

Signal over Noise #14: The AI "super-exponential" meets its speed limits

Global equity strategy

Authors: Ulrike Hoffmann-Burchardi, Chief Investment Officer Americas and Global Head of Equities, UBS Financial Services Inc. (UBS FS); Nikolaos Fostieris, Global Equity Strategist, UBS Switzerland AG; James Dobson, CFA, CIO Equity Strategist, US Energy & Utilities, UBS Financial Services Inc. (UBS FS)

- AI capabilities appear to be improving at an accelerating pace, supporting our "super-exponential" framing. While models tackle increasingly sophisticated tasks, investors should avoid pricing a similarly rapid rise in earnings into valuations, as technical progress does not automatically translate into commercial returns.
- Value capture may accumulate along a more linear path, as power supply and infrastructure cannot readily expand at exponential rates, while enterprises need time to adapt their operations. In our view, these bottlenecks create selective buying opportunities within the broader AI investment theme, particularly in essential inputs whose supply is slow to expand and difficult to replace.
- Narrowing this gap will likely require a combination of more compute devoted to inference, broader on- and off-grid power supply, and targeted permitting and grid reforms. Alongside more efficient hardware and adapted business processes, these measures could accelerate commercialization.
- Our investment focus remains on: 1) the bottlenecks across the AI and power and resources value chains, through suppliers of scarce inputs and providers that help ease these constraints; 2) use cases in the application layer where we see the highest potential return on investment for enterprises; 3) areas such as pharma, biotechnology, and health care services where we believe long-term benefits are not yet fully appreciated.



AI capabilities seem to have no ceiling


The capabilities of AI models continue to improve at rates faster than what simple exponential terms would imply. Earlier this year, we introduced the notion of "AI super-exponential" to describe this trajectory. Although used descriptively rather than as a strict mathematical claim,

the term describes an exponential fit whose growth-rate coefficient increases as new observations are incorporated, implying a progressively shorter doubling time. The pattern appears in the METR metric, where new model releases are increasingly placed above the exponential trendline.

This report has been prepared by UBS Financial Services Inc. (UBS FS), UBS Switzerland AG. **Please see important disclaimers and disclosures at the end of the document.**

Figure 1 - Model capabilities are improving at an accelerating pace

The longest human task (by duration) that an LLM can complete with at least 50% success rate



Source: Model Evaluation and Threat Research, UBS, as of Apr 2026
Note: Further model updates are excluded as measurements above 16 hours are not reliable based on the current methodology

The latest development with state-of-the-art models resolving aspects of the Navier-Stokes problem is not just pointing toward this direction, but rather highlighting that intelligence on specific aspects has reached and is expanding the frontier of human capabilities. When we hosted NVIDIA’s CEO Jensen Huang, back in 2024, and discussed the topic of breakthroughs at the level of Millennium Prize Problems, he noted: "If you told me that within five years, some of these areas are going to see some big breakthroughs, I would not be surprised."

At the time, his answer sounded far-fetched and overly optimistic. The fact that we reached levels of intelligence able to tackle some of the world’s hardest unsolved problems at speeds that not even the most optimistic voices could predict suggests that even after four years into this cycle, the pace of progress may still remain underappreciated.

Whether labeled AGI, AI singularity, or just another indicator toward increasing AI capabilities, it strongly signals that bottlenecks in the value chain increasingly lie outside the intelligence layer. In our view, the question of AI value capture likely becomes dependent on two major bottlenecks: first, the availability of sufficient physical infrastructure, including data center space and power and compute capacity; and second, enterprises’ ability to resolve organizational barriers and integrate AI effectively into their workflows.

The binding constraint is the input hardest to replace and slowest to expand

The expansion of AI computing capacity increasingly depends on how quickly capital can be converted into operational infrastructure. This requires sites, permits, grid connections, power generation, electrical equipment, cooling, chips, and networking, ultimately delivering servers that can be used economically. Delays across these

interdependent stages can leave capital deployed but unable to earn a return.

The key analytical point is that the binding constraint is not simply the item with the longest lead time, but rather the input on the critical path with both the longest time to resolve and the lowest substitutability. Long lead times can be managed where purchasers can switch supplier, technology, or geography. By contrast, a constraint that cannot easily be switched around can hold up every subsequent layer of the value chain while supporting pricing and margins.

