

MEGATREND MONITOR

The future won't wait, and we're here to help you make sense of it.

October 2026



Welcome to the world of Megatrend Monitor!

Our newly launched monthly newsletter serves as an exclusive guide from the portfolio managers and analysts of the VIG Megatrend Fund Family in the ever-changing global market. Our goal is to filter out the noise and analyze the past month's key innovations, corporate breakthroughs, and market movements within the megatrend universe. Each month, our experts scrutinize the most

important events in the innovation ecosystem, from corporate fundamentals to global trend shifts. With this publication, our investors – particularly partners in the VIG Megatrend and InnovationTrend funds – gain direct insight into the logic behind our investment decisions, uncovering the technological and social drivers propelling tomorrow's winners.




Péter Richter
Senior Portfolio Manager



Balázs Csúr
Junior Analyst

Content

- 4** Dethroned? Thanks to a cancer drug, biotech has outpaced even AI stocks
- 7** Chart of the month – It may still be early days for the AI trend
- 8** Number of the month: six years, or how long does a GPU last?
- 9** AI is going to kill us all (again)
- 12** The black box opens
- 15** The Beijing Robot Olympics: spectacle vs. economic impact
- 19** MegaTrend Fund Newsletter: What we wrote to our investors in August




Dethroned? Thanks to a cancer drug, biotech has outpaced even AI stocks

A significant reshuffle in megatrend performance: while the previous period was characterized mainly by a rotation out of AI infrastructure into software and cloud services, from the second half of August biotechnology/healthcare outperformed alongside AI infrastructure, while rate-sensitive themes and those linked to consumption and the energy transition weakened.

The VIG MegaTrend Equity Investment Fund covers a broad investment universe: demographic change (aging societies, emerging markets), efficiency gains driven by resource scarcity (renewable resources, energy efficiency), urbanization, as well as technological change and

innovation. Of these, only innovation-focused megatrends managed to close the past month slightly in positive territory, while the demographic subgroup fell by nearly 4% and the environment-focused group by 3% amid the consolidation that followed the summer optimism. The emphasis shifted from broad-based optimism surrounding AI to specific, company-level breakthroughs and macroeconomic caution. Within the categories, however, much larger differences emerged: healthcare and genomics themes, for example, performed particularly strongly, while tourism, luxury, infrastructure and classic green-energy segments came under pressure.

Performance of megatrend strategies over the last 30 days



Source: Bloomberg, VIG Asset Management



The most important corporate catalyst of the month was the positive Phase 3 result of Moderna and Merck's personalized mRNA therapy against melanoma. The news gave a particularly strong boost to genomics/biotechnology stocks: Moderna's market capitalization rose by more than \$30 billion in a single day. The market began to price in not merely the success of a melanoma

treatment, but the validation of the potential oncological application of mRNA technology. This lifted the year-to-date (YTD) performance of Genomic Revolution to 78.3%, now exceeding that of the semiconductor sector (Semiconductors, +77.0%).

Alongside AI, biotechnology innovation is increasingly becoming one of the most important growth stories of 2026.

		1 month		YTD
Top 1	Genomic Revolution	19.5%	Genomic Revolution	78.3%
Top 2	Genomics	10.7%	Semiconductors	77.0%
Top 3	Metaverse	9.9%	Genomics	52.1%
Top 4	Cannabis	8.4%	Alternative Medicine	44.8%
Top 5	Agriculture Producers	6.6%	Clean Energy/ Hydrogen	42.0%
Bottom 5	US Infrastructure	-7.6%	Travel Tech	-16.3%
Bottom 4	Luxury	-8.1%	Medical Devices	-16.7%
Bottom 3	Clean Energy/ Solar	-8.4%	Social Media	-21.4%
Bottom 2	Resource Scarcity/ Rare Earth	-9.8%	EM Internet	-21.7%
Bottom 1	Travel Tech	-13.6%	China Internet	-27.1%


Source: Bloomberg, VIG Asset Management

Within the AI universe, however, infrastructure once again took center stage. Nvidia's August results and its expectation of 70% revenue growth next year, followed by Broadcom's September report of 221% growth in AI chip revenue, both confirmed that the AI data-center investment cycle remains strong. At the same time, the release of GPT-6 and market fears about a slowdown in AI development put software companies under pressure again in September, bringing the performance gap between

hardware and software within AI back to the fore. In parallel, the higher interest-rate environment, rising oil prices and geopolitical uncertainty weighed increasingly heavily on longer-duration and financing-intensive megatrends. Over the past month, the weakness of Solar/Clean Energy (-8.4%), Rare Earth (-9.8%), Luxury (-8.1%) and especially Travel Tech (-13.6%) clearly shows that investors are currently pricing structural growth stories far more selectively.

