M.M.WARBURG & CO
BANK



Economic Situation and Strategy

13. November 2026

AI Investments, Bubble, Frenzy or Opportunity?

We are at the beginning of 2028. Anthropic is the most valuable AI lab and the most valuable AI model developer in the world. Its models routinely review contracts, advise patients, enable breakthroughs in physics, write software and replace entire departments. For new data centers, Anthropic has taken on commitments of one trillion US dollars through 2045, with a capacity of ten gigawatts, even though it only brings in a fraction of that amount and its revenue growth is running behind plan. Fresh capital is still flowing, and nobody is asking yet who will ultimately foot the bill.

Then comes a second DeepSeek moment. A few years after 2025, a Chinese model developer releases another model that is rated as highly capable but appears to use computing power more efficiently than anyone thought possible.

At the same time, an engineer in Mexico notices that a supposedly fully leased data center is drawing only a third of its planned power consumption. While people at Anthropic nervously run the numbers on how many quarters the money will last, the financial industry is slowly realizing that financing data centers carries systemic risk. There are only a few tenants, most of them loss-making. The technology is short-lived, but the financing runs for decades. On top of that, the financing is often circular and opaque. Data centers are frequently held by special purpose vehicles, and the debt on the balance sheets of the major players is not fully visible.

So the first cracks appear, and nobody sees that the whole structure is starting to wobble. It reads like the blurb of an economic thriller, and it could become reality within a few quarters.

But we are not there yet, and there are good reasons to believe it will not come to that. Still, it is well worth taking a closer look at the AI ecosystem of language models, hyperscalers, startups and hardware makers. Economic growth in many countries is shaped directly and, in part, indirectly by investment in this sector. Many stock indices would look far less spectacular without a handful of companies that have profited enormously from this development. And anyone who lived through the bursting of the internet bubble in 2000 and the financial crisis of 2008 cannot help but keep looking for the next systemic crisis. In theory, at least, the AI hype has the potential to trigger one. So let us take a look behind the scenes of the AI hype.

The AI ecosystem can be pictured as a pyramid through which money flows from the top down. At the very bottom sit the chipmakers, such as Nvidia, AMD or Broadcom. Like the people who sold picks and shovels during the gold rush, they are profiting handsomely and relatively predictably right now. Above them are the so-called hyperscalers, namely Amazon, Microsoft, Google, Oracle and their peers. They buy chips by the hundreds of thousands, build data centers along with their power supply (mostly gas-fired power plants or even reactivated nuclear plants, since solar and wind cannot deliver electricity reliably around the clock), and rent out their computing power.

Their biggest customers are the developers of AI language models, such as OpenAI or Anthropic. They rent the computing capacity, train their models on it, and then sell "intelligence per token" or subscriptions. At the top of the pyramid sit the second-tier startups, such as providers of coding tools, customer service bots or legal assistants. They usually do not build their own AI model. Instead, they buy one through an interface and develop a customized product on top of it. These companies often grow fast but operate on thin margins, because a large share of their revenue goes straight to the model provider. It is not uncommon for them to lose money, since they sell tokens for less than they pay for them in order to win market share. Interchangeable providers with no proprietary data and no deep integration into customer processes are especially vulnerable. Anyone who also subsidizes their customers' usage depends heavily on further venture capital.

And there is a catch. Part of the money goes round in circles. Chipmakers and hyperscalers invest in AI model developers, which then spend that money on chips and cloud services. For the entire pyramid to hold up over time, companies and consumers would have to generate enormous revenue growth, and that has not happened yet. So let us look at the numbers more closely, starting with the chipmakers.

Companies like Nvidia have grown phenomenally in recent years because their revenues and profits have risen just as phenomenally. Share price setbacks, such as the current one caused by an affair worthy of a James Bond plot involving a Nvidia executive and a Chinese spy, which led to the resignation of Jensen Huang, are seen by many investors as a bargain and an opportunity. And perhaps they are right. However, it should not go unmentioned that the rapid rise in production figures is not directly linked to the rapid expansion of data centers. Instead, hyperscalers have recently bought millions of chips and stockpiled them to guard against possible shortages in the future. This very stockpiling is likely to cause growth to slow noticeably for a few quarters.

Perhaps even more critical is the fact that the hyperscalers' main customers are two companies, Anthropic and OpenAI, and these two in turn generate a large share of their revenue from very few individual customers. Given the immense spending and investment, one can hardly speak of granular risk diversification.

Moreover, it must be assumed that the extreme power demand of data centers cannot always be met. Many people do not realize how much additional electricity is at stake. To generate the power needed by the data centers expected worldwide by the end of 2028, you would need almost 100 nuclear reactors of the Isar 2 type. Isar 2 was one of the most powerful reactors in the world before it was shut down. Converted into large wind turbines, you would need roughly 15,000 units, assuming the wind blows steadily, day and night, at maximum strength. In reality, even 30,000 very large wind turbines plus enormous storage capacity would fall far short. Anyone who downplays the power problem has not grasped the scale of this challenge.

