



The Resale Standard: Forward-Looking GPU Valuations

Transaction-based, stochastic forward GPU collateral valuations

The importance of Resale-based GPU valuations

Hard asset-based lending requires confidence in an asset's future valuation. Three criteria are required to establish a thriving lending market: revenue-generating, large market size, and verifiable resale values. AI-Infrastructure meets the first two; the third has been lagging until now.

Lenders and broader financial markets rely on valuation methodologies developed over decades that meet strict criteria and government-recognized standards when assessing the value of hard assets. IFRS 13, developed by the international body IASB, in partnership with the US FASB, created a 3-level standard. All carry dependence on observable market data, with identical asset transactions bearing the largest weight.

For GPUs, this data has historically been thin and hard to aggregate. This is why we've seen multiple attempts to use more abundant proxy data, such as rental rates, to establish forward-looking valuations. This fails the IFRS 13 level 1 requirement, and as outlined in the Barkr report from April 2026, the rental and resale markets do not move in lockstep. While revenue is an important part of underwriting, rental rates are opaque, volatile and a poor indicator of future asset value.

At Barkr, our decades of experience across multiple asset types, combined with our unique ability to gather large data sets, means we can deliver highly accurate, resale-based valuations that lenders can trust. Meeting IFRS standards.

Key findings

This research report builds on our prior reports (April, August) and develops forward-looking valuations for enterprise GPU systems. The report focuses on two models: an eight-GPU H100 system, a more mature platform, and the GB300 NVL72, a newer rack-scale architecture. Barkr has built a formidable dataset of observable prices for identical assets, broader historical price evidence, technical attributes and forward-looking market-participant assumptions. Consistent with IFRS 13 and USPAP, the basis of the valuation comes from the resale or realized price of identical physical assets, rather than the compute service generated by that asset. Forward values are generated through 100,000-path Monte Carlo simulations in which prices evolve across four persistent market regimes: correction and rebalancing, balanced market, demand squeeze and realized supply shock.

Correlated normally distributed demand and supply-inelasticity innovations alter monthly transition probabilities, while each regime exhibits different statistical moments. The lower valuation curve is the conditional median of the low-range regime, what we at Barkr consider 'marketable cash value' against which we have issued almost \$300 million in valuation guarantees; the upper curve is the conditional median of the realized supply-shock regime. Indicative fair market value is the median of the complete evolving regime mixture. The analysis yields three key findings. First, prior-generation systems can retain substantial economic utility because a single architecture runs every phase of the AI lifecycle, supporting workloads as they migrate from frontier training toward inference, fine-tuning and batch processing. Second, chip and system values can remain relatively decorrelated from compute-access prices because hardware is transferable across operators, whereas rental prices embed utilization, service terms, power and operating execution. Third, short- and medium-term supply is relatively inelastic across semiconductor manufacturing, advanced packaging, critical minerals, electrical infrastructure and data-center commissioning, while demand for tokens and AI capital expenditure continues to expand. These conditions create persistent scarcity premia for immediately deployable hardware and support the financialization of GPUs through equipment facilities, asset-backed structures and usage-linked financing.

1. Introduction and valuation problem

The rapid expansion of artificial-intelligence infrastructure has created a valuation problem. Enterprise GPU systems are physical assets, but their observed economic output has been largely expressed through rental rates for compute. Those two objects are related but meaningfully different. A GPU asset transaction includes its accelerators, processors, memory, interconnect, power architecture and associated configuration. A compute contract transfers temporary access to processing capacity and often embeds utilization, software, networking, storage, power, service quality and counterparty-specific terms. A defensible collateral framework must therefore isolate the value of the asset from the durable revenue characteristics of the service it produces.

