of interdependent balance sheets and expected cash flows.

Such arrangements are useful since they can strengthen creditor protection: long-term contracts make revenues more predictable, non-recourse structures isolate assets and liabilities, and guarantees can shield senior creditors. Yet apparently separate protections may depend on the same underlying economics. Collateral values, customer creditworthiness, infrastructure utilization, and contractual cash flows can all deteriorate together if a technological breakthrough, a change in model architecture, or a revision in expected AI profitability affects several links in the structure at once. Diversification across firms then offers less protection when exposures share the same technological frontier.

Asset sales can further amplify the adjustment. Secondary markets for specialized equipment may work well when buyers and sellers arrive gradually, but a common shock can bring many similar assets to market simultaneously. If owners attempt to dispose of the same generation of GPUs, servers, or related equipment, prices may reflect the urgency of sellers rather than longer-run values in alternative uses. The event that makes creditors interested in liquidation can therefore make liquidation less attractive. This is the familiar logic of fire sales and collateral amplification (Kiyotaki and Moore 1997; Shleifer and Vishny 2011; Brunnermeier and Pedersen 2009), with a Schumpeterian trigger. Lower secondary-market prices weaken collateral values and expected recoveries, inducing further deleveraging

3 Remaining performance obligations (RPOs) represent contracted revenue that has not yet been recognized under existing customer contracts.

and potentially more sales. Technological progress has not physically destroyed the capital; it has changed the price at which balance-sheet-constrained owners can turn it into cash.

The sheer volume of financing broadens the channel through which balance-sheet adjustments transmit across the financial system. Through interconnected instruments, private credit, banks, insurers, asset managers, and securitization vehicles share exposure to a common investment cycle – whether via direct loans, senior tranches, or residual-value guarantees. A single technological shock can thus manifest simultaneously as a collateral problem for one lender, a spread shock for another, and a valuation loss for a third. Credit ratings compound this vulnerability: downgrades mechanically shift capital requirements, collateral eligibility, and investment mandates, compelling portfolio reallocations regardless of underlying asset liquidity.

These mechanisms explain how technological obsolescence can turn collateral losses into coordinated balance-sheet adjustment. They still assume, however, that investors understand the claims they hold and react through familiar contractual and portfolio constraints. A different problem arises when those claims have come to be treated as sufficiently safe that close scrutiny of their collateral appears unnecessary. A technological shock can then change not only recovery values, but the value of information about the claims themselves.

5 Creative Destruction and the Fragility of Safe-Like Debt

A more subtle vulnerability emerges when technology-backed claims acquire sufficiently safe-like characteristics to circulate beyond specialist lenders and enter institutional fixed-income portfolios. At that point, the relevant question is no longer only the expected loss embedded in the debt, but how much information investors need in order to hold it.

What makes safe debt economically valuable is precisely its information insensitivity: holders have little incentive to acquire private information about the assets supporting it (Holmstrom 2015; Gorton and Ordoñez 2023). When repayment is sufficiently likely across a wide range of states, detailed investigation of collateral quality has little private payoff. Such claims can circulate without continuous screening of the underlying assets. Their safety rests not on the absence of risk, but on the fact that information about that risk is normally not worth producing.

Financial structuring can make a claim safe even when the assets ultimately supporting it are risky. Seniority, junior capital, long-term leases, guarantees, and

other forms of credit enhancement create a buffer between fluctuations in the value of the underlying assets and losses borne by senior creditors. As long as that buffer is comfortably large, holders of the senior claim have little reason to investigate continuously what a particular GPU generation would fetch in liquidation. This is precisely what allows the debt to remain information-insensitive. The problem arises when technological obsolescence is large enough to erode that buffer. A new generation of equipment can sharply reduce the resale value of existing GPUs while leaving them physically intact. If the same technological shift also weakens customers' expected cash flows and floods secondary markets with similar equipment, several layers of protection can deteriorate together. At that point, information about the collateral, contracts, and counterparties becomes relevant to creditors who previously had little reason to acquire it.

This is where creative destruction acquires an informational dimension. A technological breakthrough changes relative capital values and, with them, expected recoveries on debt backed by older vintages. Once losses threaten the buffers that made senior claims safe-like, investors have reason to distinguish exposures they previously treated as equivalent: which assets can be redeployed, which contracts remain valuable, which counterparties are sound, and how much protection remains before senior creditors take losses. Screening takes time and expertise. Until strong claims can be separated from weak ones, uncertainty itself commands a price: spreads widen, liquidity deteriorates, and some investors may sell simply because they cannot establish quickly enough that a claim remains safe.