The biggest sticking point, however, may be the business model of the data centers themselves. Hyperscalers are, after all, ordinary companies that make investments in order to earn profits later. Ultimately, those investments have to pay for themselves, or it would have been better not to make them. Payback periods are among the first things a business student hears in an introductory course. Or perhaps one should say they used to be, 20 or 30 years ago, because it seems that these basic business principles no longer carry much weight. Otherwise it is hard to explain why such enormous investments, which are still met by very little revenue, are being greeted with a shrug and blind faith.

Here are the numbers behind that. Hyperscalers are spending roughly 700 to 800 billion US dollars in 2026 on data centers and chips. So far, there is no sign of a structural break in this trend. The latest figures from the US cloud provider CoreWeave, for example, look quite good. Next year, the figure will very likely exceed 700 to 800 billion US dollars, and most estimates put it at 1,000 billion US dollars or more. For comparison, Microsoft's AI revenue in 2026 is around 35 billion US dollars, mostly paid by OpenAI. Amazon is in roughly the same range, mainly paid by OpenAI and Anthropic.

It is nearly impossible that the hyperscalers' AI revenues in 2026 will exceed 200 billion US dollars. You do not need an MBA to see that this math can only work if revenues rise sharply in the coming years and investment levels fall. Alternatively, investment can keep rising, but then revenues would have to surge.

How demanding the math is can be shown with a calculation that fits on the back of an envelope. An investment

has paid off once the returns have recovered the capital employed. Returns are not revenues. They are the surplus left after power, staff, maintenance, sales and other running costs. Capital providers also expect a return on their money. Data centers add a complication: buildings and chips have very different lifespans. For AI chips, an economic useful life of four to six years is often assumed, while buildings typically last around 20 years. That does not mean older chips become unusable. They do, however, have to compete with ever more powerful successors. Unfortunately, much of the money is tied up in this shorter-lived technology.

Let us take the 700 billion US dollars in investment for 2026 and assume, for simplicity, that 60 percent goes into chips and servers with a five-year useful life, and 40 percent into buildings and infrastructure with a 20-year life. Without any return on capital, roughly 100 billion US dollars a year would have to flow back during the first five full years of operation for the math to work. Assuming a 10 percent cost of capital, the required annual surplus rises to around 144 billion US dollars to recover the capital and pay a return. After five years, the chip portion and part of the buildings would be paid off. The remaining building and infrastructure costs would then require around 33 billion US dollars a year for another 15 years.

Now the back of the envelope turns into a revenue calculation. If, after running costs, a rather ambitious and in fact unrealistic 60 percent of revenue remains as cash surplus, the facilities would need to generate about 240 billion US dollars in revenue per year during the first five operating years. At 40 percent, it would be about 360 billion US dollars. If another 1,000 billion US dollars is invested in 2027 under the same assumptions, the annual hurdle during the overlapping full operating years rises to around 350 billion US dollars in cash surplus, or 580 to 870 billion US dollars in revenue.

To be fair, these amounts cannot simply be set against currently estimated AI revenues to explain the gap with the proven financing shortfall. The total investment does not serve only the separately reported AI business. It can also support traditional cloud services or generate additional income from advertising and software. What matters is how much of the additional return can actually be attributed to these assets. Still, the calculation shows the scale of the challenge: a great deal of money must be earned over many years. How much of that is explicitly "AI revenue" is secondary. Whether it actually reaches the bank after ongoing costs is not.

That does not necessarily mean money has been thrown away. Artificial intelligence is the big theme of the coming years and very likely of the next decade. In the last 200 years, perhaps no other development has transformed humanity and, with it, the economy so profoundly. Explosive growth in AI-generated revenue at the hyperscalers is therefore not unthinkable. It is even likely. One thing is just as clear, though. If that revenue fails to materialize, the pyramid described above will collapse like a house of cards.

The AI providers that offer customized solutions, but mostly build on models from Anthropic or OpenAI, do not help the pyramid's stability either. Their actual technological capabilities are often fairly limited, and their business model only works if they win a large market share. Insolvencies are practically built in, and they add another revenue risk for the developers of the large AI models. Revenue growth is vital for companies like Anthropic. In 2025, it generated less than 5 billion US dollars in revenue while posting an operating loss of 8 billion US dollars.

