since 1977
PASI NEWS

Volume 50 - Number 2 Summer 2026

Professional Advisory Services, Inc.
2770 Indian River Blvd. – Suite 204 • Vero Beach, FL 32960
(800) 847-7274 • (772) 778-0552 • fax (772) 770-2979

Market Update – *A Quarter for the Record Books*

by David A. Jaffe, M.D.

The stock market just notched its best quarter in six years, with the Market-Cap Weighted (MW) S&P 500 gaining 15.20%. This dramatic recovery from the declines of the first quarter was fueled by strong corporate earnings, and optimism that U.S. – Iranian negotiations would stabilize global oil flows disrupted by the war. This second quarter surge left the MW S&P 500 up 10.21% for the year. The Equal-Weight (EW) S&P 500 added 11.39% for the quarter, standing up 12.13% year-to-date. Meanwhile, the PASI composite stock portfolio gained 15.61% in the quarter, and sits up 10.21% for 2026 YTD, matching the MW S&P 500 return. [1,2]

The most noteworthy change in the character of the market may be the broadening of the rally. While the Mega cap "Magnificent 7" have driven the index gains of the last three years, many of those stocks have languished in 2026. Shifting investor sentiment in the second quarter spurred advances in industrials, materials, utilities, financials, health care, and consumer staples. Small company stocks surged, as captured by the Russell 2000 index gain of 22.57%. Broadening of the stock market advance is seen as a healthy phenomenon, suggesting that investors are looking for opportunities outside of the sphere of businesses expected to benefit from the growing adoption of Artificial Intelligence (AI).

That's not to say that excitement over AI no longer dominates allocation of investor dollars. The semiconductor sector in particular, chip stocks and suppliers of semiconductor manufacturing equipment, saw meteoric price appreciation in the second quarter. One beneficiary, PASI portfolio stock Advanced Micro Devices (AMD), is highlighted by Jeremy Goldberg later in this newsletter. Performance of some of the other participants is captured in the chart included in my discussion of 2026 capital gains, found on page nine.

While the second quarter recovery was impressive, one must wonder about the fragility of the catalysts buoying investor confidence. As of this writing, shortly after the close of the quarter, the Memorandum of Understanding forged between the U.S. and Iran is crumbling, and

[1] All quoted returns include reinvested dividends.
[2] Please see performance disclosure on page 11.

commerce through the Strait of Hormuz is currently in doubt. Data center expansion, a cornerstone of forecast U.S. company profit growth, faces challenges as outlined by Nathan Polackwich in the following article. Further, new data center construction is grappling with a cascade of local opposition, residents worried over environmental and public health issues, and the cost of electricity. This is a non-partisan issue; the number of opposition groups has doubled in 2026 compared to 2025, with participation in 49 states. At least 75 large data center projects had been delayed or cancelled by the end of the first quarter this year. Last, the anticipated interest rate tailwind has turned into a headwind. A rate cut, widely expected early in 2026, now seems unlikely with rising energy costs, and the June consumer price index (CPI) reported at 3.5% year-over-year, moderating but still well above the Federal Reserve's 2% target.

Predicting the near-term direction of the financial markets is always fraught. Our strategy remains unchanged – time has proven that maintaining a well-diversified portfolio of high-quality stocks, balanced with a conservative mix of investment-grade bonds, will provide long-term growth and safety for our clients.

Where's the Gold?

by Nathan Polackwich, CFA

There's an old saying – often and likely apocryphally credited to Mark Twain – that advises, "in a gold rush, sell picks and shovels." The idea came from the observation that in the mid-19th century California gold rush, the largest fortunes were made not by miners but by those who supplied them with tools and equipment, clothes (the origin of Levi Strauss), alcohol, transportation, banking, and other goods and services.