"Barkr's work valuing illiquid, hard assets is vital to understanding the risk involved in lending against those assets. Munich Re supports Barkr in this effort through aiSure™ performance warranty, bringing accuracy and confidence to asset-backed lending",
Tom Van den Brulle, Head of New and Green Tech Solutions, Munich Re

IFRS 13 defines fair value as the price received to sell an asset in an orderly transaction between market participants at the measurement date [1]. It further establishes a hierarchy that gives priority to observable inputs. Level 1 inputs are unadjusted quoted prices in active markets for identical assets accessible at the measurement date. Barkr incorporates such inputs where the relevant criteria are satisfied, while also using observable historical transactions, comparable-system prices and technical characteristics. Forward-looking statements and market-participant assumptions also enter through the stochastic model.

The IVSC market approach similarly directs the valuation toward identical or comparable assets for which price information is available [2]. USPAP supplies complementary discipline in asset identification, scope of work, data analysis, documentation and the treatment of extraordinary assumptions [4]. These standards establish the conceptual boundary of the analysis: current fair value is grounded in market evidence; future values are conditional distributions derived from explicit assumptions.

*"Uncertainty is one of the biggest constraints on capital formation in asset-backed lending, and asset valuation is often at the center of that uncertainty. Independent valuation frameworks built on resale data play an important role in improving transparency and confidence across the AI infrastructure market", **Joshua Stevens, Macquarie.***

Barkr's dataset consists of 300 million observations across hard-asset categories, including more than 1.6 million GPU-specific data points. The dataset combines proprietary primary and secondary transactions, configured-system indications, asset attributes, condition and market context. This breadth is important because public GPU listings often conflate individual accelerator cards with integrated enterprise systems. The latter incorporate processors, high-bandwidth memory, NVLink or other interconnect, networking, cooling, power delivery, storage and validated software configurations. System-level valuation therefore requires matching the economic unit that a buyer can actually deploy.

The research addresses three questions. First, how should fair value be modeled when the asset is technically complex and transaction evidence is distributed across primary and secondary markets? Second, how should useful life and forward price uncertainty be represented across relatively mature and frontier architectures? Third, how do manufacturing constraints, infrastructure lead times and the emerging financialization of GPU ownership affect the distribution of recovery values? The remainder of the research develops a regime-switching valuation method and interprets the resulting ranges in the context of current market structure.

2. Illustrative assets and economic life

This report is grounded in observable NVIDIA transaction data, with the H100 and GB300 NVL72 illustrating the framework across a mature and a current generation. Barkr's capabilities extend to multiple brands across many hard-asset classes wherever comparable transaction data is available. We selected the NVIDIA H100 and GB300 NVL72 models because they illustrate different stages in the lifecycle and have deep observable transaction data due to NVIDIA's dominant market share. NVIDIA describes DGX H100 as an integrated eight-GPU enterprise platform [5]. In the present analysis, H100 represents a relatively mature generation: it has an established installed base, a deeper transaction history and broad compatibility with the CUDA software ecosystem. GB300 NVL72 represents a newer NVIDIA architecture. NVIDIA describes it as a liquid-cooled rack-scale system integrating 72 Blackwell Ultra GPUs, 36 Grace CPUs and high-speed networking [6]. Its value reflects frontier performance, density and scarcity, together with greater dependence on specialized power, cooling and facility design.

Barkr's September 2026 reference values are approximately \$320,000 for a configured eight-GPU H100 system and \$4 million for GB300 NVL72 (based on observations before NVIDIA's Aug 23, 2026 price increases). Public configured-system and procurement indications provide external reasonableness checks [7, 8]. The valuation itself is based on Barkr's transaction evidence and proprietary valuation modeling. The starting values are therefore measurement-date estimates, whereas the curves shown below are forward distributions conditional on the model parameters.

Barkr estimates a five-to-six-year useful life for NVIDIA H100 and a nine-to-ten-year useful life for GB300 NVL72. Economic life is determined by the period during which the system remains productive with a workload cascade enabled by architectural fungibility. A platform that is displaced from frontier training can remain valuable for inference, fine-tuning, retrieval, batch processing, simulation and other workloads for which software compatibility, immediate availability and total cost are more important than peak benchmark performance. J.P. Morgan cites continued use and repurposing of older accelerators, including positive A100 operating margins beyond the first two or three years [9]. Secondary-market reporting likewise records continuing values across earlier accelerator generations [10].