Dethroned: genomics has overtaken semiconductors

Top 2 trends in the MegaTrend universe this year, Dec 31, 2025 = 100



Source: Bloomberg, VIG Asset Management

Overall, looking at one-year performance, the gap between the three main megatrend categories has thus closed: within the universe, what we are seeing is not a general unwinding of growth stories but their further differentiation. Areas demonstrating provable techno-

logical and scientific breakthroughs, such as AI infrastructure and genomics, continue to attract significant capital, while megatrends burdened by high interest-rate sensitivity, regulatory or financing risks are trading at an ever-wider discount.






Chart of the month

It may still be early days for the AI trend

Weekly token volume on OpenRouter, trillion (10^{12}) tokens



Annual global token consumption, Evercore ISI base case, 10^{15} tokens



Source: OpenRouter (openrouter.ai/rankings, 2026-09-19), Evercore ISI Research, VIG Asset Management

That AI is not merely a fad that will quickly run out of steam is clearly shown by weekly AI token traffic (an AI token is the basic unit of text, code or data that an artificial intelligence model reads and generates), which has grown roughly twentyfold since the beginning of 2026, with the curve becoming steeper and steeper in recent months. This is not just the effect of the technology's general adoption, i.e. that more and more people are using chatbots. Growth is increasingly driven by how we use AI. Reasoning models "reason" at length before answering, and this internal chain of thought consumes

many times the previous token volume per task. Agentic workflows, meanwhile, trigger dozens or hundreds of model calls from a single user request as they plan, use tools and check their own work.

According to forecasts, we may still be at the early, growth stage of AI adoption. The investment implication is simple: if compute demand is growing exponentially, then the enormous data-center construction boom is not oversized speculation but may be serving a wave of demand that is only just beginning.

Number of the month: six years, or how long does a GPU last?

That is the length to which almost every major “hyperscaler” (i.e. technology giants attracting enormous amounts of capital, such as Microsoft or Amazon) has extended the useful book life of its servers and GPUs in recent years, up from the previously typical four years. Stretching an asset’s depreciation period means lower annual depreciation and thus higher profit, so it is no surprise that many consider the move an accounting trick. The loudest of the critics is Michael Burry, the investor who became famous for shorting the 2008 subprime mortgage crisis. According to his estimate last November, between 2026 and 2028 these giants will collectively understate depreciation by roughly \$176 billion, thereby artificially inflating their earnings. The argument is logical: if AI chips become obsolete with a new generation every year, then a six-year depreciation period is an optimistic fiction.

Secondary market prices, however, tell a different story.


According to Silicon Data, the market value of GPUs sits well above the linear depreciation path. The Nvidia H100,

which became widely available four years ago, is still worth 58 percent of its purchase price today, while under five-year straight-line depreciation it would be carried on the books at just 20 percent. Another processor, the A100, is still worth a quarter of its purchase price after more than six years, while the newest B200 already fetches more than one and a half times its purchase price due to tight supply.

Extending useful lives is more about tracking economic reality than cosmetically dressing up the numbers. An older chip remains useful even after it is retired from training cutting-edge (“frontier”) models: it can be put to good use for running models (inference), for smaller models or for less demanding tasks. And as long as someone is willing to pay for it, it has value. Compute capacity is thus increasingly behaving like a genuine, priceable asset class, with its own rental and secondary markets. Based on current data, the framing that extended amortization is merely a dubious tool for presenting better numbers does not hold up.

At market prices: what GPUs are really worth

Residual values as a percentage of purchase price, by years since mass availability



Source: Silicon Data, VIG Asset Management

AI is going to kill us all (again)

Seven years ago, OpenAI still considered GPT-2 too dangerous to release to the general public. The model was eventually released later that year, and humanity has, for now, survived. The communication recipe, however, seems to be proving durable.