Today's gold rush is Artificial Intelligence (AI), and once again it's the companies selling the "picks and shovels" that are thriving. The profitability of the semiconductor industry, in particular, has grown to almost incomprehensible levels. Memory chipmaker Micron Technology's net income is anticipated to be around $130 billion this calendar year. In 2022 it made $9.5 billion and in 2023 it actually lost $4.8 billion (the memory chip business has historically been extremely cyclical and low margin). Taiwan Semiconductor, the world's largest contract semiconductor manufacturer, is anticipated to earn roughly $80 billion in profit this year vs. $32 billion in 2022 and just $12 billion in 2019. And finally Nvidia, the biggest AI chip designer, is expected to make *$220 billion* in profit in 2026. Back in 2022, before the AI boom began, its annual profit was only $4 billion. PASI semiconductor and semiconductor equipment stocks like Advanced Micro Devices (AMD), Analog Devices (ADI), Lam Research (LRCX), and ASML are also benefiting mightily from this boom, with revenue and profits surging.

And it's not just the semiconductor industry that's profiting heavily off the AI gold rush. Companies involved in the building of the huge data centers needed to power AI – such as those in construction and electrical equipment, grid infrastructure, cooling systems, networking hardware, power management solutions, and industrial automation – are also experiencing explosive demand.

The thing about a gold rush, however, is that it requires some gold to actually be there! In California between 1848 and 1855, an estimated 750 tons was mined. At the time this was

equivalent to around 25% of U.S. Gross Domestic Product (GDP, which is the total dollar amount of economic activity in a year). Many prospectors and mining companies did indeed strike it rich, though most barely covered their expenses.

In the AI boom, the "gold" would be profitable AI-based businesses, but so far at least that's proven elusive. The "miners" in the AI gold rush – from the major generative AI model developers like OpenAI (maker of ChatGPT) and Anthropic (maker of Claude) to the AI application companies built on top of their models – are all profoundly unprofitable. Worse, their losses seem to be accelerating even as their revenue grows.

Traditional software has historically been an extremely high-margin business because once the code is written, the cost to distribute and run it is close to zero. AI models are different. Both training the models and running them are exorbitantly expensive, as they demand massive computing power and electricity and costs rise linearly with usage. Specialized AI data centers must be constructed costing tens of billions of dollars, and once demand exceeds capacity, more must be built.

While impressive, the capabilities of AI haven't been great enough (or useful enough for businesses) that model makers and AI application companies can charge prices that come anywhere close to offsetting the models' cost. The Financial Times recently reported, for instance, that in 2025 OpenAI lost $20 billion on just $14 billion in revenue. Revenue might double this year to perhaps $30 billion, but the company's operating margin remains deeply negative at a reported -122%, suggesting 2026 losses could surpass $30 billion! Anthropic's operating margin was just as poor as OpenAI's in 2025, and it's expected to incur billions more in losses this year. Keep in mind that both OpenAI and Anthropic are generating these prodigious losses despite receiving preferential pricing on computing power from the companies funding much of the industry's massive data center buildout – namely, Microsoft, Google, and Amazon. Incredibly, along with outside venture capitalists, these three companies have also been directly financing OpenAI and Anthropic through equity investments. In effect, it's questionable whether OpenAI and Anthropic would still be in business without the largesse of venture capitalists and these three tech giants.

In 2026 Microsoft, Google, and Amazon as well as other data center builders (collectively called the "hyperscalers") like Oracle, SpaceX, and Meta (formerly Facebook), are expected to spend something like $650-$700 billion constructing AI data centers. In total, since 2023, they'll have spent around $1.6 trillion. What kind of return on these massive investments are they receiving? Not as much as you'd expect. Microsoft's annualized AI-related revenue is reportedly $37 billion, Amazon and Google's are each estimated to be around $15 billion, and Meta's is zero (its AI spending is internally directed towards improving its core business of delivering ads on Facebook and Instagram). Finally, xAI (part of SpaceX) and Oracle might add another $15 billion and $10 billion, respectively, to industry-wide AI revenue this year.