All of this can only work if enormous amounts of fresh capital keep flowing into the AI business in the coming years. In 2026 alone, roughly 500 billion US dollars of new debt is likely to have flowed to AI companies, because internal cash flow at very many firms is nowhere near enough to finance them and keep them alive. It all works as long as belief in the long-term economic sense of these ventures holds, but that belief can evaporate quickly. Then there is the need for equity. Here too, nothing goes up forever. In September 2026, OpenAI wanted to raise 30 billion euros at a valuation of 1,500 billion US dollars. So far, the result has been a capital increase of 11 billion US dollars at a valuation of 1,250 billion US dollars.

If capital increases, and perhaps later IPOs, fail to deliver the desired financing volumes, debt financing becomes all the harder. This is especially true because the providers of the large AI language models have made enormous, binding commitments to build out computing capacity over the coming years. According to Anthropic's prospectus of November 2 this year, its firm minimum commitments alone come to around 300 billion US dol-

lars through 2030, out of total committed spending of almost 900 billion US dollars as of September 30, 2026. With cash of roughly 130 billion US dollars in mid-2026 and still fairly modest revenues, the air could get thin.

This applies to Anthropic to a certain extent, but perhaps even more so to the many companies building data centers that, realistically speaking, have little experience doing so. A great deal of equity and debt has flowed in on the assumption that very little can go wrong. Reality is often more complicated. In some cases, the experienced civil engineers needed to deliver such massive projects smoothly simply do not exist. Permitting procedures rarely run as smoothly as one would ideally hope, and the required grid connections sometimes arrive later than planned.

So what can go wrong? Quite a lot, to be honest. Anthropic and OpenAI have committed to spending many times their annual revenues. The hyperscalers are pouring hundreds of billions into chips that will be scrap within a few years. The second-tier AI startups are on life support twice over. They depend on the prices charged by model providers and on those providers' goodwill in not moving into their business themselves. All it takes is for revenues to grow more slowly than hoped. And because almost everyone is at once a customer, supplier and investor of everyone else, the stability of this ecosystem is lower than often assumed.

Still, there will be winners who should stay in business for years to come, above all the chip industry. In the gold rush of the Wild West, it was mainly the shovel sellers who made a fortune, and they did it handsomely. The same is happening again. Nvidia achieves a gross margin of around 75 percent on its data center chips, and payment is due on delivery. Whether the purchase ever pays off is the customer's problem. What makes life hard for buyers is exactly what makes the business model work for manufacturers: the short life span of the chips. Those who survive the shakeout have to keep upgrading. The real danger for manufacturers arises when they lend their customers the money to buy the chips or take stakes in them. That is what happened with Lucent and Nortel, and it will probably happen again.

Still, the world will not end. Some hyperscalers may run into temporary economic trouble. The developers of the large language models may have to come down to earth in the coming years. Investors may lose money. It only becomes critical if this leads to a systemic crisis, and there is currently no compelling reason to expect one. With every day that the buildout of data centers slows, the risk of a systemic crisis shrinks a little. Seen this way, and as paradoxical as it may sound, the Democrats' tendency toward tighter regulation of data center construction in the US may help take some of the drama out of the development. An orderly slowdown of new projects could limit overinvestment. An abrupt halt to investment or long delays to projects already financed, however, would be counterproductive.

Finally, a bit of comfort. For one thing, the entire AI ecosystem in the US counts as critical infrastructure, and the US will not allow a complete collapse. Just as the German state would step in if Rheinmetall were on the brink of collapse, the US government would be expected to step in if its technological leadership in AI were put at risk across the board by corporate missteps.

For another thing, there have always been temporary exaggerations and overheated hopes around technological developments, yet the technologies themselves endured. After the dotcom crash, plenty of unused fiber optic cable lay in the ground, and Netflix and YouTube later ran over it. This time, the chips, buildings, power lines and power plants will remain. So the payback math may still work out after all, quite certainly for the chipmakers and in the long run even for the economy and society as a whole. But not necessarily for every investor who was there from the start and did not get out in time, or could not get out in time.