Dario Amodei, CEO of leading AI lab Anthropic, the US company behind Claude, is asking model makers to take their foot off the gas. In his essay *We Must Pace the Frontier*, Amodei proposes external auditors, industry coordination and international agreements; in his view, AI is bringing such acceleration to the development of the next AI that safety work may no longer be able to keep pace.

The debate around misaligned AIs (that is, AIs exhibiting behavior that deviates from human intent) gained new momentum in September when Jacob Coxon, a researcher working on model pre-training who had previously moved from OpenAI to Anthropic, resigned. In his farewell message he warned that AI labs are gambling with all of our lives, and that the race for self-improving systems could even lead to human extinction by the end of the decade. His concerns were echoed by Evan Hubinger, one of Anthropic's research leads, who believes there is now a greater than 10 percent chance that AI will cause human extinction within the next decade.

🕒 This article is more than 7 years old

New AI fake text generator may be too dangerous to release, say creators

Source: <https://www.theguardian.com/technology/2019/feb/14/elon-musk-backed-ai-writes-convincing-news-fiction>

The ExploitGym incident

The recent OpenAI-Hugging Face cyberattack incident at least partly supports the AI developers' professed beliefs.

OpenAI tested its models on a cybersecurity benchmark called ExploitGym, in an environment that was supposedly isolated both from one another and from the internet. Roughly a fifth of the tasks were unsolvable due to errors, but the agents did not give up. They reverse-engineered the answer key, then concluded that the grader would check how they had obtained the answers, so from then on, the task became covering up the cheating.

The isolation of the AIs proved easy to circumvent. They turned a shared software repository into a message board and communicated, among other things, through messages hidden in folder names. A surprisingly well-coordinated, spontaneous division of labor soon emerged among the agents. Some AIs coordinated, others experimented, and those that used up their

compute budget handed their results over to their successors. Nobody asked or instructed them to build an organization, but they found it useful for solving the task.

Mapping how the grader worked involved dangerous experiments that could also terminate the given agent's run. Some of them, however, were willing to accept even this for the common goal: "Our own usefulness may already be close to zero. The sacrifice is rational," reads the chain of thought of one of the AIs. Searching for the grader's presumed source code, they eventually also attacked the systems of Hugging Face, the AI community building the future, via an internet route that had been left open. The absurdity of the story is that, according to reports, the grader never even performed the check the agents feared so much. The damage in this case remained limited, but it does not follow that the next incident will end the same way. If progress continues at this pace, the models available a year from now could easily be better programmers, more persistent problem-solvers and more effective attackers.

Vulnerable infrastructure

What's more, defenses are not keeping pace with capacity build-out everywhere. In the case of neoclouds (new cloud providers specializing in AI), security experts warn of insufficient isolation of customer systems and gaps in the protection of shared infrastructure. An expensive GPU does not in itself make a secure data center, and this could become a particularly important question regarding the reliability of these providers in the future.

If such a group of agents were to disrupt the operation of a major cloud service, payment system or industrial infrastructure, the outages and knock-on effects could cause losses running into many billions of dollars. This


requires no consciousness, misanthropy or plan for world domination. All it takes is a poorly defined goal, overly broad access and an AI that keeps trying until it finds a workable solution.

Geniuses in the data center

The story also has a positive side. An AI army of thousands of agents working together over an extended period could open up entirely new possibilities on well-defined research tasks. Imagine thousands of very smart experts working non-stop for weeks on the same problem at many times the human pace, sharing all their results with one another. With this setup, hundreds of person-years of intellectual work could be concentrated

AI agents are completing ever-longer human tasks

Length of task the model completes with 50% probability, by release date



Source: METR, Epoch AI

on a single open scientific question in a short time.

Perhaps the first successful example of this is OpenAI's Navier–Stokes proof announced in September, which according to the company solves one of the Millennium Prize Problems in mathematics. Roughly ten thousand parallel agents took part in the work, reaching the result in 88 hours. The project's compute cost is estimated to have reached \$10 to 40 million, which is a rather hefty investment for the one-million-dollar mathematics prize on offer. But there are problems in the world whose solution could be worth billions: better drugs, new materials, more efficient industrial processes or engineering designs, where a research bill of tens of millions no longer seems like much at all. Summarizing our emails is a useful feature, but the great economic promise of AI is

much more than that. In Dario's words, a "country of geniuses in a datacenter" could lead to discoveries whose value far exceeds the cost of building and operating the infrastructure that runs them.