Adding it all up, the hyperscalers will produce around $92 billion in AI-related revenue in 2026, most of which will come from the two biggest and wildly unprofitable AI model makers, OpenAI and Anthropic. Meta and Google are also using AI to improve ad delivery, but it's hard to say how much this has directly increased their revenue. Still, even $150 billion in total hyperscaler AI revenue on $1.6 trillion in spending (9.4%) is pedestrian. It would be an

extremely disappointing outcome, by comparison, if a plumber added a $100,000 truck to his fleet, and it only generated an additional $9,400 in annual sales.

Moreover, revenue isn't profit. AI data centers are far more expensive to run than traditional ones and so operate at much lower margins. And remember the hyperscalers are renting their data center computing power to OpenAI and Anthropic at discounted rates. My guess would be that from a pure profit standpoint, the hyperscalers so far have little, if anything, to show for their $1.6 trillion in spending. They're underwriting the biggest gold rush in history, but so far none of the miners have struck gold.

Clearly, this isn't sustainable. Either AI models must improve enough that OpenAI and Anthropic can charge a price that substantially exceeds their cost, or the expense of training and operating the models themselves must fall dramatically. Neither outcome looks likely anytime soon.

Earlier this year, in an attempt to better match their pricing with underlying expenses, OpenAI and Anthropic began moving customers away from their generous flat-rate subscription plans and toward pricing based on actual usage. Heavy users, especially software engineers relying on AI for coding, saw their costs explode from a few hundred dollars a month to thousands or even tens of thousands. Uber, for example, reportedly spent its entire AI budget for 2026 in just the first four months of the year, forcing the company to limit how much its engineers could use the AI models going forward.

In addition to usage caps, companies are increasingly routing more routine tasks to less powerful but open source and lower-cost models like China's DeepSeek while only using OpenAI and Anthropic's models for more demanding work. And despite already operating deeply in the red, OpenAI has recently been considering cutting prices to better compete with Anthropic and the other model makers. It may be that, even if AI model capabilities exceed expectations, competitive pressures could still make it difficult to turn a profit.

Will the models improve? Certainly, but as I've discussed in previous newsletter articles, they suffer from a structural limitation – AI technology is probabilistic, not deterministic. AIs like ChatGPT work by predicting the next token (word or part of a word) in a sequence based on statistical probability. This is why they will produce different answers to the same question and are inherently vulnerable to errors (hallucinations). In contrast, traditional software programs like Microsoft Excel are deterministic and follow strict sets of rules with a predictable, repeatable, and verifiable process. So long as you input the correct data into its spreadsheets, Excel's calculations will always be right.

Unfortunately, a calculator that's correct even 97% of the time isn't useful for most business applications. You can't manage inventory with it, payroll, draft legal documents or do anything that requires precision without verifying the model's output. In contrast, tasks like computer coding, ad targeting, customer service, or drafting an email can still create value despite occasional errors and are thus practical use cases for AI. But it's difficult to see how those applications alone can generate sufficient returns to justify the trillions of dollars hyperscalers will ultimately spend on AI data center construction.

OpenAI and Anthropic have committed to purchasing almost $900 billion of computing power from the major hyperscalers in the next five years. This figure represents the majority of the hyperscalers' reported backlog increase over the last year or so. Given the two major AI model makers' combined revenues of $60-$70 billion, tens of billions in operating losses, and increasing competition between the major models and cheaper, open-source ones like DeepSeek, it's not clear where the money will come from to honor these obligations.

But even if the major AI model makers do come up with the money, the hyperscalers themselves can't deliver on their end, as building new AI data centers has proven much more difficult than expected. A mystery I'd been trying to solve was that despite the hyperscalers spending over a trillion dollars on AI infrastructure, OpenAI and Anthropic are still struggling to secure the computing power they need to run their models. How can this be?

The biggest issue is not that the models are incredibly expensive to run (though they are). Rather, it's that the hyperscalers are running into all kinds of bottlenecks to complete (and in many cases even begin) construction on their AI data center projects. Globally, 4.6 gigawatts of AI data center capacity came online in 2024 – Note that one gigawatt is enough energy to simultaneously power 750,000 homes. In 2025 another 5.8 gigawatts was reportedly completed.