Table 1 - Compute: The bottleneck is moving upstream

Segment	Substitutability	Pre-AI benchmark	Current lead time
Chip manufacturing			
Semicap equipment	Very low	3-10 months	1-2 years
Foundry	Very low	2-3 years	3-4 years
Advanced packaging / substrates	Low	3-6 months	1-1.5 years
Semiconductors			
GPUs / AI accelerators	Low	5-6 months	6-12 months
Memory (HBM / DRAM)	Very low	3-6 months	6-12 months
CPUs	Medium	ca. 3 months	3-6 months
IT infrastructure			
Optical transceivers / Networking switches	Low	1-3 months	3-9 months
AI servers	Medium	ca. 1 month	1-2 months
Storage systems	Medium	1-3 months	1-2 years

Source: Company reports, UBS, as of Sep 2026
Note: Figures are indicative and based on company-reported project lead times

The semiconductor cycle has already illustrated this dynamic. Early in the AI buildout, the visible shortage was accelerator compute. As hyperscaler demand expanded, constraints also migrated upstream into high-bandwidth memory (HBM), advanced packaging, and leading-edge foundry capacity. These pressures are now extending further into semicap equipment and process-control capacity.

That does not mean that foundries and semicap tools are scarce at all times. Rather, the economics of AI have moved upstream. This also helps separate scarcity from high demand: When accelerator demand exceeds the ability of HBM, packaging, foundry capacity, and equipment to expand, value accrues to the inputs that enable marginal supply. Accelerators can remain a large and profitable market even as alternatives proliferate. In our view, however, the more durable scarcity premiums may currently lie in select upstream inputs where switching suppliers, technologies, or production routes is particularly difficult. Those premiums can weaken as capacity expands or alternative technologies become viable.

Power is becoming the physical equivalent of the upstream semiconductor constraint

The same logic applies to the power stack. The first phase of the buildout centered on accelerator availability and data center shells. Constraints are increasingly extending

upstream through the power chain: access to land with usable interconnection, transmission capacity, substations, high-voltage transformers, switchgear, generation, fuel supply, and skilled labor.

Table 2 - Physical infrastructure: The current scarcity map


Segment	Substitutability	Pre-AI benchmark	Current lead time
Data center operators			
Permitting / site selection	Medium	n/a	6–18 months
Power supply			
Utility power availability / Grid interconnection	Low	2-3 years	3-7 years (5 year median)
Behind-the-meter power (fuel cells, reciprocating engines, gas turbines for on-site generation)	Medium	n/a	1-3 years
Gas turbines	Low	1-2 years	3- 5 years
Battery storage (BESS)	Medium	3-6 months	6–12 months
Industrial equipment			
Substations / Transformers	Low	6-12 months	2-4 years
Back-up generators	Medium	6-9 months	1-2 years
Electrical distribution & cooling (UPS / PDU / busway / cooling)	Low-Medium	1-3 months	6–18 months

Source: Company reports, UBS, as of Sep 2026
Note: Figures are indicative and based on company-reported project lead times

This is why power should not be treated as a single bottleneck. Aggregate generation capacity does not guarantee dependable power at the required site, on a specific timetable, and under a financeable commercial structure. Interconnection data illustrate this challenge—median lead times from connection requests to commercial operation increased by more than 60% for projects built in 2015-2025 compared to 2005-2014.

Figure 2 - Interconnection queues have lengthened considerably in recent years

Median duration to move from interconnection request to commercial operations date, in years



Source: Lawrence Berkeley National Laboratory, UBS, as of Sep 2026

Similar to the hardware stack, the most persistent scarcity premiums arise where capacity is both slow to expand and difficult to substitute: grid access, transmission, and select large-scale equipment, particularly gas turbines for projects relying on new gas-fired generation.

That was explicitly stated by Microsoft’s CEO in late 2025, noting that “If you can’t do that [i.e., secure access to

power], you may actually have a bunch of chips sitting in inventory that I can’t plug in. In fact, that is my problem today. It’s not a supply issue of chips. It’s actually the fact that I don’t have warm shells to plug into.”


Utilities are advancing capacity plans based on available demand visibility. However, it is worth highlighting that power remains a highly capital-intensive segment. AI data center power needs can range from hundreds of megawatts for large campuses to gigawatt-level clusters for the largest planned developments. As a result, supplying demand at this scale requires substantial infrastructure investment, firm contracts, or durable customer commitments.