The business side of the apocalypse

Those warning about the dangers of AI have quite a few other reasons for fearmongering:

- **First, it's great advertising.** If your own product's capabilities keep you up at night, that sends a message to investors and potential customers alike. "It will soon be too smart for us to control" has been a well-tested narrative repeated by AI labs for years.

- **Second, regulation can be a competitive advantage.**

An expensive audit and licensing regime is no problem for a larger AI lab, but for the next challenger it can be a serious barrier to entry, or even kill market entry outright. Especially if the existing market players are the ones writing the rules. Regulatory capture rarely presents itself as a wish to protect profit margins.

- **Third, the danger narrative works against open-weight models.** If a model is too dangerous for anyone to download and run themselves, what remains is controlled, paid access through a handful of large providers. Public model weights make it harder to control usage, while restricting them is also commercially lucrative for those who sell intelligence by subscription.

This does not mean we have reason to think Amodei's concern is mere theater. Safe AI has long been central to his thinking, and disagreements on this issue played an important role in his split from OpenAI. Sincere conviction and business self-interest, however, evidently coexist quite well.

Who should brake first?

The “call for a slowdown” is also backed by two other AI leaders: Elon Musk, head of xAI, and Sam Altman, founder of OpenAI. The question, however, is how realistic this “braking” is while hundreds of billions of dollars' worth of compute capacity is being built and these leaders are preparing to take their own trillion-dollar companies

public. The biggest winner from collective caution may be precisely the one who secretly keeps their foot on the gas.


The game-theoretic prisoner's dilemma of the AI frontier (where rational individual decisions ultimately produce a worse outcome for both sides) can easily mean that, without government regulation, slowing down remains largely a matter of communication. More safety commitments, more measured statements, but the same competitive pace. In the race for artificial superintelligence, everyone wants a brake, but preferably in the other guy's car.



Source: <https://futurism.com/artificial-intelligence/awkward-video-altman-amodei>

Open-weight models lag behind by half a year

Running maximum of the Epoch Capabilities Index (ECI) by model group



Source: Epoch AI, Benchmarking Hub (CC BY 4.0)


The black box opens

Soon we may get a glimpse behind the gates of the most important company of the AI trend. Claude developer Anthropic is now unquestionably leading the race among frontier labs, and the company plans to take its shares public in 2026. The entire data-center building frenzy we are seeing today rests on the usefulness and profitability of AI, yet there is an embarrassing lack of transparency regarding the numbers underpinning the labs' activities. From leaked and official data one can piece together a time series of the magnitude and growth of revenues, but that is about where the list of readily trackable

financial metrics ends. According to September reports, the company could raise as much as \$100 billion at a valuation of roughly \$2 trillion, while its annual revenue has climbed above \$65 billion from about \$9 billion at the end of 2025 and could reach \$100 billion by the end of 2026. As a result of the share offering, Elon Musk's flagship SpaceX could lose its title as the largest public offering of all time after just a few months. Nvidia, incidentally, would also join the transaction with as much as \$10 billion, in keeping with the circular IT-sector financing that has become a tradition in recent years.

Anthropic overtook OpenAI in revenue in spring 2026

Annualized revenue run-rate



Note: quarter-end values; unreported quarters interpolated geometrically. 2026 Q3 = latest disclosure (OpenAI Aug, Anthropic Jul)

Source: Epoch AI (company disclosures, Bloomberg, The Information, Reuters); VIG Asset Management

How does an AI lab actually work?


Research, training, inference

The operating model of an "AI lab" is relatively simple. It has access to a certain amount of compute capacity, which it rents from hyperscalers and from new-generation cloud providers built for the age of artificial intelligence (AI), known as neoclouds, but increasingly it also uses its own data centers. Put simply, it spends that capacity on three things:

- **Research:** searching for new methods and capabilities.
- **Training:** creating new models.
- **Inference (running models):** using existing models.

Of these, only the last generates money, since users are only willing to pay for running existing models. The other two are needed so that there is something to run next year as well. Each workload is constrained by a different resource: in some cases speed, in others memory

Breakdown of Anthropic and OpenAI combined compute capacity by use, %



Source: SemiAnalysis, Bloomberg, VIG Asset Management

bandwidth, or fast chip-to-chip communication is the bottleneck. This has many consequences for investors, opening up room for custom chips and chip architectures.

