AI Exposure

Implications Across Asset Classes

CONTRIBUTORS
Frank Benham, CFA, CAIA
Steve MacLellan, CFA

WHITEPAPER | MAY 2026

Institutional investors have become increasingly focused on concentration in their public equity portfolios, particularly among diversified US technology companies (e.g., the "Magnificent Seven"). While some investors may regard this concentration as simple company-specific risk, others are looking at this phenomenon as one channel of what has become a cross-asset driver: artificial intelligence ("AI"). A recent survey of 800 institutions found that 63% identified artificial intelligence as the single biggest megatrend influencing their investment decisions over the next five years.[1] This exposure to AI that is prevalent in the US stock market also appears, through different mechanisms, in private equity, private credit, real estate, infrastructure, and fixed income. A driver that large, propagating through that many asset classes, is worth understanding at a scale that matches its reach.

This paper makes the case that artificial intelligence exposure should be evaluated at the total portfolio level rather than just within individual asset classes. It outlines the channels through which AI-linked exposure has entered a typical institutional portfolio, examines the considerations that any measurement framework needs to address, presents an illustrative example, and closes with the implications for long-horizon investors. The exercise cuts both ways: for some investors the relevant finding will be that cumulative exposure is larger than intended; for others it will be that a diversified, cross-asset AI position is smaller than a growth-focused, long-horizon portfolio should carry.

Key Takeaways

- **AI exposure is a whole-of-portfolio question, not an equity-only question.** The driver propagates across public equity, fixed income, private credit, real estate, infrastructure, and private equity through different mechanisms, each with its own risk-return profile. Evaluating it within any single asset class understates its importance to the total portfolio.
- **The sum of locally rational allocation decisions can produce cumulative exposure that no one has explicitly chosen.** Portfolios governed asset class by asset class lack a natural aggregation layer for thematic drivers, and the cumulative exposure that results is often invisible in standard risk reporting.
- **Scope choice is consequential and must be applied consistently.** A tight definition and a wide definition applied to the same index can produce AI exposure estimates that differ by a factor of two or more. Whatever scope is chosen, applying it differently across asset classes produces a picture that is internally inconsistent regardless of whether individual measurements are accurate.
- **Public equity exposure is easier to measure, while private market exposure generally must be estimated.** Any useful measurement approach should make this gradient explicit rather than produce a single number that conceals it. Ranges and confidence intervals are more informative than false precision.
- **The goal of measurement is deliberate decision-making, not a mandate to act.** An investor who measures their cumulative AI exposure and concludes it is within their intended range has made a conscious choice. The governance obligation this exercise creates is not to change the portfolio; it is to decide whether to, rather than simply accepting whatever the portfolio contains.

The Definitional Problem

Before mapping where AI exposure appears across a portfolio, a definitional question must be addressed: what do we mean when we say "AI"? Reasonable investors can draw the boundary between AI and non-AI in different places, and the boundary choice has substantial implications for the resulting measure of portfolio exposure. A tight definition and a wide definition, applied to the same portfolio, can produce AI exposure estimates that differ meaningfully (e.g., by a factor of up to 1.6 times for public equities).

A rough continuum runs from tight to wide:

- A tight definition includes only direct participants in the AI value chain: the firms designing and manufacturing computational resources ("compute"). These types of firms include but are not limited to semiconductor designers, foundries, the companies building and deploying foundation models, and the hyperscalers[2] providing the cloud infrastructure on which AI runs.
- A moderate definition adds firms that benefit from capital expenditures of the "tight" AI companies: the power generators, transmission operators, grid equipment makers, specialized data center developers, and critical materials producers whose revenues depend on the AI buildout continuing.
- A wide definition further includes indirect beneficiaries and adversely affected parties: enterprise software firms integrating AI capabilities, industrial automation and robotics companies, AI-enabled incumbents across sectors whose competitive positioning depends on adoption, and business services firms whose productivity may accelerate meaningfully with AI deployment or be disrupted via creative destruction.

One widely-referenced approach to categorizing AI exposure distinguishes four phases of the value chain: 1) creators (foundation model labs such as OpenAI, Anthropic, and Google DeepMind); 2) enablers (the semiconductor, cloud, and data center infrastructure that makes AI possible, including NVIDIA, TSMC, and the hyperscalers); 3) monetizers (firms charging for AI features, such as Microsoft Copilot, Salesforce Agentforce, and ServiceNow); and 4) beneficiaries (businesses that gain productivity from AI adoption without selling AI products directly).[3] A useful extension distinguishes a fifth category of potentially disrupted businesses, namely traditional seat-based software firms whose pricing models face structural pressure from AI-driven alternatives (see Figure 1).