The Global Financial Crisis offers a clear structural precedent. The lesson of 2007–08 was not merely that mortgage collateral could lose value, but that structured, rated claims allowed large parts of the financial system to treat them as default-free. Once doubts materialized, investors were forced to scrutinize exposures that previously required no monitoring, triggering a systemic withdrawal of liquidity (Coval et al. 2009; Gorton and Metrick 2012). Technological collateral gives this mechanism a Schumpeterian twist: innovation itself supplies the shock, as advances expanding the production frontier simultaneously erode the recovery value of recently installed capital. Here, the critical threshold requires neither mass default nor worthless collateral; it is simply the point at which asset deterioration makes information acquisition worthwhile for a broad set of holders. Claims previously traded on coarse public signals – such as ratings and seniority – suddenly become differentiated by collateral quality and counterparty strength, stressing haircuts, liquidity, and refinancing terms. An investment-grade rating does not guarantee this transition occurs. Rather, broad institutional distribution of rated debt creates the necessary preconditions for an informational shock, while the extent to which holders treat these claims as truly information-insensitive remains an empirical question.

The underlying tension emerges clearly. Financial systems economize on information by engineering claims that require no continuous scrutiny; economic growth advances by replacing old capital with superior vintages. Each mechanism is efficient in isolation. Their interaction turns fragile when the capital rendered obsolete by innovation is the very collateral that enabled financial claims to circulate without scrutiny. The macroprudential question is whether the system can absorb the transition when creative destruction makes previously unexamined collateral worth investigating again.

6 Macroprudential Implications

Technological obsolescence, by itself, provides no case for a new regulatory regime. Rapid depreciation is not a market failure, and investors financing specialized capital should bear the consequences of misjudging its economic life. Asking supervisors to value technologies would simply replace market risk with a public information problem.

The macroprudential problem begins when losses cease to be confined to the investors that priced and accepted the original risk. A lender may correctly price its own exposure while having no reason to internalize how a common collateral shock affects other institutions through refinancing pressure, fire sales, or simultaneous portfolio adjustment. The problem becomes more consequential as technology-linked debt migrates from specialist lenders into highly rated claims held across institutional portfolios, where a deterioration in perceived safety can trigger responses well beyond the originating credit relationship.

This leaves an active role for macroprudential policy. Its purpose is not to shield investors from technological losses, but to prevent individually rational exposures from accumulating in forms that make creative destruction unnecessarily costly for the financial system and, ultimately, for investment and growth. In this sense, macroprudential policy is central to reconciling financial resilience with an innovation process that must be allowed to render existing capital obsolete. Supervisors therefore need to know where technology-linked claims ultimately reside, how concentrated the exposures are, and how their holders may respond when perceived safety deteriorates.

Exposure mapping is a natural starting point. Central banks and prudential supervisors should trace direct loans, securitizations, project vehicles, guarantees, residual-value commitments, bank lending to private-credit funds, and institutional holdings of rated technology-linked debt. The perimeter should follow the risk rather than the legal form. Ratings deserve particular attention because they influence

portfolio eligibility, capital treatment, collateral acceptance, and monitoring intensity.

Stress tests should treat technological change as a scenario rather than a forecast. Authorities can ask how balance sheets respond when recovery values fall sharply while cash flows weaken and secondary-market liquidity contracts. The exercise should extend beyond direct credit losses to higher haircuts, rating migration, tighter refinancing, and portfolio constraints. Information sensitivity cannot be observed directly, but its potential consequences can be tested where similar claims are widely held, collateral is concentrated, markets are thin, and investors face common rating or portfolio constraints.

Such exercises must cross institutional boundaries. A private-credit fund may have locked-up capital but rely on bank financing; an insurer may face higher capital requirements after a downgrade; a bond fund may confront outflows. Looking at each institution separately can miss a shock transmitted through different constraints at different points in the system. For the same reason, technology-specific capital charges would be misplaced. Prudential attention should follow concentration, leverage, maturity mismatch, opacity, common collateral, rating dependence, and broad institutional distribution – not an AI label.

Industrial policy has a narrower, complementary role. The ability to redeploy technological capital after obsolescence supports recovery values and limits the losses borne by creditors. Interoperability, modular infrastructure, functioning secondary markets, and the repurposing of data-center capacity can all increase that flexibility. Where governments already shape standards, energy connections, procurement, or infrastructure design, they can therefore favor adaptability without attempting to steer technological choice. The distinction is important. Policies designed to preserve the value of particular assets risk subsidizing obsolete capital, protecting incumbents, encouraging leverage, and socializing private losses. Public intervention should instead remain technologically neutral where possible and concentrate on the market infrastructure that makes capital easier to adapt and redeploy.