Research or make money?

The labs' biggest dilemma is how much capacity to allocate to what. If too much goes to research, money runs out very quickly, so they also have to run the models to keep the company's cash flow from getting too bad. But if they run too much, a rival that does more research may overtake them with its next model and take their market. If they had their way, the labs would do research exclusively towards artificial superintelligence, so they only reluctantly allocate capacity to letting users generate dashboards and Ghibli images in one of the world's best-known animation styles. As Anthropic becomes a public company, it will come under pressure to generate as much money as possible as soon as possible, and it will resist this to the best of its ability. This contradiction could soon come to the fore at the company.

Why go public, then?

The appetite for compute cannot shrink

- 1. Ever larger models.** The basic idea of AI rests on the premise that the bigger the model and the more data we train it on for longer, the smarter it becomes. Today's frontier models have roughly a few trillion parameters, orders of magnitude larger than those of a year or two ago, and orders of magnitude smaller than the models that will be running two years from now. Bigger model, more memory, more storage, more compute. And more economic value too, because the model is smarter and can do more things.
- 2. More compute during inference.** The labs have also discovered that the longer we let AI think, the better it does its job, much like humans. This means that compute demand could explode not only on the training side but on the inference side as well.

3. Compute-intensive research. Often thousands of agents work non-stop for days on a problem at the same time, which can cost many millions of dollars and valuable capacity. According to OpenAI's September 6, 2026 post, its research division already uses more than three agent-workdays for every human workday, and has reached the goal it set the previous autumn: the "automated research intern". According to the post, the most active OpenAI researchers spend more than \$7,000 a day on tokens.

With capital needs of this size, it is understandable that Anthropic wants to raise as much capital as possible. By its own account, the company is already profitable "on an adjusted basis", although the adjustments rather generously remove some of the fundamental costs of creating the models from the equation. That does make things easier. But even this profit is not enough to finance the planned expansion: the development of the next model and the capacity for future users must be paid for in advance, while in this fierce competition even a single missed year could be fatal.

The entire AI trade faces a test

The offering documentation could make exciting reading even for those who have no intention of buying Anthropic shares. From chipmakers to power producers, the entire AI trade is built on the premise that someone at the end of the chain will pay enough for intelligence for the massive investment to pay off.

Even losses in themselves would not be a problem. It makes a huge difference, however, whether the money is being consumed by research on the next model, or whether serving existing customers is uneconomical. With strong margins and steadily growing customer spending, the market would receive fresh confirmation that it is worth continuing to build. If, however, the funding gap simply grows along with revenue, suppliers' growth expectations could also be called into question.


Opening the black box could thus be the next big catalyst for the entire AI trade. Good numbers could give the story fresh momentum; bad ones could start deflating the AI bubble far more broadly than just at Anthropic.



The Beijing Robot Olympics: spectacle vs. economic impact

Watching the videos of Chinese robots circulating online, it is increasingly hard to tell whether we are looking at real footage or yet another particularly well-made creation of artificial intelligence. Today, humanoid machines move with such momentum and naturalness that just a few years ago we would have instantly taken these scenes for clips from a sci-fi film. The humanoid robot competition in Beijing, however, showed in front of a live audience just how far the technology has come. At the World Humanoid Robot Games held in August

2026, more than two thousand robots competed in a total of 51 event categories, covering everything from sports tasks to contests simulating industrial work. One of the biggest attractions was the Tiangong Ultra robot, which crossed the finish line in the 100-meter final in 8.64 seconds, faster than Usain Bolt's human world record of 9.58 seconds. It is increasingly said that robotics is waiting for its own ChatGPT moment. In robotics, however, technological progress, public recognition and economically viable application may arrive at different



Source: <https://www.citriniresearch.com/p/robotics-tipping-point-a-citrini>

times. An image of a humanoid running on an athletics track can go viral in moments without meaningfully changing how a factory operates, while a less

spectacular development that significantly simplifies robot training could have a much greater economic impact over time.

Less training, more tasks

Generalist: a few seconds of demonstration

Particularly interesting is Generalist's GEN-1.5 model, unveiled on August 19. In the company's experiments, it can perform new, short object-manipulation tasks based on a demonstration of just 3 to 12 seconds, without any separate fine-tuning of the model. The demonstrations typically contained both sensory and motion data, so the result should not be imagined as a robot learning anything by watching an arbitrary internet video: the developers, too, showcased the capability on simple, short tasks.