Figure 1
A Five-Category Framework for AI Exposure

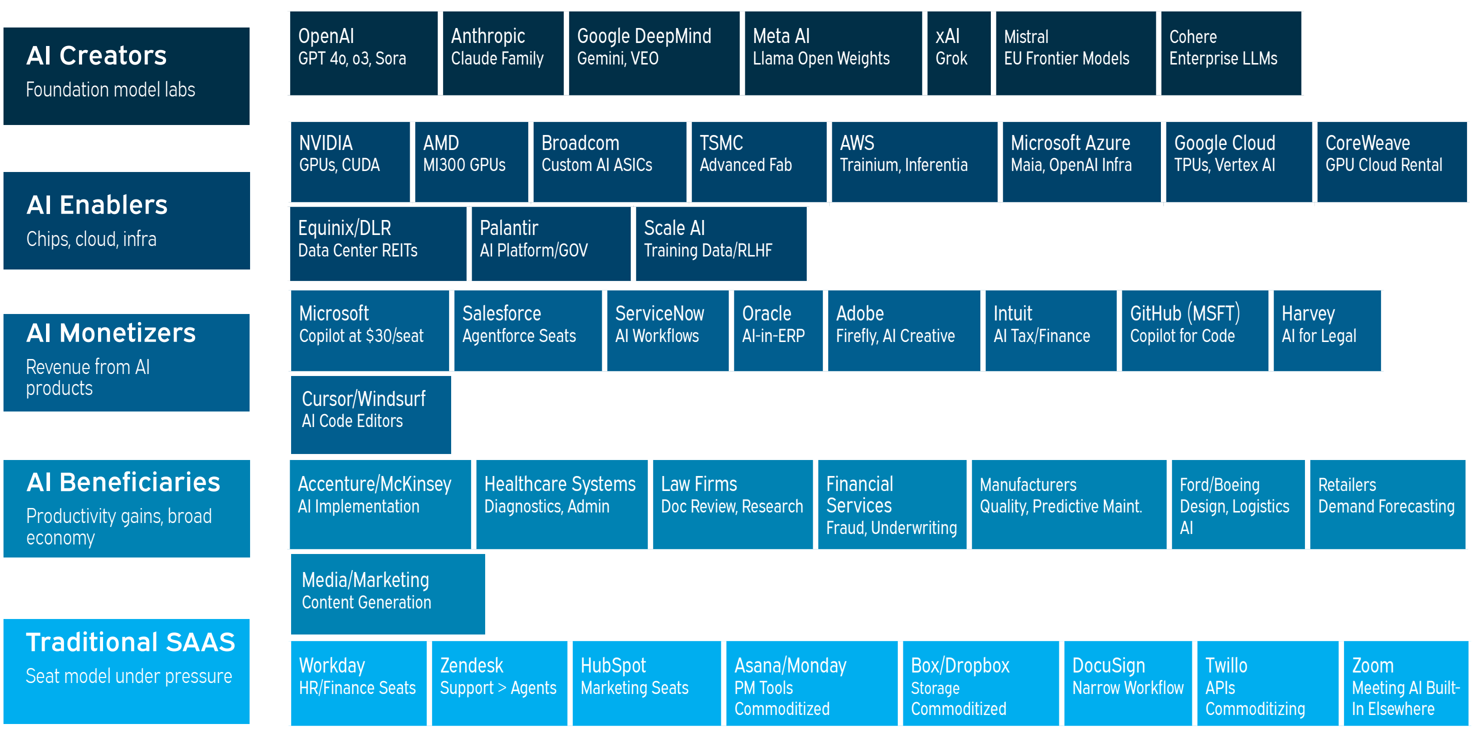

Source: Meketa Investment Group, 2026.

This five-category structure is one reasonable way to determine an answer to the "What is AI?" question.[4] Importantly, though, it is not the only way. Research from the Bank for International Settlements, for example, restricts the AI-firm universe to a handful of hyperscalers for macroeconomic analysis.[5] Index data providers use their own taxonomies with different boundaries as well.[6]

Two observations emerge from applying this framework to a rapidly evolving field. First, the same company often appears in multiple categories. Microsoft is an enabler through Azure, a monetizer through Copilot, and, through its OpenAI partnership, a partial creator. Alphabet acts as an enabler through Google Cloud and a creator via Google DeepMind. A framework that does not accommodate these overlaps may misclassify the largest and most consequential names. Second, the category boundaries are porous and likely to shift as the landscape evolves. Firms that look like monetizers today may become beneficiaries as AI features become widely available and seat-based pricing comes under pressure. Firms that act as enablers today may see their moats erode as alternative chip architectures and open-weight models gain ground. Any framework that treats categorizations as permanent will age poorly.

No framework is fully adequate across all circumstances. The relevant scope of a framework should be determined by the investor based on what governance question is being asked. An investor trying to understand concentration risk from the AI capex cycle may want a narrower scope centered on direct participants and capex beneficiaries. An investor trying to understand total exposure to productivity gains and potential disruption from AI adoption may want a wider scope. Different questions demand different scopes.

Ultimately whatever scope is chosen, it is important to apply it consistently across asset classes. A tight definition applied to public equity and a wide definition applied to private markets produces a picture that is internally inconsistent regardless of whether the individual measurements are accurate. Much of the difficulty in comparing published AI exposure estimates across different investors (e.g., asset owners, asset managers, strategic investors) stems from using different definitions that are themselves applied unevenly. While it is not feasible (and perhaps not desirable) to have all investors use perfectly consistent measures of AI exposure, an investor doing this work for their own portfolio can and should be disciplined about consistency.

Channels of Exposure Across an Institutional Portfolio

Public Equity

Public equity is often the most visible channel and the asset class where analysis is most frequently conducted. The top seven companies in the S&P 500 index now represent roughly 30% of the US stock market's capitalization, with information technology earnings growth driving much of recent index performance.[7] The specific names concentrated in the top of the index (the Magnificent Seven plus a handful of semiconductor designers and AI-adjacent firms) are the direct public-market expression of the AI capex cycle.

Public equity exposure to AI is measurable at the security level, and the definitional question is correspondingly consequential. Applied to the MSCI USA Index, a tight-scope estimate that includes only direct AI value chain participants (e.g., AI-direct semiconductors and semiconductor equipment, the hyperscaler cloud providers, and the foundation model developers) produces AI-linked exposure of approximately 34% of the index.[8] A moderate-scope estimate that adds capex-cycle beneficiaries (e.g., power generation tied to data center demand, data center real estate, AI-monetizing software platforms, networking and IT services) produces exposure of approximately 48%. A wide-scope estimate that further includes AI-enabled incumbents across industrials and health care can reach approximately 54%. All three figures describe the same index.

For most institutional investors, public equity is also the asset class where exposure has grown most visibly over the past several years, driven less by deliberate allocation decisions than by market-cap-weighted benchmarks. An investor who has held a consistent allocation to US large cap equities has seen their AI-linked exposure rise substantially without taking an active decision to increase it.

Non-US public equities carry AI exposure that is structurally different from the US case and is often underappreciated. In China, the AI buildout has proceeded largely through domestic hyperscalers (e.g., Alibaba, Tencent, Baidu) and a wave of foundation model developers (e.g., DeepSeek). Institutional investors with allocations to Chinese equities likely hold exposure to this value chain, though the regulatory, geopolitical, and data-access constraints on measurement are more significant than in the US case. Elsewhere in Asia, South Korean and Taiwanese names (e.g., Samsung, TSMC) represent substantial semiconductor and memory exposure that is directly linked to global AI capex. European equity benchmarks carry more modest direct AI exposure but meaningful indirect exposure through ASML and AI-enabled industrials. In aggregate, non-US large cap exposure to direct AI participants is meaningfully less concentrated than in the US.

Combining US and non-US equities into a global portfolio proxied by the MSCI ACWI Index produces lower figures across all three scope definitions than for the US alone: approximately 26% under the tight scope, 37% under the moderate scope, and 41% under the wide scope.[9] The reduction relative to the US-only index reflects the lower AI concentration in non-US developed markets, where pharmaceuticals, banks, and consumer staples represent a larger share of total market capitalization. A globally diversified equity allocation will have meaningfully lower AI exposure than a US-only portfolio, though the absolute level remains substantial under any of the three scope definitions.

Public Fixed Income

Public fixed income is the channel many institutional investors have not considered as AI-relevant, but this is changing. Hyperscalers turned to corporate debt markets at scale to help meet their capex needs starting in 2025.[10] Within investment-grade corporate credit, hyperscaler and AI-related issuers have become a meaningful and growing share of new issuance.[11]

An investor holding a broad investment grade portfolio now has hyperscaler credit exposure that was not there three years ago. The credit quality of the underlying issuers remains high, but concentration has risen, and credit default swap spreads on lower-rated hyperscalers widened noticeably through late 2025, reflecting both the volume of supply and uncertainty around eventual project returns.[12]

A typical institutional fixed income allocation with Bloomberg US Aggregate index exposure will have perhaps 2-4% of the asset class holdings in hyperscaler and related AI-issuer credit, depending on scope and index selection.[13] This is a small share of the total portfolio but a non-trivial share of what has historically been the ballast component of that portfolio.