But keep in mind that these 2024/2025 vintage data centers broke ground and ordered their equipment from 2021-2023. Back then, lead times for high-power electrical transformers (what connects a data center to an electrical grid) were 24 months. Today, the wait time is five years! Further, electric utilities had the extra grid capacity to accommodate the data centers' massive energy consumption. Today, that spare capacity is mostly exhausted, which means additional generation and grid infrastructure will have to be built to power future data centers – a process likely to take years and in many cases more than a decade.

Moreover, the data centers completed in 2024/2025 used earlier generation Nvidia chips that required far less power with server racks drawing only 10 to 15 kilowatts that traditional air-cooling and electrical systems could handle. Today, a single modern rack (powered by new chips like Nvidia's Blackwell) can demand up to 140 kilowatts. This huge increase in power density means modern data centers require entirely new, incredibly complex cooling and heavy-duty electrical setups to prevent the chips from melting.

In fact, Nvidia's Blackwell chips generate so much heat they cannot be cooled by air alone and require liquid cooling with a steady stream of fluid flowing directly over the processors. This needs industrial grade piping networks and high-pressure pumps with perfectly leak-proof plumbing, as even a tiny rupture could destroy tens of millions of dollars' worth of hardware. All this supporting equipment means that a fully loaded Blackwell server rack weighs more than 6,000 pounds, so heavy that new data centers must have specialized reinforced concrete slabs and structural steel framing to keep the floor from literally collapsing. I've read that building a modern AI data center is akin to constructing an industrial nuclear cooling facility with the precision requirements of a semiconductor manufacturing plant.

Data center construction bottlenecks already began to bite in 2025. Although 5.8 gigawatts of new capacity was brought online (and it's questionable how much of this is fully operational) another 3.2 gigawatts was delayed. This year, 16.2 gigawatts of capacity was supposed to be

built, but only 5 gigawatts is currently under construction! An outlandish 31.2 gigawatts is planned for 2027, yet only a fraction of these planned facilities have even broken ground.

This begs the question if much of the planned data center capacity remains unbuilt due to roadblocks like power and grid constraints, equipment shortages, and the sheer complexity of construction, where is all the money the hyperscalers are spending actually going?

Back in November 2025 Microsoft's CEO Satya Nadella made an interesting admission saying, "The biggest issue we are now having is not a compute glut, but it's power. It's sort of the ability to get the builds done fast enough close to power. *S*o, if you can't do that, 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." Yet despite this concern, over the last nine months, Microsoft spent *another* $50 billion or so on semiconductor chips! Given that the lack of progress with AI data center construction is an industry-wide phenomenon, I suspect that other hyperscalers like Amazon, Google, Meta, Oracle, and SpaceX also have tens of billions of semiconductor chips gathering dust somewhere in warehouses.

So we have a gold rush where no one has yet found gold (a profitable AI business model). Meanwhile, the cost and time required to build new mines (AI data centers) have risen so rapidly that most planned projects remain stuck on the drawing board. Undaunted, the mining companies (the hyperscalers) continue to splurge on equipment (semiconductors), much of which is just being stockpiled in warehouses awaiting, perhaps indefinitely, the infrastructure needed to deploy it.

Where does this leave us? Fifteen years ago I wrote a newsletter article titled "What the Future Doesn't Hold." In it I discussed a forecasting method described by the author and securities trader Nassim Taleb in his book *The Black Swan.* Taleb explained that most people try to predict the future by imagining what will exist *in addition to* what already exists – for instance, new technologies or political trends. But this approach is exactly backwards. Rather than adding to the present to make predictions, the would-be prophet should instead look at the world and try to imagine the current systems that *won't* persist in the future. Predicting what won't continue is much simpler because, as Taleb wrote, "What is fragile will eventually break; and, luckily, we can easily tell what is fragile."