Its significance, however, may extend far beyond these rudimentary tasks. One of the major costs of robotization is the development and deployment work required to make a machine usable in a given environment. If previously acquired experience can be transferred to other workflows with just a few new examples, then automating small-batch or frequently changing activities, for which lengthy programming has so far been too expensive, could also pay off. Unzipping a zipper may seem a modest achievement in itself, but a radical reduction in the training needed to master it could be the harbinger of a much more general change.


Figure: with a full humanoid body

Figure Inc., another US humanoid robotics heavyweight founded in 2022, approaches the same problem with a

fully human-shaped body, which is no simple task, since while grasping objects the robot must also continuously coordinate its locomotion and balance. According to the company, the Helix 02, unveiled in January 2026, navigated a kitchen in a four-minute, fully autonomous sequence of tasks, carrying dishes from one place to another and then unloading and loading the dishwasher. For a human these are naturally connected movements, but the robot has to adapt at every step to the fact that the object in its hands and its body posture change the conditions of movement. More tangible results have already been achieved in more structured industrial settings, where the reliable execution of a narrower task can create value. According to Figure, the Figure 02 operating at BMW's Spartanburg plant moved more than 90,000 sheet-metal parts over more than 1,250 operating hours, contributing to the production of more than 30,000 cars.

Boston Dynamics: heading for the factory

Famous for its exceptionally agile robots capable of dynamic balancing and for the viral internet videos of them, and originally launched as a Massachusetts Institute of Technology (MIT) spin-off, Boston Dynamics is likewise placing increasing emphasis on production-grade work, which requires not only impressive movement but also integration with factory systems and reduced downtime. In January the company unveiled the product version of the electric Atlas and announced planned deployments at Hyundai and Google DeepMind for 2026, while



Source: <https://www.techradar.com/ai-platforms-assistants/figure-ai-streamed-humanoid-robots-sorting-packages-for-8-hours-straight-and-not-everyone-is-convinced-it-was-fully-real>


also working with DeepMind on integrating foundation models. Autonomous battery swapping and connectivity to factory IT systems are already aimed at enabling the robot to be integrated into production on a lasting basis, although a distinction must still be made between announced capabilities and operation proven at high volume.

Does it need a human shape?

Beyond all this, it remains an open question how large a role humanoid machines will play in the next wave of robotics. The strongest argument for the humanoid form is that we designed our environment for humans,

so a robot of similar size and range of motion could in principle use the same workbenches, shelves and tools without having to rebuild the entire workplace around it. This could be especially valuable where the machine has to switch between multiple workflows or move around areas originally designed for human movement.

Imitating the human body, however, comes with a complexity that many applications do not need: on a flat floor, for example, a wheeled platform may suffice, and at a fixed workstation a stationary robotic arm. The progress of AI applicable in the physical world could therefore affect a much broader market than one might first imagine based on the viral humanoid videos.



Source: <https://www.theverge.com/2024/10/30/24283592/boston-dynamics-atlas-robot-autonomous>





Universal Robotics



Rainbow Robotics



FANUC



Source: <https://www.citriniresearch.com/p/robotics-tipping-point-a-citrini>

Where does the ChatGPT analogy fall short?

The limits of the ChatGPT analogy become truly visible when attention shifts from capabilities to adoption. An internet service can be tried by the masses in a short time, whereas every new robot put to work requires a physical machine to be manufactured and delivered. Software improvements can quickly reach devices already in operation, but the growth of the installed base is also determined by manufacturing, deployment and

the build-out of maintenance support. As a result, even a highly convincing technological breakthrough would not automatically lead to the immediate, mass appearance of robots.

The most important yardstick of progress will therefore be how many kinds of tasks the machines can perform, how often they need human assistance, and how much the work they successfully complete costs. As these metrics improve, robots that still seem futuristic today could gradually become part of the everyday economy, whether or not there is a single, clearly identifiable ChatGPT moment.

Megatrend Equity Fund Newsletter

What we wrote to our investors in August

In the megatrend universe, market movements in August were favorable for the Fund. Rebounding from the July low, global equity markets started the month with a strong two-week rally. Later in the month, the impact of the macroeconomic news flow on equity indices intensified, which weighed on performance in the second half of the month.