Private Equity

Private equity exposure to AI is more heterogeneous than public equity and meaningfully harder to measure. The asset class is often decomposed into three segments (buyouts, growth equity, and venture capital), each with different AI exposure profiles.

Buyout funds are the primary component of most institutional private equity programs and typically invest in mature, cash-generating businesses. AI influences these portfolios mainly through its impact on margins and business model disruption rather than through direct AI investment. Portfolio companies may use AI to expand margins, face competition from AI-enabled peers, or incorporate AI into their operations. The degree of AI involvement depends on the specific company and sector. Many mature buyout portfolios contain significant software holdings, particularly in seat-based and back-office SaaS, which may be more likely to be affected by disruption risk than to benefit directly from AI. Mature buyout portfolios, which reflect commitments going back at least ten vintage years, may understate both their positive AI-adjacent exposure and their AI disruption risk exposure, especially as the typical lag in buyout valuations implies that these exposures have not been fully priced in.

Growth equity captures mid-stage direct AI companies and has seen substantial AI-directed deployment over the past two years. Investors with meaningful growth equity allocations have more current AI exposure than those without.

AI exposure in venture capital (VC) is heavily vintage-dependent. Programs with meaningful commitments from 2022 onward are far more likely to carry current AI exposure than older programs: the foundation model era effectively began with the public release of GPT-3.5 in late 2022, and most funds raised before 2021 had already deployed the bulk of their capital in other opportunities before the current cycle. Recent vintages have become more heavily weighted in the AI theme. Nearly half of global venture funding in 2025 flowed into AI-related companies, with OpenAI and Anthropic alone capturing approximately 14% of global venture investment.[14]

VC exposure also provides access to a segment of the AI value chain that is not available through buyout or directly through public market channels. Several leading foundation model developers (e.g., OpenAI, Anthropic, xAI) have remained private well into their scaling phase, with valuations that have moved from venture-scale to large-cap-equivalent while still outside public markets. The same is true for the most AI-native application companies and for the infrastructure layer below the hyperscalers. For an investor seeking diversified AI exposure, VC is the primary route to this part of the value chain. That access comes with the usual illiquidity, J-curve, and return-dispersion characteristics of venture investing, but on a whole-of-portfolio basis, excluding VC from the AI exposure picture means excluding some of the most direct exposure to where the technology is being built.

Because many institutional private equity programs are buyout-heavy and VC exposure to AI is primarily in the most recent vintage years, the aggregate AI exposure in a typical mature private equity portfolio is lower than headline technology-sector allocations might suggest. An investor might find that 30-50% of their private equity portfolio is classified as technology but only 10-25% is meaningfully AI-adjacent under a moderate-scope definition, with the balance being legacy software and technology-enabled services that predate the current capex cycle.[15]

Private Credit

Private credit has grown more rapidly than almost any other institutional asset class over the past decade and has become a major channel for AI-related financing. Outstanding private credit loans to AI-related firms grew from near zero to more than $200 billion, with projections suggesting $300-600 billion by 2030.[16] The share of private credit funds with AI exposure is approximately 20% today, and the average AI share within those funds is approximately 5% of fund volumes.[17] More broadly, private credit lending to software companies could be as high as 30% of private credit exposure, depending on how it is defined.[18]

Two sub-channels matter. Direct lending to software companies and technology-enabled buyouts overlaps substantially with AI exposure, given that software is both where private credit has been most active and where AI disruption and adoption both concentrate. Data center asset-backed financing is the faster-growing sub-channel. Issuance of data center debt and related structured products surpassed $25 billion in 2025, exceeding the combined total of the previous three years.[19] The collateral and credit support structure also matters in a way it does not for equity, since a data center loan backed by the underlying real estate or supported by minimum lease commitments carries a meaningfully different risk profile than an unsecured loan to the same borrower.

Private credit AI exposure is challenging to measure. Fund strategy descriptions rarely specify AI share. Manager-reported sector breakdowns typically classify loans by industry rather than thematic driver. A bottom-up review is possible in principle but will require outreach to GPs and perhaps some estimation. An investor with a private credit allocation should assume there is meaningful AI exposure and estimate it with substantial uncertainty. Direct lending to software and technology-enabled buyouts plus data center asset-backed financing may represent 10-30% of a diversified institutional private credit portfolio, depending on manager selection and vintage.[20]

Real Estate

Data centers have transitioned from a niche sub-sector in real estate (and infrastructure) to a core theme over the past three years. Listed data center REITs (e.g., Digital Realty, Equinix) are often held in broad REIT portfolios. Exposure in the FTSE Nareit All Equity index to data centers is roughly 10%,[21] and some active REIT strategies could hold more. Private data center funds have proliferated rapidly, with transaction sizes consistently exceeding $10 billion in 2024 and 2025.[22] Broader private real estate funds increasingly include data center assets as well, though the average allocation in the NCREIF ODCE is less than 2%.[23]

Counterparty concentration in the channel is striking. A small number of hyperscalers account for most of the new hyperscale leasing volume. Single-tenant hyperscaler leases dominate new deliveries. An investor holding data center real estate exposure is, economically, holding exposure to hyperscaler credit quality and renewal intent, which means exposure to the same names that appear in their public equity and investment grade credit allocations.

Given the contrast in data center exposure between many REITs and core private market funds, the exposure will be driven by how much an investor allocates to REITs, core private funds, and funds focused on digital infrastructure. A typical institutional real estate portfolio with exposure to both listed REITs and private real estate funds may have 2–15% of the asset class in data center and related digital infrastructure, depending on manager selection and vintage.[24]

Infrastructure

Infrastructure may be the most underappreciated AI exposure channel in institutional portfolios (even setting aside direct data center holdings). The energy demands of the AI buildout are meaningful enough to have shifted power sector dynamics materially. US data centers are projected to account for nearly half of US electricity demand growth through 2030, and global data center electricity consumption is projected to approximately double from 2024 to 2030.[25]

The infrastructure required to support this demand is being built now, financed partly through institutional infrastructure funds. Power generation (e.g., natural gas combined cycle, nuclear restart and small modular reactors, renewables with storage), transmission and distribution, fiber networks, and data center platforms all appear in infrastructure portfolios.