What makes a system fragile? Taleb identified a few major red flags. First, it's relatively new. For instance, if I had to predict which political system will likely still be around in a century – the Eurozone (began in 1999) or America's constitutional framework (1789) – the latter, which has already stood the test of time, would be the clear favorite.

The second marker of a system's fragility is interconnectedness and concentration. Democracies may appear extremely chaotic on the surface but the diversity they embody is actually far more stable long term than, say, a dictatorship, where power is concentrated in a single source.
The strength inherent in diverse, non-interconnected systems can be seen in the simple example of dominoes on a table. If the dominoes are numerous and placed at random, tipping one over likely won't cause many others to fall. But if the dominoes are arranged in an orderly fashion, knocking one down could cause a catastrophic collapse. Similarly, a table with just a few enormous dominoes – even if randomly distributed – is also more vulnerable to disaster.

The AI industry, in our view, fully meets the criteria of a fragile system. It's a brand new industry dominated by a handful of highly interconnected companies built on tenuous economic fundamentals – a table with just a few large dominoes standing in a row. While it's hard to say when they might topple, to us a cascade looks inevitable. Accordingly, while we continue to maintain some exposure to the industry, prudence dictates that we focus solely on those companies with the strongest fundamentals and limit the concentration of your stock dollars in this area to a level we believe strikes the right balance between risk and reward.

AMD: Investing in Inference

by Jeremy Goldberg, CFA, CFP®, MSF

Over the last two years, few businesses have captured more headlines than Nvidia. Its chips train large language models like ChatGPT and Claude, its products power much of the artificial intelligence (AI) infrastructure being installed in today's data centers, and its software has become a major part of the AI developer ecosystem. Nvidia is also one of the largest companies in the world and has grown to more than 7% of the S&P 500 index.

At PASI, we understood why investors were drawn to Nvidia, but in late 2024, we believed Advanced Micro Devices (AMD) offered a more attractive path. Despite being smaller and easy to overlook, it was perfectly positioned for the next phase of AI.

There are two fundamental stages of AI: training and inference. Training uses massive datasets to teach a model how to recognize patterns. Inference comes afterward, when that trained model is put to work, like when ChatGPT generates a response to a question. Training requires enormous computing power and is typically done using graphics processing units, or GPUs, because GPUs are especially good at running many calculations at the same time. Nvidia has been the clear leader in this market. Inference also often uses GPUs, especially at large scale, but the priorities are different: cost, memory, and efficiency. That's where we felt AMD had an opportunity to make inroads.

So far, that decision has been incredibly rewarding. From our initial purchase in December 2024 through the end of the second quarter of 2026, AMD's total return was nearly seven times Nvidia's.[3] That doesn't mean AMD will outperform Nvidia forever. It simply highlights an important investment lesson: the right opportunity is not always the most obvious one, even in a powerful long-term trend.

AMD designs advanced semiconductors used in data centers, personal computers, gaming consoles, and industrial applications. Semiconductors are the tiny electronic components that make modern computing possible. They act like microscopic switches, controlling electrical signals so computers can process information and perform calculations at extraordinary speed. In AMD's case, that can mean powering a laptop, running an Xbox or PlayStation, helping a company process data in the cloud, or helping an AI model generate a response.

[3] From our initial AMD purchase on 12/18/2024 through 6/30/2026, AMD shares generated a total return of approximately 378.50% vs. approximately 55.40% for Nvidia. Source: Bloomberg L.P.

Historically, AMD was known mostly as a competitor to Intel in PC and server processors. But it was a much smaller company than it is today. In 2017, AMD generated approximately $5.3 billion in revenue and had very little presence in the server market. By 2025, revenue had soared to $34.6 billion, driven in large part by AMD's growing role in data centers.

A major factor for AMD's data center growth has been customer adoption of EPYC, AMD's line of server processors. EPYC processors are central processing units, or CPUs. If GPUs are the specialized chips doing the heavy mathematical work behind AI, CPUs are more like system managers. A data center may use GPUs to generate an AI response, but it still needs CPUs to route the request, move the right information, assign work to the right chips, and keep thousands of tasks running at the same time.