Barkr's current observations are consistent with this mechanism. Selected prior-generation configurations have maintained, and in some periods increased, transaction

value as demand for production workloads expanded and immediately available systems remained scarce. Generation age is therefore an incomplete sufficient statistic for price. Installed compatibility, memory configuration, interconnect, location, operating condition, software support, removal cost and the breadth of the buyer universe jointly determine the value of the system.

3. Stochastic valuation methodology

Forward valuation is implemented as a straightforward state-dependent Monte Carlo process. Let V_t denote the market value of the configured system and $S_t \in \{C, B, D, S\}$ the prevailing market regime: correction and rebalancing, balanced market, demand squeeze or realized supply shock. Conditional on $S_t = j$, the monthly log-price change follows $\ln(V_{t+1}/V_t) = \mu_j + \sigma_j \varepsilon_t$, where ε_t is standard normal. The resulting conditional price distribution is lognormal, remains positive and permits both expected return and volatility to vary with market structure. This specification follows the regime-switching logic introduced by Hamilton and subsequently developed for financial prices and conditional volatility [22–25].

Regime transitions are endogenous to simulated elasticity conditions. At each month, Barkr draws a standardized demand innovation D_t and a supply-inelasticity innovation E_t from a correlated bivariate normal distribution with correlation 0.25. Their time-varying means represent the current imbalance: elevated hyperscaler demand and constrained effective supply are assumed to moderate gradually as fabrication, packaging and commissioned data-center capacity improve. Transition probabilities are calculated through a multinomial-logit transformation of a persistent baseline transition matrix. Positive D_t increases the probability of demand squeeze; positive E_t increases the probability of realized supply shock; simultaneous positive realizations reinforce the probability and persistence of scarcity.

This architecture matters economically. A favorable percentile drawn from a single process with negative drift ultimately falls below the starting value because the drift term grows linearly with time while dispersion grows with the square root of time. The revised specification therefore separates scenario centers from the unconditional distribution.

Table 1. Illustrative regime calibration used in the Monte Carlo simulation.

Market regime	Annual central path	Annualized volatility	Economic interpretation
Correction / rebalancing	67.4% value retention	8.7%	Demand normalizes faster than committed capacity; financing and utilization weaken.
Balanced market	93.1% value retention	6.2%	Production improves and demand remains orderly; technical aging is the principal value driver.
Demand squeeze	4.9% appreciation	8.7%	Hyperscaler and enterprise requirements absorb incremental production and commissioned capacity.
Realized supply shock	12.7% appreciation	12.1%	Fabrication, packaging, minerals, power or data-center constraints produce a persistent scarcity premium.

For each asset, Barkr generates 100,000 monthly paths using a fixed random seed to support reproducibility. The initial regime allocation is 8% correction and rebalancing, 42% balanced market, 32% demand squeeze and 18% realized supply shock. These allocations are parametric assumptions informed by observed market behavior, historical transaction patterns and forward-looking market-participant expectations. Although the model extrapolates beyond currently observable data, each regime corresponds to a market state for which tangible evidence has emerged during the preceding years.

The correction and rebalancing regime reflects conditions already observed when new product generations entered the market, buyers deferred procurement, financing conditions tightened or available inventory temporarily exceeded immediate deployment capacity. Under this regime, technological substitution and weaker utilization progressively reduce transaction values. The low-range conditional median exhibits approximately 67.4% annual value retention. It represents a sustained correction rather than a single-period price decline.

The balanced-market regime reflects periods in which expanding production capacity is broadly absorbed by underlying demand. Transaction values decline gradually as systems age, but immediate availability, software compatibility, established operating performance and the migration of older accelerators toward inference, fine-tuning and batch-processing workloads provide continuing support. Evidence from the preceding two years indicates that selected prior-generation systems have remained economically productive and, in certain periods, maintained or even increased transaction values despite the introduction of newer architectures.