The counterparty structure resembles real estate. Many of the largest infrastructure deals involve hyperscaler consortium members, hyperscaler offtake agreements, or hyperscaler-related tenant exposure. A typical institutional infrastructure allocation may have 10–30% of holdings in AI-adjacent themes, depending on manager selection and vintage, as well as how widely the scope is defined (e.g., to include public utilities that provide power generation for AI-driven load growth).[26]

Natural Resources and Commodities

The AI buildout has meaningful implications for copper (grid expansion, data center power distribution), uranium (nuclear generation for data center demand), natural gas (dispatchable power for data centers), and critical minerals used in semiconductor manufacturing. Institutional investors with dedicated natural resources allocations may have modest AI-linked exposure through these channels, depending on how widely defined that exposure is. For most investors, their allocation is small enough that the exposure does not materially affect total portfolio estimates.

Hedge Funds

Hedge fund exposure to AI has been meaningful but varies substantially by manager and mandate. Long/short equity managers have held large long positions in AI-related names. Systematic trend-following strategies have been persistently long US equities, which have been AI-driven. For most institutional investors, hedge fund AI exposure is better treated on a portfolio-by-portfolio basis, determining exposure by underlying asset type and security-level holdings.

Shared Counterparty, Shared Driver

The channels described above are distinct in their mechanisms and risk-return profiles, but several share the same ultimate counterparties. The hyperscalers appear as equity holdings in the public equity portfolio, as investment grade debt issuers in fixed income, as tenants and counterparties in the real estate portfolio, as consortium members and offtake counterparties in the infrastructure portfolio, and as off-balance-sheet sponsors in private credit. A portfolio that holds each of these positions at sizes deemed appropriate for that specific asset class may find they hold concentrated name exposure across asset class boundaries.

The same logic extends beyond formal counterparties to the broader supply chain: exposures can be embedded in equipment vendors, contractors, component suppliers, and critical input providers whose order books and pricing power ultimately depend on the same small set of AI spenders, even when there is no direct contractual relationship with a hyperscaler.

More broadly, the different asset class exposures to AI share a common driver: the AI capex cycle and the revenue it eventually generates. Public equity valuations, investment grade credit spreads, data center lease economics, infrastructure offtake contracts, and private credit collateral values all depend, in varying degrees, on the capex cycle meeting expectations. If it does not, the correlations among these positions in a drawdown will be higher than their location in separate asset classes might imply.

Considerations for a Measurement Framework

Scope

The first question a framework has to answer is where to draw the boundary. The earlier discussion of the definitional problem introduced the continuum from tight to wide definitions and noted that different scope choices can produce exposure estimates that differ by a factor of two or more. A tight scope that counts only direct participants in the AI value chain will produce a substantially lower exposure count than a wide scope that includes capex beneficiaries and AI-enabled incumbents.

For most institutions, a moderate scope may be the appropriate starting point: direct participants in the AI value chain plus their capex-cycle beneficiaries. This scope captures the exposure most relevant to concentration risk and drawdown scenario analysis without requiring judgment calls about which incumbents are sufficiently AI-enabled to include. Investors who want to understand portfolio-level exposure to AI-driven productivity gains can extend to a wider scope that includes enabled incumbents, but should treat that as a separate exercise with a separately reported number. Mixing scope levels within a single aggregate figure produces a number that is neither analytically clean nor comparable across periods. Whatever scope is chosen, it must be applied consistently across asset classes.

Exposure Type

The second question is what kind of exposure is being measured. The three types introduced earlier (direct ownership, counterparty exposure, and thematic exposure) have different behaviors and should not be aggregated without distinction.

Direct ownership is the most straightforward. For example, an investor holding shares of Microsoft has a direct claim on Microsoft's cash flows and participates in both the upside and the downside of the AI buildout through that claim.

Counterparty exposure is different. For example, an investor who has lent to a special-purpose vehicle financed by Microsoft, or who owns real estate leased to Microsoft, does not have a direct equity claim. The investor's return is capped at the agreed interest rate or lease yield. The downside, however, is exposure to Microsoft's willingness and ability to continue servicing its obligations, which may be affected by an AI disappointment even if Microsoft as a corporate entity remains solvent and continues to grow. Counterparty exposure has an asymmetric return profile that equity exposure does not.

Thematic exposure represents a third distinct category. For example, an investor holding equity in a utility or infrastructure fund that has committed to building gas-fired generation for data center load is not directly exposed to AI monetization, but is exposed to data center demand continuing.[27]

These three exposure types have different volatilities, different correlations with other assets, and different behaviors in stress scenarios. A framework is more useful if it measures them separately rather than summing them into a single number.

Opportunity and Risk Symmetry

A framework should be as useful for identifying exposure that exceeds what an investor intends as for identifying exposure that falls short of it. For many institutional investors, particularly those with large, mature private equity programs, the more interesting finding from a whole-of-portfolio measurement exercise may be the gap between their public equity AI exposure and their private markets AI exposure, with private markets likely running lower. That gap may be what the investor wants to see, or it may not. Either way, the framework needs to reveal it.

The opportunity side of the measurement exercise includes identifying the firms and themes most likely to benefit from continued AI investment, whether through direct participation, through capex cycle spillovers, or through productivity gains captured by AI-enabled incumbents. The risk side includes identifying concentrated exposure to a disappointment in the AI capex cycle, disruption exposure (i.e., businesses whose current revenue models may be eroded by AI-driven alternatives), as well as investments where these risks are attenuated through collateralization or capital structure positioning. Both sides matter. A framework that focuses only on opportunity treats AI as a directional bet to be sized. A framework that focuses only on risk treats AI as a threat to be hedged. Neither framing captures what a whole-of-portfolio view should show, which is the cumulative position the investor holds across both dimensions.

Measurability and Estimation

Data for public equity and public fixed income portfolios tends to be readily accessible at the security level. An investor with access to holdings data and a defined scope can calculate AI-linked exposure with reasonable precision for either asset class. For investors holding index funds rather than individual securities, the look-through exercise is the same in both cases: the AI-relevant names need to be identified and their weights in the index confirmed.

A category-based approach treats AI exposure as binary, classifying each holding as either in or out of a defined AI universe and summing the resulting weights to produce a portfolio-level number. An investor could choose to take an exposure-based approach instead, which treats AI risk as a continuous variable that varies across firms within the same category. Under this approach, a firm whose AI-related revenue, earnings, or capital expenditure represents a large share of its business carries proportionally greater exposure than a firm with only incidental AI involvement, even when both fall within the same category. The exposure-based approach is meaningfully harder to implement, since it requires firm-level data on AI-related revenue and capex that is rarely disclosed in a consistent or comparable form.