Delays in adding capacity are often amplified by hyperscalers’ approach, which creates a vicious cycle: Longer connection queues prompt them to pursue multiple locations in parallel, potentially adding to queue congestion and making underlying demand harder to assess. While negotiations over risks and shared liabilities are complex, utilities cannot commit capacity to every prospective project.

Even after contracts are signed, new plants take years to build, since physical infrastructure cannot match the rapid scaling investors have grown accustomed to in software in recent years. Meanwhile, utilities are pursuing other ways to ease supply constraints while new capacity is developed. These include expanding battery storage, both on and off grid, to smooth demand peaks and improve utilization of installed generation capacity. Where new capacity is expected to lag demand, delayed retirements and operating-life extensions are also becoming increasingly important to preserve existing generation.

Figure 3 - Expanding power-generation capacity requires new installations and longer operating lives for existing plants


US added total power capacity and plant retirements, as a % of installed capacity



Source: US Energy Information Administration, UBS, as of Sep 2026

Figure 4 - Battery storage has been growing at a 65% annual rate over the last decade

Battery energy storage system installed capacity in the US, in GW



Source: Energy Institute - Statistical Review of World Energy, UBS, as of Sep 2026

Permitting is part of time-to-power


Political acceptance is becoming an additional constraint on converting capital into usable compute. In the US, federal support for AI leadership remains strong, particularly efforts to maintain a technological advantage over China. However, federal support does not automatically translate into permission to build at a particular site. Washington can facilitate federal reviews and support energy infrastructure, but state regulators and local authorities retain substantial influence over electricity pricing, cost allocation, zoning, water use, and project approvals.

Without a uniform federal framework governing data center development, state and local governments have pursued moratoriums that temporarily pause construction or expansion. The underlying tension comes from the public’s perception that many costs are local and immediately visible, while benefits are broader and less tangible, prompting political opposition. Developers also share part of the responsibility, as the push to move quickly can leave community engagement lagging and limit communication about local benefits and how project risks will be managed.

As a result, the Bipartisan Policy Center’s review recorded 54 local moratoriums enacted nationwide and proposals for state-level moratoriums in 15 states. Similarly, Data Center Watch identified 75 projects representing approximately USD 130bn of planned investment blocked or delayed by local opposition, almost matching the combined total for 3Q and 4Q25.

Figure 5 - Opposition to data center buildout is growing in the US

Opinions on the construction of a data center in the local area in the US, in %



Source: Heatmap News, JP Morgan, UBS, as of Sep 2026

Headline investment figures and moratorium counts do not equate to lost near-term capacity. Also, SemiAnalysis’s project-level assessment estimates that local moratoriums directly delay only around 7.6% of capacity within restricted areas. Nevertheless, while they do not establish a brake on nationwide deployment, political acceptance remains an important site-selection and execution risk.

Generation projects can face multiyear interconnection queues, while site selection and permitting can take 12-18 months. These represent different development stages, not directly comparable end-to-end schedules. Where equipment delivery or grid works already determine the completion date, a permitting pause may be absorbed without delaying energization, provided these activities can proceed in parallel. Where approvals determine when subsequent work can begin, however, permitting can extend already lengthy deployment timelines.

Workarounds shift rather than remove the critical path

With time-to-power evolving into a key commercial variable, limited access to powered sites is prompting developers to use bridge power and behind-the-meter generation to reduce the risk of idle compute while awaiting grid connections.

Developers are deploying natural-gas fuel cells, reciprocating gas engines, and aeroderivative turbines, including repurposed aircraft engines, to bypass grid delays. Certain fuel-cell installations can be deployed within months, subject to equipment availability, permitting, and site readiness, while new engine and turbine packages can require several years. Some of these technologies can provide permanent power rather than simply temporary bridges, but they introduce two main constraints.

First, bypassing the electricity-connection queue shifts the critical path toward equipment delivery, fuel infrastructure, permitting, system integration, and commissioning. These

dependencies can limit deployment even where a grid connection is no longer the immediate bottleneck.

Second, variable AI power demand and the financial consequences of delays or downtime raise the importance of reliable delivery and operation. Developers need providers with proven execution, strong balance sheets, integrated engineering and construction capabilities, and dependable maintenance support, along with systems capable of managing rapid load changes. Yet very few behind-the-meter providers meet all these requirements, contributing to complex contracts and lengthy negotiations.