The demand-squeeze regime represents circumstances in which hyperscaler procurement, sovereign AI initiatives, cloud-platform expansion and enterprise adoption grow faster than deliverable supply. Elements of this regime are already visible in large capital-expenditure programs, expanding order pipelines, sustained demand for GPU financing and the absorption of incremental production before it becomes broadly available in secondary markets. Under these conditions, buyers assign a premium to systems that can be delivered, installed and monetized immediately. The last two years have leaned towards that scenario, with secondary market value of several GPU classes (including NVIDIA A100 and NVIDIA H200) increasing rather than depreciating.

The realized supply-shock regime reflects constraints extending beyond accelerator production. During the preceding two years, market participants have encountered limitations in advanced packaging, high-bandwidth memory, networking equipment, power components, transformers, cooling systems, construction capacity and grid interconnection. Critical-mineral concentration, export controls and geopolitical uncertainty have added further sources of supply risk. Other supply-chain disruptions such as the ones observed during the 2020 pandemic can also contribute to immediate shocks. A completed GPU system becomes productive only when the entire infrastructure stack is available, disruption at any critical layer can therefore constrain effective supply and support transaction prices. The upper conditional median models the continuation or recurrence of these observed conditions that can lead to an annual appreciation of 12.7% in prices for existing models.


The displayed lower and upper curves are scenario-conditioned medians. They show the central simulated trajectories if the correction or realized supply-shock regime governs the valuation horizon. The indicative fair-market-value line is the median of the complete switching mixture: the value above and below which one-half of all evolving-regime outcomes fall. This central estimate allows individual paths to transition among regimes as demand, production, financing and infrastructure conditions evolve.

This presentation distinguishes economic scenarios from statistical coverage. The conditional medians preserve the original lower and upper projections, while normally distributed innovations generate dispersion around each regime trajectory. The measurement-date fair-value conclusion remains anchored to observable transaction evidence, including Level 1 inputs when the IFRS 13 criteria are met, and to documented market-participant assumptions under IFRS 13 and USPAP.

The framework constitutes a structured, forward-looking sensitivity analysis grounded in market states that have already been partially or temporarily realized. Normal innovations and Markov transitions provide a tractable representation of discontinuous events such as export controls, product-generation transitions, concentrated buyer demand, project cancellations and infrastructure outages. Barkr periodically refreshes starting values, regime parameters, transition probabilities, elasticity-shock assumptions and useful-life estimates as new transactions, technical information and supply-chain evidence become available. This measurement-date discipline preserves practical business relevance while maintaining an auditable separation between historical observations and forward-looking assumptions.

Evolution of simulated market-regime allocation

Monthly transitions respond to correlated Gaussian demand and supply-inelasticity shocks




Source: Barkr regime-switching Monte Carlo valuation model.

Figure 1. Evolution of the simulated regime allocation. Shares represent the proportion of paths assigned to each market state after transition probabilities respond to normally distributed demand and supply-inelasticity shocks. Improving capacity gradually reduces but does not eliminate, the scarcity states.

H100 simulated fair-market-value distribution

Illustrative useful life: 5–6 years | 100,000 simulated paths




Source: Barkr Monte Carlo valuation model.

Figure 2. Simulated H100 valuation from August 2026. Curves represent the low-range conditional median, the median fair-market-value reference across the full evolving mixture and the realized supply-shock conditional median.

GB300 NVL72 simulated fair-market-value distribution

Illustrative useful life: 9–10 years | 100,000 simulated paths



Source: Barkr Monte Carlo valuation model.

Figure 3. Simulated GB300 NVL72 valuation from August 2026. Curves represent the low-range conditional median, the median fair-market-value reference across the full evolving mixture and the realized supply-shock conditional median.