The resulting division of labor is clear. Markets select technologies and price their risks; financial institutions bear the losses when those choices prove wrong; industrial policy can improve the adaptability and redeployability of capital without protecting particular technologies. Macroprudential authorities have a different task: addressing the systemic effects that decentralized decisions cannot internalize when individually rational exposures become concentrated and widely distributed across the financial system. The objective is not to shield finance from creative destruction, but to prevent financial amplification from turning technological renewal into an unnecessarily severe contraction in credit, investment, and growth.

7 Conclusions

Creative destruction has a financial side. Innovation raises productivity by displacing older capital; when that capital serves as collateral, it can also erode the asset values supporting outstanding credit. Technological collateral is not new, and rapid obsolescence is not inherently systemic. What matters is the financial architecture built around it.

The AI infrastructure boom makes that architecture visible. Physical investment has surged, corporate debt, private credit, and securitization have expanded, and some claims backed directly or indirectly by data-centre and computing infrastructure now carry investment-grade ratings and reach asset managers, insurers, and other institutional investors. This does not mean that technology-linked debt is already a systemic safe asset, It does show that technological risk is migrating beyond specialist lenders into financial structures whose resilience matters well beyond the original credit relationship.

That migration creates two related vulnerabilities. Lower collateral values can weaken recoveries, complicate refinancing, activate contractual constraints, and induce asset sales. More consequentially, claims previously treated as safe can cease to be information-insensitive. Investors then investigate collateral and counterparties they had little reason to scrutinize, and liquidity can deteriorate because uncertainty about where losses reside suddenly becomes valuable information. The narrow lesson from the Global Financial Crisis is that this informational shift can amplify the underlying collateral shock.

Technological collateral adds a distinctive source of fragility because innovation itself can undermine the assets supporting financial claims. The same technological advance that raises productivity and creates new investment opportunities can rapidly erode the recovery value of earlier capital. When that capital supports debt treated as safe, creative destruction reaches directly into the financial system: it destroys collateral and, with it, the basis on which those claims could circulate without close scrutiny.

A proper division of labor is therefore needed. Markets should price technological risk, and investors should bear the losses when those judgments prove wrong. Macroprudential policy enters where common exposures can turn collateral revaluation into collective balance-sheet adjustment, with consequences extending beyond the original risk takers. The deeper challenge follows from the nature of innovative growth itself. An economy driven by creative destruction must allow capital to become obsolete while maintaining a financial architecture capable of absorbing the losses that obsolescence brings. Financial stability ultimately depends

on making finance resilient to creative destruction, not on making creative destruction less destructive.

References

Aghion, P., and P. Howitt. 1992. “A Model of Growth Through Creative Destruction.” *Econometrica* 60 (2): 323–51.

Aghion, P., and P. Howitt. 2009. *The Economics of Growth*. Cambridge: MIT Press.

Bank for International Settlements. 2026. *Progress and Peril. Chapter I: Macroeconomic Headwinds and AI Infrastructure Concentration, Annual Economic Report 2026*. Basel: BIS.

Brunnermeier, M. K., and L. H. Pedersen. 2009. “Market Liquidity and Funding Liquidity.” *Review of Financial Studies* 22 (6): 2201–38.

Coval, J., J. Jurek, and E. Stafford. 2009. “The Economics of Structured Finance.” *The Journal of Economic Perspectives* 23 (1): 3–25.

Gorton, G., and A. Metrick. 2012. “Securitized Banking and the Run on Repo.” *Journal of Financial Economics* 104 (3): 425–51.

Gorton, G., and G. Ordoñez. 2023. *Macroeconomics and Financial Crises*. Princeton: Princeton University Press.

Holmstrom, B. 2015. Understanding the Role of Debt in the Financial System, BIS Working Paper No 479.

International Monetary Fund. 2026. *Global Financial Stability Report. Chapter 2: The Rise and Risks of Private Credit*. Washington: IMF.

Kiyotaki, N., and J. Moore. 1997. “Credit Cycles.” *Journal of Political Economy* 105 (2): 211–48.

Shleifer, A., and R. Vishny. 2011. “Fire Sales in Finance and Macroeconomics.” *The Journal of Economic Perspectives* 25 (1): 29–48.

Soto, P., M. Thieu, and J. Allen. 2026. “The AI Buildout and the Economy: Publicly Available Data to Assess AI’s Impact - FEDS Notes”. Washington: Board of Governors of the Federal Reserve System.

Van Nieuwerburgh, S. 2026. “Financing the AI Buildout.” Brookings Papers of Economic Activity, forthcoming.