Private equity and debt exposure likely has to be estimated. Fund strategy descriptions do not consistently specify AI share at the company/loan level. Manager-reported sector breakdowns classify investments by industry or sub-sector (e.g., technology or software) rather than by thematic driver. A bottom-up review is possible in principle but requires substantial manager engagement and analytical effort. Even so, investors may need to estimate exposure for some funds based on coarser signals (e.g., fund strategy, vintage, manager focus).

Real estate and infrastructure exposure sits on a similar measurability gradient but with a distinct structure. Listed REIT and infrastructure exposure is measurable at the security level. Determining private real estate and infrastructure fund exposure will vary based on what the investor is trying to measure. For example, a real estate fund may report their exposure to data centers, but learning the counterparty structure for those data centers (e.g., hyperscaler tenants) will likely entail direct engagement with the manager. Likewise for power generation in infrastructure funds, where the offtake agreement may be directly tied to AI capex.

A useful framework makes this measurability gradient explicit rather than producing a single number that conceals it. Ranges and confidence indicators are more informative than false precision. An investor who reports their AI exposure as 18% when 12% is measured and 6% is estimated has produced a number that could potentially be misleading. An investor who reports a range of 15% to 22% with an indication of which components are measured and which are estimated has produced a number that more accurately reflects what is known.

Time Horizon and Evolution

The AI landscape is changing quickly, far faster than comparable shifts in previous technology revolutions. A framework built in 2026 and applied unchanged in 2028 is likely to misclassify a meaningful share of what it measures.

This has two implications for how a framework should be designed. The first is that the categorizations themselves should be treated as provisional. A firm classified as an AI enabler today may be classified as an AI monetizer tomorrow, or as an incumbent facing disruption the day after. The framework has to accommodate reclassification without forcing the investor to rebuild from scratch. The second implication is that the measurement exercise is not a one-time snapshot. It should be re-run periodically, and the governance response should be calibrated to changes over time rather than to the point-in-time number.

This consideration argues against heavy investment in any specific taxonomy in favor of a framework that addresses consistent dimensions of exposure. The dimensions themselves should hold steady across measurement periods so that exposures remain comparable over time. The specific classifications and category memberships within each dimension may shift, with some categories ceasing to apply and new ones emerging as the underlying landscape evolves.

Total-Portfolio Aggregation

The measurements for each asset class should be added together. That sum, expressed either as cumulative AI exposure as a share of the total portfolio or in the investor's risk-budget units, is the number the whole-of-portfolio argument rests on. A framework that stops at the asset class level leaves the work undone.

Aggregation is harder than it looks. Simple summation across asset classes can double-count if the same underlying exposure appears in multiple forms. For example, holding shares in Microsoft and owning a data center lease with Microsoft as tenant are different exposures, but they share a counterparty, and an aggregate number that does not recognize the overlap will misstate the level of concentration.

A useful framework provides at least a rough aggregation convention and flags the places where that convention oversimplifies. This is the point at which the exercise connects to existing portfolio risk reporting. Cumulative AI exposure expressed as a share of total assets is easy to compare to other concentration metrics the investor already tracks. Cumulative exposure expressed as a contribution to drawdown risk under a defined scenario is harder to produce but more directly relevant to governance decisions.

A Framework Designed Around These Considerations

For many institutional investors, the first iteration of this exercise may be a qualitative one. The measurement effort required to produce precise numbers across all asset classes is substantial, and for many investors the marginal value of precision may be lower than getting the qualitative picture right. An investor who has established that their cumulative AI exposure is meaningfully larger than, and different from their US equity exposure alone, without yet having a precise number, has already captured much of the governance benefit the exercise offers.

Correlated Exposure and Drawdown Scenarios

Traditional portfolio construction treats asset class diversification as a risk reducer, at least when asset classes are assumed to exhibit less-than-perfect correlations with each other. That assumption holds when the assets are responding to different economic drivers. However, it breaks down when a single driver dominates across all of them simultaneously. That is, the correlations among assets might prove higher than those assumed in a traditional modern portfolio throry(MPT) framework for those asset classes.[28] This uptick in correlations would be due to positions appearing diversified at the asset class level that carry concentrated exposure to the same underlying factor. Hence they behave more similarly due to this shared exposure, particularly when that factor is under stress.

Consider the following example of how an AI disappointment scenario could propagate through a portfolio. Public equity valuations of AI-linked names compress as revenue and earnings growth expectations reset. Private equity marks lag public markets but follow within a few quarters as comparable company valuations feed into manager valuations. Hyperscaler credit spreads widen as the scale of outstanding debt and the decline in free cash flow coverage become more visible. Data center lease economics come under pressure as hyperscaler tenants reassess renewal intent and as new supply continues to come online under leases signed in a more optimistic period. Infrastructure valuations compress for assets tied to data center power demand. Private credit may have to mark down net asset values as loans to AI-related borrowers show signs of stress and as the value of data center collateral declines. While this scenario is not certain, it is plausible and illustrates a potential cascade of impacts on positions in six or seven different asset classes that might otherwise behave more disparately.

The governance implication is not that investors should reduce their AI exposure. It is that investors whose reporting shows apparently diversified AI exposure in each asset class should be aware that diversification at the asset class level does not imply diversification of the underlying driver. A meaningful share of a portfolio's AI-related holdings across asset classes likely includes exposure to a single set of counterparties and a single capex cycle.

Investors who see a modest total-portfolio AI exposure number should consider whether that reflects a deliberate decision or simply an artifact of asset class decisions. For a long-horizon investor, a diversified cross-asset position in a structural productivity driver is not necessarily something to minimize. Measuring exposure is as useful for identifying opportunities to add exposure as identifying excessive concentration.

This measurement exercise is most valuable when the directional outcome is uncertain. Whether the AI buildout produces durable, outsized profits for the firms financing it or primarily benefits the businesses and consumers using AI tools downstream is a question that reasonable investors may answer differently. The whole-of-portfolio view supports either conclusion by giving the investor a clearer picture of what they hold.

An Illustrative Example

To illustrate the whole-of-portfolio approach, we apply it to a stylized large US public pension plan. The plan composition below is broadly representative of the asset allocation mix of a typical mid-to-large public pension plan, though individual plans vary substantially. The exposure estimates are directional and methodology-dependent. The purpose of the exercise is not to produce a number that any specific plan should rely on but to illustrate the gap between asset class-level and total-portfolio views.

The stylized plan holds approximately 40% public equity, 20% fixed income, 15% private equity, 10% real estate, 5% infrastructure, 5% private credit, and 5% hedge funds.

Public equity (40% of plan). Using a moderate-scope definition produces AI-linked exposure of approximately 48% of a US large cap equity allocation. A globally diversified equity portfolio has meaningfully lower exposure of approximately 37%. The actual figure depends on the plan's specific US versus non-US allocation. Applied to the 40% allocation, the contribution to total plan exposure is approximately 12-18%, with the lower end of the range reflecting a more globally diversified sleeve and the upper end reflecting a US-tilted sleeve.