4. Simulation results and asset-value interpretation

At the six-year NVIDIA H100 horizon, the low-range conditional median, full-mixture median and realized supply-shock conditional median values are approximately \$30k, \$384k and \$655k, respectively, from a \$320k starting valuation. At the ten-year GB300 horizon, the corresponding values are approximately \$78k, \$4.6m and \$13.2m, from a \$4m initial valuation. The upper path remains above the measurement-date value because it represents a durable, realized supply shock. The median fair-market-value path incorporates transitions among all four regimes. The lower range represents the minimum expected value, or the 'marketable cash value' price that Barkr has historically guaranteed.

The NVIDIA H100 result illustrates why an older generation can remain financeable. Its physical asset value is supported by an established software environment, known operating characteristics and a large and growing set of applications that do not require frontier performance. The GB300 result illustrates a different combination: high initial performance, greater infrastructure dependence and a longer potential migration path across workloads. The comparison is informative because it shows that useful life and market value arise from workload compatibility and deployability rather than from a simple chronological ranking of chip generations.

A central implication is that chip value and compute value can be relatively decorrelated. Chip value is the transaction price of transferable hardware. Compute value is the price of temporary service output and varies with utilization, contract term, location, energy price, networking, storage, software and operator execution. Barkr's prior April report documents that rental and equipment-price signals can move at different speeds [3]. A decline in rental price can weaken borrower cash flow while the underlying hardware continues to command a resilient price among buyers able to redeploy it. Conversely, high rental prices can reflect temporary utilization or service scarcity without establishing an equivalent increase in asset value.

This distinction has direct implications for credit analysis. Debt-service capacity should be evaluated using productivity metrics such as operating revenue and utilization as well as total opportunity of contractable offtake.

Collateral recovery should be evaluated using durability, fungibility and standardisation metrics such as equipment transactions, condition, remaining useful life, removal costs and operator and buyer depth. Combining the two measures is appropriate at the portfolio level, collapsing one into the other is not. The chip is a productive, durable and fungible asset, while compute is the valuable service flow generated by that asset.

5. Supply elasticity, demand growth and scarcity premia

The supply of deployable AI capacity is governed by a vertically integrated production function. It begins with mineral extraction and refining and extends through wafer fabrication, high-bandwidth memory, advanced packaging, boards, networking, power electronics, cooling systems, construction and grid interconnection. As outlined in our August report, a GPU system becomes economically productive only after the entire stack is assembled, installed, energized and commissioned. The lead time of the slowest component therefore constrains the effective supply elasticity of the completed system.

Production capacity is improving. NVIDIA reports that Vera Rubin is ramping into full production across hundreds of supply-chain partners and factories [19]. This expansion is economically important, and it occurs against a rapidly growing demand base. Goldman Sachs estimates annual AI capital expenditure of approximately \$765 billion in 2026, rising to \$1.6 trillion by 2031 [20]. Hyperscalers, AI laboratories, cloud providers and sovereign programs compete for the same accelerators, memory, network equipment, transformers, cooling systems and construction resources. Incremental production can therefore be absorbed before it creates excess deliverable supply.

The timing mismatch is material. Semiconductor fabs, packaging lines, mines, refineries, substations, transmission systems and large data centers require multi-year development. Demand for chips and commissioned capacity can accelerate over quarters. Goldman Sachs estimates that only 50 to 60% of data-center capacity scheduled over the next one to two years will arrive on time and cites an 18–24-month construction period after permitting [11]. JLL reports grid-connection timelines of approximately four years in the United States and close to five years across major markets [12]. The IEA estimates that grid constraints could delay approximately 20% of planned global capacity through 2030 [13].

Barkr's September 2026 capacity analysis produces a consistent result. Expected U.S. live capacity in 2027 is estimated at 75.1 to 80.9 GW, compared with 94.3 GW required to preserve the 2025 operating cushion. The implied 13.4 to 19.2 GW gap remains after a 14% reduction in operating demand [14]. The relevant scarcity is thus not confined to the number of chips fabricated. It concerns commissioned capacity capable of converting those chips into marketable compute.