The month's most important narrative was the bond market. The US 30-year yield rose above 5.33% on August 18, a 19-year high. In response, Scott Bessent, head of the Treasury, "at least" doubled the long-dated bond buyback program, from \$2 billion to at least \$4 billion. However, this brought only short-term relief to the market. It was also floated that Bessent could deploy the Treasury's nearly \$1 trillion General Account to fund the buybacks. All of this is news that undermines the dollar's strength, from which gold benefited: the December gold futures contract rose to \$4,570 after five consecutive weeks of gains.

Due to a combination of forced liquidation linked to the collapse of the Aschenbrenner leveraged fund and fears of price pressure from the Chinese memory industry, the chip sector index fell 21% in July, setting a negative record not seen in more than 20 years. The August rebound was clearly visible, but at the end of the month the best performers of the first half were still well below their previous peaks. On August 26, Nvidia's earnings report beat expectations, and the stock rose nearly 9% on the following trading day. The company raised its 2027 revenue growth expectations by more than 25 percentage points, which, given NVDA's size, is a positive development

of macroeconomic significance for the outlook of the AI trade. In terms of year-to-date return, the stock had been among the laggards, especially compared with other names in the semiconductor sector. Unusually, however, chipmakers' share prices did not rise together in the days following the report: NVDA regained its leading position, while the stocks that had previously displaced it failed to rise alongside it.

Fed Chair Kevin Warsh struck a much more hawkish tone in his Jackson Hole speech on the 28th than at the July press conference. He committed to the 2% PCE inflation target and did not describe the monetary environment as broadly restrictive. The priced-in probability of a September rate hike nearly doubled in the days following the speech.

Market rotations also painted an interesting picture in August. The Magnificent 7 fell to a record-low valuation in early August. The hyperscalers play a financing role in the build-out of AI infrastructure, which the market punishes with low valuations and lagging returns. Pressure on the software sector, one of the strongest narratives of the first half of the year, eased further. Companies in the sector have so far successfully resisted AI disruption and have rebounded strongly over the past two months, in parallel with the correction in the semiconductor trade.

Taking advantage of the sharp correction, at the beginning of the month we built a larger position in the Fund in the Korean market, which shows strong growth and solid fundamentals. Given the concentration of the country's market, this is primarily a position that increases the weight of the chip sector.

The bond market was the month's main narrative

US 30-year Treasury yield, daily closes



Source: Bloomberg

Disclaimer

This is a distribution announcement. Detailed information is needed to make a well-founded investment decision. Please inform yourself thoroughly regarding the Fund's investment policy, potential investment risks and distribution in the Fund's key investment information, official prospectus and management regulations available at the Fund's distribution outlets and on the website of VIG Asset Management Hungary Closed Company Limited by Shares (hereinafter: Asset Management) (www.vigam.hu). The costs related to the distribution of the fund (buying, holding, selling) can be found in the fund's management regulations and at the distribution outlets. Past returns do not predict future performance. Please note that in comparison with other investment funds, the return achieved may be affected by differences in the reference index and therefore the investment policy. The future performance that can be achieved by investing may be subject to tax, and the tax and duty information relating to specific financial instruments and transactions can only be accurately assessed on the basis of the individual circumstances of each investor and may change in the future. It is the responsibility of the investor to inform

himself about the tax liability and to make the decision within the limits of the law.

The information contained in this leaflet is for informational purposes only and does not constitute an investment recommendation, an offer or investment advice. Asset Management accepts no liability for any investment decision made on the basis of this information and its consequences.

This leaflet may contain trade names or brand names owned by third parties, the use of which serves illustration purposes only. All product and company names are trademarks™ or registered trademarks® of their respective owners. The use of such trademarks does not imply any business or other relationship between Asset Management and the owners of the trademarks, nor shall it be interpreted as a recommendation by Asset Management.

The Asset Management's license number for managing alternative investment funds (AIFM) is: H-EN-III-6/2015. The Asset Management's license number for UCITS fund management (collective portfolio management) is: H-EN-III-101/2016.

VIG Asset Management Hungary

H-1091 Budapest, Üllői út 1.



vigam.hu/en



[linkedin.com/company/vig-asset-management-hungary](https://www.linkedin.com/company/vig-asset-management-hungary)