Fixed income (20% of plan). Hyperscaler investment grade issuance and related AI-issuer credit represent perhaps 5% of a broad investment grade index, with additional exposure through corporate and agency holdings tied indirectly to the AI buildout. Applied to the 20% allocation, the contribution is approximately 1%.

Private equity (15% of plan). Most mature buyout-heavy private equity portfolios have 15-25% of holdings that are meaningfully AI-adjacent under a moderate-scope definition, with the balance being legacy software, services, and industrials that predate the current cycle. A plan with a meaningful growth equity or venture allocation would run toward the higher end. Applied to the 15% allocation, the contribution is approximately 2-4%.

Real estate (10% of plan). Depending on manager selection and allocation to REITs, data center and related digital infrastructure holdings may represent 2-15% of a diversified real estate portfolio, with the balance in traditional property types. Applied to the 10% allocation, the contribution is approximately 0.2-1.5%.

Infrastructure (5% of plan). AI-adjacent themes (e.g., power generation serving data center demand, transmission and grid, fiber, data center platforms) may represent 10-30% of an infrastructure allocation, depending on vintage and manager focus. Applied to the 5% allocation, the contribution is approximately 0.5-1.5%.

Private credit (5% of plan). Direct lending to software and technology-enabled buyouts plus data center asset-backed financing may represent 10-30% of a diversified private credit portfolio, though the share is growing rapidly as data center financing expands. Applied to the 5% allocation, the contribution is approximately 0.5-1.5%.

Hedge funds (5% of plan). AI-related long exposure through long/short equity and trend-following strategies is likely modest on a net basis but non-zero. We note this qualitatively rather than attempt a numerical estimate.

Total Portfolio Aggregation

Summing the contributions under a moderate-scope definition, the stylized plan has cumulative AI-linked exposure in the approximate range of 18-26% of total plan assets, compared to a public equity headline figure of 12-18%. As a result, the aggregate number could be more than 10% higher than when viewed solely through a US equity lens.

The ranges are wide, and they should be. Much of the exposure outside public equity is estimated rather than measured, and the estimation methodology is explicit rather than precise. A narrower-scope definition would produce a lower aggregate number. A wider-scope definition would produce a higher one. An investor performing a similar analysis for their own plan would need to choose a scope, apply it consistently, and accept that the result has meaningful ranges around it.

Figure 2

Asset Allocation vs AI Allocation for Illustrative Example

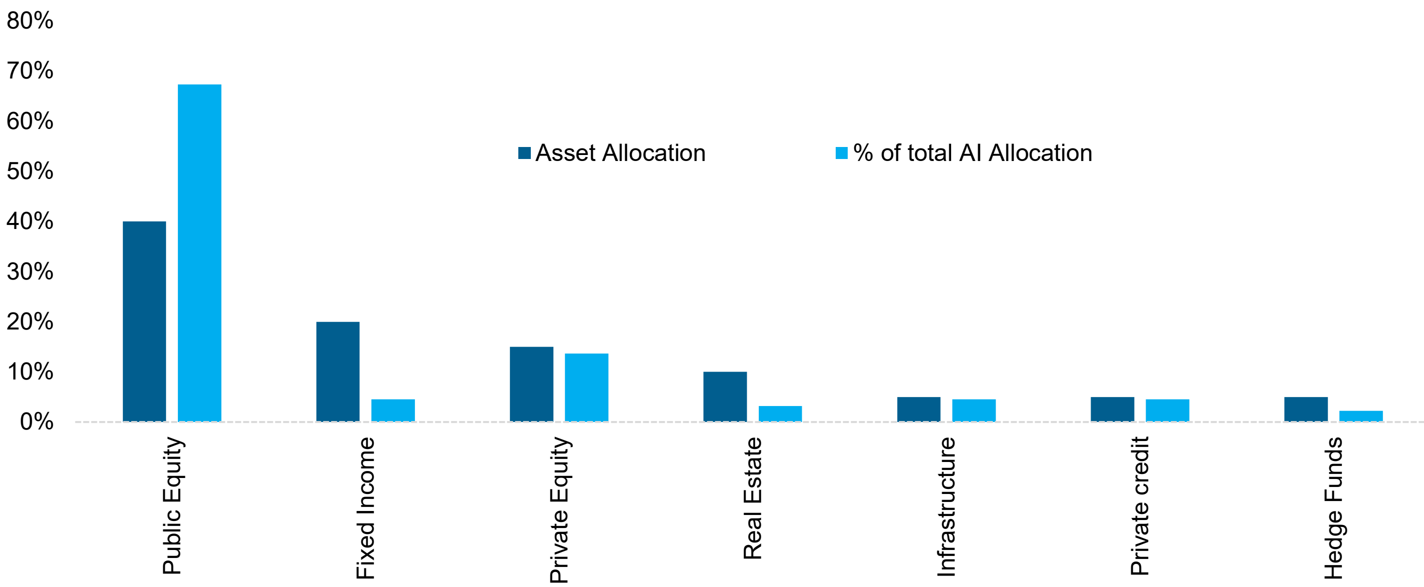

Source: Meketa Investment Group, 2026.

What the Exercise Shows

First, the gap between asset class-level and total-portfolio views is meaningful. A plan that considers AI exposure only through its public equity allocation sees only about two-thirds of what the whole-of-portfolio view shows. Most of the gap occurs in asset classes that are often evaluated separately from public equity in standard risk reporting.

Second, the private markets contribution is substantial but hard to pin down precisely. The real estate, infrastructure, private credit, and private equity allocations together contribute approximately 5-10% of total plan exposure under the stylized assumptions, which is less than the public equity contribution, but with ranges that reflect real uncertainty about the underlying positions. An investor who wants precision in this area will have to engage managers directly and build the exposure picture from the bottom up.

Third, the counterparty concentration point is visible only at the total-portfolio level. The Magnificent Seven names that drive the public equity contribution are also the counterparties to the data centers in the real estate portfolio, much of the infrastructure exposure through offtake agreements, and a meaningful share of the private credit exposure through special purpose vehicles. A total-portfolio view reveals that these positions represent a concentrated exposure to a small number of ultimate counterparties.

Investor Implications

The measurement exercise produces a picture. The question that follows is what, if anything, an investor should do with that picture. The answer depends on the investor.

For All Investors

Most institutions have already made their fundamental equity risk decision through policy asset allocation, and they rely on other asset classes to provide ballast in drawdowns. When a thematic driver is concentrated enough to affect multiple asset classes simultaneously, that ballast may not perform as expected. Thematic exposure across the portfolio should therefore be evaluated within the total portfolio's risk budget and rebalancing policies, not addressed piecemeal within individual asset classes. The whole-of-portfolio view of AI exposure supports that evaluation.