Upstream constraints add a second layer. AI infrastructure is copper-intensive and depends on strategic minerals used in semiconductors, power electronics, networking, motors and grid equipment. The IEA projects a continuing copper supply deficit through 2035 and reports increasing concentration across critical-mineral refining [15]. Rare-earth supply is particularly concentrated: China accounted for approximately 60% of magnet rare-earth mining and 91% of refining in 2024 [16]. Gallium, germanium, graphite and other strategic materials exhibit comparable geopolitical exposure. Long mine-development cycles, technically complex refining and export controls restrict short-term adjustment.

These conditions generate a scarcity premium for immediate control of deployable hardware. The relevant competitive advantage is the combination of physical access and balance-sheet liquidity: firms able to secure systems today avoid the uncertainty embedded in future production, construction and interconnection schedules. Barkr currently observes strong demand for credit to finance GPU infrastructure, while market participants report a backlog of equipment and capacity requirements that current production and commissioning capacity cannot fully satisfy.

A correction can preserve part of this scarcity mechanism. Lower utilization may moderate near-term demand, while tighter financing and project cancellations simultaneously reduce future capacity additions. Because prospective supply exits before commissioned assets do, realized supply can remain below announced pipeline figures. In the simulation, these interactions appear through evolving transitions: weaker demand innovations raise the probability of correction or balance, while simultaneous demand pressure and adverse supply-elasticity realizations move paths toward demand squeeze and realized supply shock. The upper valuation boundary is therefore generated by coherent histories of persistent scarcity rather than by a favorable draw around an otherwise declining trend.

6. Financialization and the emergence of GPUs as an asset class

The institutional structure of the market is evolving alongside the physical asset. GPU-backed loans, equipment facilities, special-purpose vehicles, usage-linked contracts and asset-backed pools increasingly separate hardware ownership from compute operation. These arrangements resemble established forms of infrastructure and equipment finance: investors fund a productive, durable and fungible asset, operators monetize its offtake service flow, and lenders underwrite both contracted cash flow and collateral value.

Financialization strengthens the need for a transaction-based valuation architecture. Asset-backed lenders require a measurement-date fair value, a forward distribution of potential recoveries and a documented link between observable evidence and model assumptions. The principal monitoring variables follow directly from the model. They include equipment transaction prices, secondary inventory, hyperscaler capital expenditure, credit demand, order backlogs, fab and packaging output, critical-mineral policy, interconnection queues, transformer and cooling lead times, commissioning volumes, software support and performance per watt. Changes in these variables can affect the starting price, drift, volatility and useful-life assumptions. Continuous recalibration is therefore a constitutive part of the valuation process rather than a post-model adjustment.

*"NVIDIA AI factories are productive, durable and fungible. They turn energy into intelligence, continuously improve through software, and can serve any workload, any offtaker and every industry. That makes them an investable asset class," said **Nico Caprez, Vice President, AI Infrastructure Growth at NVIDIA.***

"Barkr is enabling the asset-backed financing flywheel for NVIDIA AI factories. It starts with data that allows lenders to underwrite the residual value of the asset—not only the long-term contract. And greater confidence in residual value attracts more capital to build the infrastructure of intelligence at scale."

7. Conclusion

NVIDIA GPU systems should be valued as productive, durable and transferable physical assets whose service output is compute. The distinction is economically consequential. Compute-access prices are high-frequency operating signals, chip and system prices are transaction-based asset signals. The NVIDIA H100 and GB300 NVL72 illustrate how that framework applies across a relatively mature and a newer generations. Older systems can remain resilient through workload additions and migrations and immediate availability, while frontier systems can command scarcity premia associated with performance density and constrained deployment.