Dr. Christian Jasperneite

Market data

Stock markets (in local currency)	As of 13.11.2026 08:10	Change versus 06.11.2026 -1 week	12.10.2026 -1 month	12.08.2026 -3 months	12.11.2025 -1 year	31.12.2025 YTD
S&P 500	7612	3,2%	-2,5%	-1,8%	11,1%	11,2%
Nasdaq	25366	3,2%	-6,9%	-4,6%	8,4%	9,1%
DAX	26772	3,7%	5,9%	1,7%	9,8%	9,3%
MDAX	30858	0,7%	2,5%	-4,6%	4,3%	0,8%
TecDAX	4109	2,6%	2,9%	1,0%	14,9%	13,4%
EuroStoxx 50	6399	3,7%	3,2%	-2,1%	10,6%	10,5%
Stoxx 50	5522	2,8%	3,4%	-0,1%	13,1%	12,3%
SMI (Swiss Market Index)	15149	2,4%	9,2%	4,8%	18,4%	14,2%
Nikkei 225	61142	6,5%	-11,4%	-9,5%	19,7%	21,5%
Brasilien BOVESPA	171324	1,7%	-18,1%	2,3%	8,7%	6,3%
Indien BSE 30	81411	0,5%	13,2%	4,4%	-3,6%	-4,5%
China CSI 300	4010	-1,1%	-6,5%	-14,5%	-13,7%	-13,4%
MSCI Welt	4870	2,8%	-1,9%	-2,8%	10,3%	9,9%
MSCI Emerging Markets	1555	3,5%	-8,4%	-7,5%	10,5%	10,7%
Bond markets						
Bund-Future	122,41	-71	145	-233	-690	-516
Bobl-Future	112,97	-71	37	-107	-518	-319
Schatz-Future	104,87	-30	2	-66	-216	-192
3 Monats Euribor	2,75	-4	12	25	71	73
3M Euribor Future, Dec 2026	2,89	4	-1	14	90	78
3 Monats $ Libor	4,32	13	7	45	37	65
10 year US Treasuries	5,25	39	1	56	118	111
10 year Bunds	3,32	9	-13	19	71	50
10 year JGB	2,87	18	-13	3	118	83
10 year Swiss Government	0,56	11	1	19	39	24
US Treas 10Y Performance	619,06	-1,9%	1,5%	-1,9%	-3,3%	-3,2%
Bund 10Y Performance	554,05	-0,5%	1,4%	-0,6%	-2,0%	-0,7%
REX Performance Index	458,08	-0,3%	0,9%	0,0%	-0,5%	-0,2%
IBOXX AA, €	4,12	5	-7	36	99	79
IBOXX BBB, €	4,48	4	-6	52	108	92
Commodities						
MG Base Metal Index	553,69	-0,7%	-4,7%	-5,2%	17,6%	7,3%
Crude oil Brent	96,44	-2,0%	-9,7%	8,2%	53,6%	58,5%
Gold	3900,82	-4,0%	-6,5%	-11,9%	-6,8%	-9,8%
Silver	50,02	-3,6%	-17,1%	-23,4%	-6,4%	-29,8%
Aluminium	3137,20	-0,5%	2,7%	-5,4%	9,4%	5,7%
Copper	13642,11	-2,3%	-7,8%	-4,6%	24,8%	9,5%
Iron ore	91,60	-1,2%	0,8%	-3,7%	-12,1%	-14,5%
Freight rates SPCI	3633	0,0%	-0,8%	10,9%	143,0%	119,4%
Currencies						
EUR/ USD	1,1069	-0,9%	-1,0%	-4,1%	-4,4%	-5,8%
EUR/ GBP	0,8358	0,2%	-1,1%	-2,2%	-5,3%	-4,3%
EUR/ JPY	176,87	0,1%	-0,3%	-3,7%	-1,3%	-3,9%
EUR/ CHF	0,9239	1,0%	-0,8%	-1,4%	-0,1%	-0,8%
USD/ CNY	6,6700	-0,1%	-0,4%	-1,1%	-6,2%	-4,6%
USD/ JPY	159,86	0,9%	0,8%	0,3%	3,3%	2,0%
USD/ GBP	0,76	0,9%	-0,2%	2,0%	-0,8%	1,6%

Source: LSEG Datastream







Carsten Klude
+49 40 3282-2572
cklude@mmwarburg.com

Dr. Christian Jasperneite
+49 40 3282-2439
cjasperneite@mmwarburg.com

Dr. Rebekka Haller
+49 40 3282-2452
rhaller@mmwarburg.com

Jan Mooren
+49 40 3282-2132
jmooren@mmwarburg.com

Martin Hasse
+49 40 3282-2411
mhasse@mmwarburg.com

This information does not constitute an offer or an invitation to submit an offer, but is intended solely for orientation purposes and to present possible business activities. This information does not claim to be complete and is therefore non-binding. It does not constitute a recommendation for the independent purchase of financial instruments, but serves only as a suggestion for a possible asset allocation. The opinions expressed herein are subject to change at any time without prior notice. Insofar as statements are made about prices, interest rates or other indications, these refer exclusively to the time the information was prepared and do not contain any statement about future developments, in particular with regard to future gains or losses. Furthermore, this information does not constitute advice or a recommendation. Before concluding any transaction presented in this information, advice tailored to the customer and the product is required in every case. This information is confidential and intended exclusively for the addressee named herein. Any use beyond that by the addressee is not permitted without our consent. This applies in particular to reproductions, translations, microfilming, storage and processing in electronic media, as well as any other publication of the entire content or parts thereof.
This analysis is freely available on our website.