These requirements do not make on-site generation inherently unsuitable for permanent operation, but they do introduce additional limiting factors for a broad-based adoption of behind-the-meter solutions.

Organizational inertia is the second major bottleneck

Another significant limiting factor for value capture sits inside large enterprises. Adoption is expanding steadily, but evidence of broad, measurable economic benefits remains uneven. The latest user survey from frontier large language model usage points to significant differences across occupations: Usage is skewed heavily toward computer and mathematical roles and managers who are disproportionately represented relative to their share of US employment, although the latter appear to use AI beyond their core management tasks.

Software development has been a prominent use case from the outset, reflecting the models' strengths in coding and the affinity of the domain to the technology itself. Across other functions, institutional restrictions, established workflows, employee habits, and budget considerations can slow deployment. Better models alone do not remove these barriers; rather, enterprises need to adapt processes, integrate data, and change how people work.

How does the “super-exponential” intersect with the linear

Putting it all together, our “super-exponential” framing describes the acceleration observed in select AI capability measures, while similarly, physical and organizational constraints imply a more linear growth in value capture. The central tension is that if these growth trajectories persist, the dichotomy is likely to sustain, widening the gap between technical potential and realized economic value.

We see three possible responses, which could develop in parallel.

1. Compute allocation could shift toward commercial deployment

The first sits in the AI value chain.

Where powered compute remains scarce, AI labs may allocate more incremental capacity to revenue-generating

inference. This could moderate training expansion relative to an unconstrained scenario, potentially slowing some capability gains while supporting near-term monetization.

Such a shift could better serve enterprises that can already capture substantial value without requiring the solution of Millennium problems. At the same time, it could also encourage longer useful lives for the existing compute fleet and shift priorities of next-generation chips to improvements in energy efficiency and cost.

The trade-off is that a sustained reduction in training investment could temporarily constrain advances at the AI frontier. That could raise strategic concerns for the US, given the administration's emphasis on AI leadership, particularly relative to China. At the same time, advancing model capabilities still remains a central objective and source of differentiation for AI labs. Competitive incentives may therefore make coordinated restraint difficult, as each lab risks falling behind rivals that continue investing in performance development.

Finally, an excessive focus on near-term applications could also delay research in areas where potential breakthroughs could deliver larger long-term benefits, such as drug discovery, nuclear fusion, and novel industrial materials.

2. Power constraints could accelerate alternative supply solutions

The second sits in the power and resources value chain.

In the absence of grid access, cloud providers may need to increase their reliance on behind-the-meter generation and off-grid solutions. Scaling these alternatives would require expanded manufacturing capacity and credible project delivery. While the latter is particularly important for the behind-the-meter market, which is largely served by smaller companies, limited access to timely grid connections could nevertheless prompt project developers to diversify their power sources in the interim, despite existing execution challenges.

3. Policy could address select deployment bottlenecks

Third, strategic policy intervention could directly target the resolution, at least partially, of the main bottlenecks. Potential measures include coordinated permitting for critical projects, interconnection reforms, screening out speculative or duplicate connection requests, and accelerated investment in grid expansion and modernization. Meanwhile, clearer rules for allocating large-load infrastructure costs and permitting on-site generation could also improve project certainty.

It is unlikely that any single response will resolve the growth rate mismatch. However, a combination of them could narrow the gap and help infrastructure deployment keep pace with AI development, which along with more energy-efficient hardware and better load management could

further improve the use of available power.

For investors, the takeaway is twofold. First, we favor building exposure across the AI and power and resources value chains to critical inputs that are difficult to substitute and whose supply is slow to expand. Second, providers that help developers ease these constraints or could benefit from faster, targeted policy support could be well positioned to benefit.

Global asset class preferences definitions

The asset class preferences provide high-level guidance to make investment decisions. The preferences reflect the collective judgement of the members of the House View meeting, primarily based on assessments of expected total returns on liquid and commonly known indices, House View scenarios, and analyst convictions over the next 12 months.

Absolute view

Attractive: We consider this asset class to be attractive. Consider opportunities in this asset class.

Neutral: We do not expect outsized returns or losses. Hold longer-term exposure.

Unattractive: We consider this asset class to be unattractive. Consider alternative opportunities.

Note: For equities, we have a five-tier rating system with two additional preferences. When equities are included with the other asset classes in the three-tier rating system, we collapse "Most Attractive" with "Attractive" and "Least Attractive" with "Unattractive."