The regime-switching Monte Carlo framework converts market-structure assumptions into conditional forward distributions. It presents the low-range conditional median, the full-mixture median and the realized supply-shock conditional median as the lower projection, indicative fair-market-value reference and upper projection. It thereby preserves the original scenario economics while separating observable measurement-date evidence from uncertainty about future market states. In current conditions, sustained demand growth, hyperscaler pipelines and the long lead times of manufacturing and physical infrastructure create upward pressure on the value of deployable systems. The concurrent growth of GPU-backed finance for NVIDIA systems reinforces the treatment of the chip as an asset and supports the emergence of NVIDIA GPU infrastructure as a distinct institutional asset class.

References

- [1] IFRS Foundation. IFRS 13 Fair Value Measurement, 2026. Fair-value definition, market-participant assumptions and fair-value hierarchy.
- [2] International Valuation Standards Council. Standards Glossary, 2026. Market approach, market value and income approach; accessed 26 August 2026.
- [3] Barkr Research. GPU Resale vs Rental: Robust, Divergent Markets. Technical report, April 2026.
- [4] The Appraisal Foundation. Uniform Standards of Professional Appraisal Practice, 2026.
- [5] NVIDIA. NVIDIA DGX H100 Datasheet. Technical report, 2022.
- [6] NVIDIA. NVIDIA GB300 NVL72. Official product and system specifications; accessed 27 August 2026.
- [7] Exxact Corporation. TensorEX 8U HGX H100 Server. Public configured eight-GPU H100 system indication; accessed August 2026.
- [8] Data Center Dynamics. Apple to Spend \$1bn on Nvidia GB300 NVL72 Systems: Report, March 2025. Reported \$3.7–\$4.0 million procurement estimate.
- [9] Michael Cembalest. Smothering Heights: Eye on the Market Outlook 2026. J.P. Morgan Asset Management, January 2026.
- [10] ServerBuyback. GPU AI Accelerator Resale Value Index, June 2026.
- [11] Goldman Sachs Research. US Data Center Power Demand Projected to Double by 2027, 20 May 2026.
- [12] JLL Research. 2026 Global Data Center Outlook; and Where Energy Meets Property, 2026.
- [13] International Energy Agency. Energy and AI, 2025; grid-constraint capacity analysis accessed August 2026.
- [14] Barkr Research. Valuing Live GPU Infrastructure. Technical report, August 2026.
- [15] International Energy Agency. Global Critical Minerals Outlook 2026, July 2026.
- [16] International Energy Agency. Rare Earth Elements and Critical Minerals Security Programme, 2025–2026.
- [17] U.S. Geological Survey. Mineral Commodity Summaries 2026 and Rare Earths Statistics and Information.
- [18] Office of the Comptroller of the Currency. Allowances for Credit Losses: Comptroller's Handbook, 2025.
- [19] NVIDIA. Vera Rubin Ramps Into Full Production to Power Agentic AI Factories Worldwide, 31 May 2026.
- [20] Goldman Sachs Research. The Assumptions Shaping the Scale of the AI Build-Out, 1 May 2026.
- [21] NVIDIA. AI Factory Compute Is Becoming an Investable Asset Class; and partnerships to mobilize over \$500 billion of third-party capital, August 2026.
- [22] J. D. Hamilton. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–384, 1989. DOI: 10.2307/1912559.
- [23] S. F. Gray. Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process. *Journal of Financial Economics*, 42(1):27–62, 1996. DOI: 10.1016/0304-405X(96)00875-6.
- [24] E. S. Schwartz. The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *The Journal of Finance*, 52(3):923–973, 1997. DOI: 10.1111/j.1540-6261.1997.tb02721.x.
- [25] H. Bessembinder, J. F. Coughenour, P. J. Seguin and M. M. Smoller. Mean Reversion in Equilibrium Asset Prices: Evidence from the Futures Term Structure. *The Journal of Finance*, 50(1):361–375, 1995. DOI: 10.1111/j.1540-6261.1995.tb05178.x.
- [26] A. Haramboure et al. Vulnerabilities in the Semiconductor Supply Chain. OECD Science, Technology and Industry Working Series 2023/05, 2023. DOI: 10.1787/6bed616f-en.