Before EPYC, Intel controlled roughly 90% of the server processor market. By the end of 2025, AMD had captured roughly a quarter of that market by units, with an even higher share by revenue. That matters because, as AI moves from *building* models to *using* them across real businesses, the supporting data center workloads become more important. Management also recently doubled its estimate for the server CPU opportunity to more than $120 billion by 2030, specifically because AI workloads are creating more demand for these processors.[4]

A key advantage for AMD has been its chiplet design. Rather than building every part of a processor on one large piece of silicon, AMD builds smaller chiplets and connects them inside the same processor package so they work together as one. This gives AMD more flexibility to add computing power, improve efficiency, and control cost. For data center customers buying thousands of processors, even small improvements in speed, power use, and cost can translate into meaningful savings. Today, AMD's EPYC processors are used by many of the largest cloud and enterprise customers in the world, including the companies that run the digital services we use every day.

The other part of the story is AMD's Instinct GPU business. Nvidia remains the category leader, but AMD is now supplying AI GPUs to some of the largest technology companies, including Microsoft and Meta. Its MI300 chip helped establish those relationships, while newer versions, including MI350 and the upcoming MI450, are designed for larger AI workloads.

There is also a strategic reason customers want AMD to succeed. The largest technology companies don't want to rely entirely on one supplier for the most important hardware of the AI era. Even if Nvidia continues to lead, companies like Microsoft, Meta, Oracle, and OpenAI benefit from having another serious alternative. It improves supply security, provides negotiating leverage, and gives them more flexibility as AI infrastructure evolves.

Taken together, AMD is more than a single-chip AI story. Its Instinct GPUs perform the heavy calculations needed for training and large-scale inference. Its EPYC CPUs help data centers route requests, move information, and keep thousands of AI tasks running at the same time. And the 2022 acquisition of Xilinx, a semiconductor company known for programmable chips, added products that can be customized for specific jobs, like helping AI-powered cameras spot defects on a manufacturing line.

[4] AMD 1Q26 Investor Presentation (5/5/2026)

Our goal is not to own the company that gets the most headlines. Our goal is to identify well-managed businesses where the fundamentals and long-term opportunity can reward patient shareholders. In December 2024, AMD looked like one of those opportunities. We believed the focus would move from building AI models to putting those models to work in the real world, and that shift would require more than one type of chip. Today, AMD's CPUs, GPUs, and specialized chips make it an essential part of the infrastructure needed to make AI work at scale.

Capital Gains

by David A. Jaffe, M.D.

A sage investment caveat: "make investment decisions first, tax decisions second." While we do our best to minimize realized capital gains when managing taxable accounts for our clients, there are times when extraordinary stock price appreciation far exceeds the fundamental business prospects of a company, leading us to trim positions and take profits.

Demand for semiconductors is expected to explode in conjunction with adoption of Artificial Intelligence and announced investment in new data centers. However, as detailed in the earlier article by Nathan Polackwich, *Where's the Gold?*, there are reasons for skepticism. Still, as illustrated graphically by the following chart, the 2026 price behavior of related stocks, including Applied Materials (AMAT), Advanced Micro Devices (AMD), Lam Research (LRCX), and ASML (ASML), has been extraordinary indeed.