For Investors With Large Private Markets Programs

The gap between public and private market AI exposure may be larger than expected, and the direction of the gap depends on vintage and strategy mix. Mature buyout-heavy private equity portfolios tend to be underweight AI exposure on the direct-participation side relative to public equity, while potentially carrying meaningful exposure to legacy software holdings. Recent infrastructure and real estate allocations tend to show heavier AI-adjacent exposure than headline asset class labels suggest, particularly when data center and power generation allocations are accumulated across multiple managers. Likewise, recent vintage years for venture capital have been heavily allocated to AI.

For these investors, the measurement exercise is most valuable and the most challenging. The public equity allocation is the most straightforward to analyze. However, the private markets allocations require manager engagement, bottom-up review, and explicit estimation assumptions. Iterating over time (i.e., refining the estimates as more data become available and as manager reporting improves) represents a realistic path to the most usable picture.

For Investors Considering Changes to Asset Allocation

Nothing in this paper argues for raising or lowering AI exposure. The measurement exercise supports whatever decision the investor is inclined to make by giving them a clearer picture of the starting position. An investor considering increasing their AI exposure through a new thematic allocation should know what they already hold. An investor considering reducing their AI exposure through changes to public equity should understand what share of their total exposure lives outside the equity allocation and will not be affected by changes there.

Governance Questions for the CIO and the Board

A useful output of the whole-of-portfolio exercise is a set of questions that the institution's investment office and board should be prepared to discuss. We offer the following as a starting point:

- What scope are we using to define AI-linked exposure, and is it applied consistently across asset classes?
- How confident are the estimates for each asset class, and what would improve the measurement quality?
- How does our cumulative AI exposure fit within our existing risk budget?
- To what extent is our cumulative exposure concentrated in a small number of ultimate counterparties (e.g., the hyperscalers) across multiple asset classes?
- How can we practically stress test our exposure, and is our rebalancing and/or hedging framework adequate for those conditions?
- How often should we re-run this exercise, and what would trigger us to do so more frequently?

These questions do not have universal answers. They are the conversation the exercise is designed to support.

Conclusion

A portfolio's risk comes from the drivers its positions share, not from the label of any individual asset class. Artificial intelligence is the most prominent example of a thematic driver that has propagated across asset classes simultaneously, through distinct mechanisms, producing cumulative exposure that asset class-level reporting may understate. While evaluating AI exposure within a single asset class is not wrong, it is incomplete.

The measurement exercise this paper describes is not technically complex, but it requires discipline: a defined scope applied consistently across asset classes, an honest accounting of what is measured versus estimated, and a willingness to aggregate across the siloes that standard risk reporting maintains. The result is not a precise number. It is a range with known uncertainty, which is more useful than a precise number that conceals it.

The analysis presented here suggests that a typical large institutional portfolio may carry AI-linked exposure in the range of 18–26% of total assets, compared to a public equity headline figure of 12–18%. That gap is not trivial, and for many institutions it has accumulated without an explicit decision to put it there.

Nothing in this paper argues for raising or lowering AI exposure. An investor who measures their cumulative position and concludes it exceeds what they intended has a different problem than one who concludes the opposite. Both conclusions are possible, and both are legitimate starting points for a portfolio decision. Understanding AI-driven exposure is a prerequisite for making deliberate decisions about that allocation.

The governance obligation this exercise creates is to decide, consciously, whether the portfolio's AI exposure is deliberate. The rest is established portfolio governance: risk budgeting, rebalancing discipline, and the patient acceptance of risk that comes with earning long-horizon returns.

End Notes

1 Source: Nuveen's sixth annual *Equilibrium Global Institutional Investor Survey*, preview results released February 4, 2026. The survey reflects responses from 800 institutional investors across 30 countries representing approximately $17 trillion in assets under management.

2 The hyperscalers are the small group of large cloud computing providers, primarily Amazon Web Services, Microsoft Azure, Google Cloud, Meta, and Oracle, that operate global networks of massive data centers and provide the underlying infrastructure on which most large-scale AI training and inference workloads run.

3 See Goldman Sachs Global Investment Research, which has articulated a four-phase AI value chain framework distinguishing creators, enablers, monetizers, and beneficiaries across multiple research publications beginning in 2024. The extension to a fifth category of potentially disrupted businesses (traditional seat-based software firms facing AI-driven pricing pressure) is not part of the original Goldman framework; it is introduced here for illustrative purposes to capture disruption risk that a purely opportunity-focused taxonomy might miss.

4 Note that this framework may already be dated, emphasizing the point that a framework for vetting AI must continue to evolve.

5 Source: Iñaki Aldasoro, Sebastian Doerr, and Daniel Rees, "Financing the AI Boom: From Cash Flows to Debt," BIS Bulletin No. 120, Bank for International Settlements, January 7, 2026. The BIS analysis defines AI firms as Alphabet, Amazon, Meta, Microsoft, and Oracle.

6 See, for example, "Tracking AI Innovation with an AI-Driven Indexing Approach", Abbie Zhang and Jason Ye, S&P Dow Jones Indices, September 2024.

7 Source: Bye, Per and Soerlie Kvaerner, Jens and Werker, Bas J.M., "Magnificent, but Not Extraordinary: Market Concentration in the US and Beyond" (January 16, 2026), covering January 1926 through December 2024, updated with FactSet data through 2025, as reported in "Market Concentration and the Case for Deliberate Exposure," Meketa Investment Group, February 2026.

8 Source: MSCI Index Constituents data, MSCI USA Index, April 30, 2026. Tight, moderate, and wide scope definitions applied at the GICS sub-industry level using market-capitalization weights. Tight scope includes Semiconductors, Semiconductor Materials & Equipment, Systems Software, Application Software (AI-direct subset), Interactive Media & Services, and the cloud-services share of Broadline Retail. Moderate scope adds Technology Hardware Storage & Peripherals, Communications Equipment, IT Services, Electrical Components & Equipment, Heavy Electrical Equipment, the data-center share of Equity REITs, and the data-center-load share of Independent Power Producers and Multi-Utilities. Wide scope adds Industrial Machinery (automation), additional Health Care Equipment & Services and Pharmaceuticals (AI-enabled incumbents), and remaining Tech Hardware. The MSCI USA Index covers 537 securities representing approximately 85% of the US equity universe by market capitalization and is broadly comparable to the S&P 500 in coverage of US large- and mid-cap equities.

9 Source: MSCI Index Constituents data, MSCI ACWI Index (2,514 constituents), April 30, 2026. Tight, moderate, and wide scope definitions applied at the GICS sub-industry level using market-capitalization weights. The MSCI ACWI Index covers approximately 85% of the global investable equity universe and is a widely-used benchmark for globally diversified large- and mid-cap equity exposure. Note that different weightings and the use of active management will likely lead to differences relative to the ACWI index.