Most Attractive: We consider this asset class to be among the most attractive. Investors should seek opportunities to add exposure.

Least Attractive: We consider this asset class to be among the least attractive. Seek more favorable alternatives opportunities.

Relative preferences

For the multi asset view, we also include relative preferences to help investors compare between items.

Preferred: We view this item more favorably than others within its asset class.*

Neutral: We see this item similarly to others within its asset class.*

Least preferred: We view this item less favorably than others within its asset class.*

*In the case of cross-asset, compared to other asset classes.

Note: The tactical asset allocation (TAA) positioning of our different investment strategies may differ from these views due to factors including portfolio construction, concentration, and borrowing constraints.

Appendix

Risk information

UBS Chief Investment Office's ("CIO") investment views are prepared and published by the Global Wealth Management business of UBS Switzerland AG (regulated by FINMA in Switzerland) or its affiliates ("UBS"), part of UBS Group AG ("UBS Group"). UBS Group includes former Credit Suisse AG, its subsidiaries, branches and affiliates. The investment views have been prepared in accordance with legal requirements designed to promote the **independence of investment research**.

Generic investment research – Risk information:

This publication is **for your information only** and is not intended as an offer, or a solicitation of an offer, to buy or sell any investment or other specific product. The analysis contained herein does not constitute a personal recommendation or take into account the particular investment objectives, investment strategies, financial situation and needs of any specific recipient. It is based on numerous assumptions. Different assumptions could result in materially different results. Certain services and products are subject to legal restrictions and cannot be offered worldwide on an unrestricted basis and/or may not be eligible for sale to all investors. All information and opinions expressed in this document were obtained from sources believed to be reliable and in good faith, but no representation or warranty, express or implied, is made as to its accuracy or completeness (other than disclosures relating to UBS). All information and opinions as well as any forecasts, estimates and market prices indicated are current as of the date of this report, and are subject to change without notice. Opinions expressed herein may differ or be contrary to those expressed by other business areas or divisions of UBS as a result of using different assumptions and/or criteria. UBS may utilise artificial intelligence tools ("AI Tools") in the preparation of this document. Notwithstanding any such use of AI Tools, this document has undergone human review.

In no circumstances may this document or any of the information (including any forecast, value, index or other calculated amount ("Values")) be used for any of the following purposes (i) valuation or accounting purposes; (ii) to determine the amounts due or payable, the price or the value of any financial instrument or financial contract; or (iii) to measure the performance of any financial instrument including, without limitation, for the purpose of tracking the return or performance of any Value or of defining the asset allocation of portfolio or of computing performance fees. By receiving this document and the information you will be deemed to represent and warrant to UBS that you will not use this document or otherwise rely on any of the information for any of the above purposes. UBS and any of its directors or employees may be entitled at any time to hold long or short positions in investment instruments referred to herein, carry out transactions involving relevant investment instruments in the capacity of principal or agent, or provide any other services or have officers, who serve as directors, either to/for the issuer, the investment instrument itself or to/for any company commercially or financially affiliated to such issuers. At any time, investment decisions (including whether to buy, sell or hold securities) made by UBS and its employees may differ from or be contrary to the opinions expressed in UBS research publications. Some investments may not be readily realizable since the market in the securities is illiquid and therefore valuing the investment and identifying the risk to which you are exposed may be difficult to quantify. UBS relies on information barriers to control the flow of information

contained in one or more areas within UBS, into other areas, units, divisions or affiliates of UBS. Futures and options trading is not suitable for every investor as there is a substantial risk of loss, and losses in excess of an initial investment may occur. Past performance of an investment is no guarantee for its future performance. Additional information will be made available upon request. Some investments may be subject to sudden and large falls in value and on realization you may receive back less than you invested or may be required to pay more. Changes in foreign exchange rates may have an adverse effect on the price, value or income of an investment. The analyst(s) responsible for the preparation of this report may interact with trading desk personnel, sales personnel and other constituencies for the purpose of gathering, synthesizing and interpreting market information.