End Notes

[10] The Bank for International Settlements documents that hyperscaler corporate bond gross issuance topped $100 billion in 2025, with most issuance at maturities over five years to lock in funding for multi-year buildouts. Source: Egemen Eren, Ingomar Krohn and Karamfil Todorov, "BIS Quarterly Review: Markets recalibrate amid shifting currents," March 16, 2026.

[11] Hyperscalers issued approximately $75 billion of US investment grade debt in September and October 2025 alone, more than double the sector's 2015-2024 average annual issuance. Source: Bank of America research, cited in Brandywine Global, "Brave New World of AI Capex: Giving Credit Where Credit Is Due," November 12, 2025.

[12] Source: Egemen Eren, Ingomar Krohn and Karamfil Todorov, "BIS Quarterly Review: Markets recalibrate amid shifting currents," March 16, 2026.

[13] The investment grade corporate bond market composed ~24% of the Bloomberg Aggregate index as of April 30, 2026. The technology sector composed ~8.5% of the corporate bond market, while other AI-related names (Amazon, Meta, various utilities and REITs) add up to another ~5%. Sources: Bloomberg, ICE BofA Corporate Index.

[14] Source: Crunchbase, "6 Charts That Show the Big AI Funding Trends of 2025," December 16, 2025.

[15] These ranges reflect Meketa's analysis of a subset of client portfolios as of December 2025 and may not be representative of any specific investor's portfolio. Actual exposure will vary based on portfolio construction, manager selection, and vintage.

[16] Source: Aldasoro, Doerr, and Rees, BIS Bulletin No. 120.

[17] Source: Aldasoro, Doerr, and Rees, BIS Bulletin No. 120; S&P Global Market Intelligence; Capital IQ. Given the sources, the data is likely reflective of publicly disclosed private credit vehicles (e.g., business development companies and registered interval funds), where portfolio-level data is available. The data may differ significantly for institutional private credit portfolios. And as with private equity, portfolio construction decisions will also alter exposure.

[18] Source: Octus, "Octus Private Credit Software Analysis Reveals Almost 30% Exposure to BDCs; 13% of Loans Have PIK Component; Fewer Than 10% Loans Mature Before 2028", February 26, 2026. Octus found that BDC investment in software companies was approximately 29% as of September 30, 2025. They note that this figure includes any company whose primary service is delivered via software.

[19] Source: JP Morgan, Deutsche Bank, Bank of America, Barclays, and Bloomberg, as compiled in Legal & General Investment Management, "Securitized Credit Financing Data Centers," February 2026.

[20] These ranges reflect Meketa's analysis of a subset of client portfolios as of December 2025 and may not be representative of any specific investor's portfolio. Actual exposure will vary based on portfolio construction, manager selection, and vintage.

[21] Source: NAREIT, "Data Centers Rebound in Q4 Actively Managed US Real Estate Fund Tracker," February 11, 2026. Data as of the fourth quarter of 2025.

[22] Source: Preqin data on private infrastructure transactions, as reported in industry press coverage through 2025.

[23] Source: Meketa analysis of NCREIF ODCE data as of June 30, 2025. Data centers are included (along with other assets) in the "other" category, which composed 1.7% of the index.

[24] These ranges reflect Meketa's analysis of a subset of client portfolios as of December 2025 and may not be representative of any specific investor's portfolio. Actual exposure will vary based on portfolio construction, manager selection, and vintage.

[25] Source: International Energy Agency, "Energy and AI", April 2025. Base Case projections are used throughout.

[26] These ranges reflect Meketa's analysis of a subset of client portfolios as of December 2025 and may not be representative of any specific investor's portfolio. Actual exposure will vary based on portfolio construction, manager selection, and vintage.

[27] An empirical illustration is available from MSCI's Thematic Exposure scores, which assign each security a continuous score from 0 to 1 across themes including disruptive technology, autonomous technology, and next-generation internet. Applied to the top 50 weighted constituents of the MSCI USA Index as of April 30, 2026, the category-based approach (which counts each AI-classified security at its full weight) produces a tight-scope figure of approximately 34%, while a simple exposure-based approach that weights each security by the maximum of its disruptive technology, autonomous technology, and next-generation internet scores produces approximately 40% across the same top 50 names. The two approaches will not always diverge in the same direction, and the gap depends on which thematic scores are used, but the example illustrates that the choice of methodology has measurement consequences worth being explicit about. Note that MSCI thematic scores are not designed as AI-specific metrics, and these three themes are the closest proxies for AI exposure in their published thematic taxonomy.

[28] Modern Portfolio Theory (MPT), developed by Harry Markowitz, serves as a traditional framework for combining asset classes to maximize expected return for a given level of risk based on assumed returns, volatilities, and correlations.

Important Information

This report (the "report") has been prepared for the sole benefit of the intended recipient (the "recipient").

Significant events may occur (or have occurred) after the date of this report, and it is not our function or responsibility to update this report. The information contained herein, including any opinions or recommendations, represents our good faith views as of the date of this report and is subject to change at any time. All investments involve risk, and there can be no guarantee that the strategies, tactics, and methods discussed here will be successful.

The information used to prepare this report may have been obtained from investment managers, custodians, and other external sources. Some of this report may have been produced with the assistance of artificial intelligence ("AI") technology. While we have exercised reasonable care in preparing this report, we cannot guarantee the accuracy, adequacy, validity, reliability, availability, or completeness of any information contained herein, whether obtained externally or produced by the AI.

The recipient should be aware that this report may include AI-generated content that may not have considered all risk factors. The recipient is advised to consult with their Meketa advisor or another professional advisor before making any financial decisions or taking any action based on the content of this report. We believe the information to be factual and up-to-date but do not assume any responsibility for errors or omissions in the content produced. Under no circumstances shall we be liable for any special, direct, indirect, consequential, or incidental damages or any damages whatsoever, whether in an action of contract, negligence, or other tort, arising out of or in connection with the use of this content. It is important for the recipient to critically evaluate the information provided.

Certain information contained in this report may constitute "forward-looking statements," which can be identified by the use of terminology such as "may," "will," "should," "expect," "aim," "anticipate," "target," "project," "estimate," "intend," "continue," or "believe," or the negatives thereof or other variations thereon or comparable terminology. Any forward-looking statements, forecasts, projections, valuations, or results in this report are based upon current assumptions. Changes to any assumptions may have a material impact on forward-looking statements, forecasts, projections, valuations, or results. Actual results may therefore be materially different from any forecasts, projections, valuations, or results in this report.

Performance data contained herein represent past performance. Past performance is no guarantee of future results.