Different areas, groups, and personnel within UBS Group may produce and distribute separate research products **independently of each other**. For example, research publications from **CIO** are produced by UBS Global Wealth Management. **UBS Global Research** is produced by UBS Investment Bank. **Research methodologies and rating systems of each separate research organization may differ**, for example, in terms of investment recommendations, investment horizon, model assumptions, and valuation methods. As a consequence, except for certain economic forecasts (for which UBS CIO and UBS Global Research may collaborate), investment recommendations, ratings, price targets, and valuations provided by each of the separate research organizations may be different, or inconsistent. You should refer to each relevant research product for the details as to their methodologies and rating system. Not all clients may have access to all products from every organization. Each research product is subject to the policies and procedures of the organization that produces it. The compensation of the analyst(s) who prepared this report is determined exclusively by research management and senior management (not including investment banking). Analyst compensation is not based on investment banking, sales and trading or principal trading revenues, however, compensation may relate to the revenues of UBS Group as a whole, of which investment banking, sales and trading and principal trading are a part.

Tax treatment depends on the individual circumstances and may be subject to change in the future. UBS does not provide legal or tax advice and makes no representations as to the tax treatment of assets or the investment returns thereon both in general or with reference to specific client's circumstances and needs. We are of necessity unable to take into account the particular investment objectives, financial situation and needs of our individual clients and we would recommend that you take financial and/or tax advice as to the implications (including tax) of investing in any of the products mentioned herein. This material may not be reproduced or copies circulated without prior authority of UBS. Unless otherwise agreed in writing UBS expressly prohibits the distribution and transfer of this material to third parties for any reason. UBS accepts no liability whatsoever for any claims or lawsuits from any third parties arising from the use or distribution of this material. This report is for distribution only under such circumstances as may be permitted by applicable law. For information on the ways in which CIO manages conflicts and maintains independence of its investment views and publication offering, and research and rating methodologies, please visit www.ubs.com/research-methodology. Additional information on the relevant authors of this publication and other CIO publication(s) referenced in this report; and copies of any past reports on this topic; are available upon request from your client advisor.

Important Information About Sustainable Investing Strategies: This terminology refers to definitions in the UBS Sustainable Investing Framework and does not refer or relate to any product specific regulatory labelling regime or naming conventions. Sustainable investing strategies aim to consider and incorporate environmental, social and governance (ESG) factors into investment process and portfolio construction. Strategies across geographies approach ESG analysis and incorporate the findings in a variety of ways. Incorporating ESG factors or sustainable investing considerations may inhibit UBS's ability to participate in or to advise on certain investment opportunities that otherwise would be consistent with the Client's investment objectives. The returns on a portfolio incorporating ESG factors or sustainable investing considerations may be lower or higher than portfolios where ESG factors, exclusions, or other sustainability issues are not considered by UBS, and the investment opportunities available to such portfolios may differ.

External Asset Managers / External Financial Consultants: In case this research or publication is provided to an External Asset Manager or an External Financial Consultant, UBS expressly prohibits that it is redistributed by the External Asset Manager or the External Financial Consultant and is made available to their clients and/or third parties.

USA: Distributed to US persons only by UBS Financial Services Inc. or UBS Securities LLC, subsidiaries of UBS AG. UBS Switzerland AG, UBS Europe SE, UBS Bank, S.A., UBS Brasil Administradora de Valores Mobiliarios Ltda, UBS (Brasil) Corretora de Valores S.A., UBS Asesores Mexico, S.A. de C.V., UBS SuMi TRUST Wealth Management Co., Ltd., UBS Wealth Management Israel Ltd and UBS Menkul Degerler AS are affiliates of UBS AG. **UBS Financial Services Inc. accepts responsibility for the content of a report prepared by a non-US affiliate when it distributes reports to US persons. All transactions by a US person in the securities mentioned in this report should be effected through a US-registered broker dealer affiliated with UBS, and not through a non-US affiliate. The contents of this report have not been and will not be approved by any securities or investment authority in the United States or elsewhere. UBS Financial Services Inc. is not acting as a municipal advisor to any municipal entity or obligated person within the meaning of Section 15B of the Securities Exchange Act (the "Municipal Advisor Rule") and the opinions or views contained herein are not intended to be, and do not constitute, advice within the meaning of the Municipal Advisor Rule.**

For country information, please visit ubs.com/cio-country-disclaimer-gr or ask your client advisor for the full disclaimer.

Except as otherwise specified herein and/or depending on the local entity from which you are receiving this report, this report is distributed by UBS Switzerland AG, authorised and regulated by the Swiss Financial Market Supervisory Authority (FINMA).
Version C/2026. CIO82652744
© UBS 2026. The key symbol and UBS are among the registered and unregistered trademarks of UBS. All rights reserved.