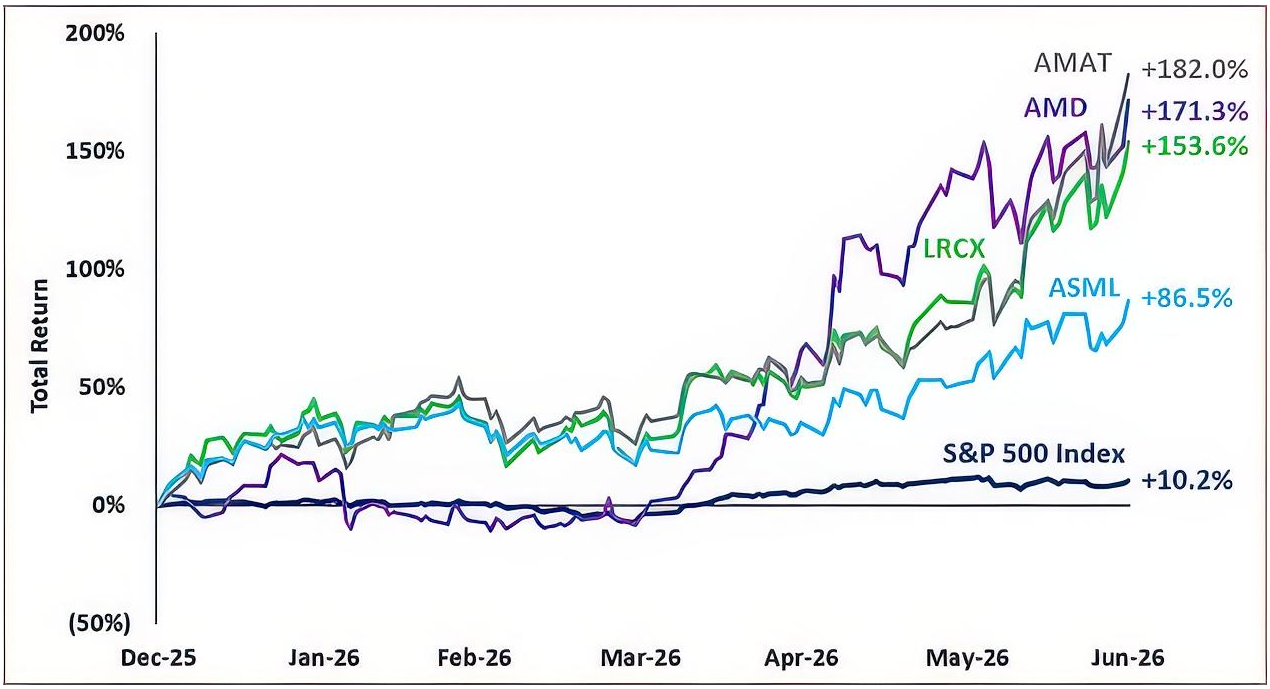

Source: Bloomberg L.P. from 12/31/2025 through 6/30/2026.

Looking for reasonably valued participants in the AI buildout, we added these companies to our portfolio in late 2024 and early 2025, and at much lower prices. Despite substantial capital gains, risk management dictates that we trim positions to maintain our target investment exposure and capture profits.

The bottom line – current realized capital gains for our taxable accounts are substantially higher than has been typical in past years. For some clients, it may be appropriate to adjust your planned tax withholding. We will be following this newsletter with individual letters to clients with taxable accounts providing current realized capital gains information. We encourage you to review this with your accountant and make adjustments if necessary.

Hurricane Season – *It's That Time of Year Again*

Officially ranging from June 1st to November 30th, Hurricane Season is once again upon us. Our hurricane procedures include preparation of our physical office to minimize the impact of water intrusion and protecting the integrity of our central computer. Most of our team will be able to plug in phones and laptops and be ready to serve your needs quickly, limited only by availability of electricity. David's Montana office serves as remote backup. BNY will be alerted to initiate contingency processing for client needs.

While PASI has settled on a staffing structure minimizing "work from home" time, we surely have the experience, as well as resources and technology, to serve client needs remotely if required.

You can read our Disaster Recovery Policy on our website www.pa-services.com. Please follow "contact us" from our main menu and choose the "emergencies" entry. You will find a link to our Disaster Recovery Policy on that page. In the event that primary communications are affected by a storm, we will post updates and any important information on our website. If you have any questions about our contingency planning, please don't hesitate to call.

Disclosure

Professional Advisory Services, Inc. may, from time to time, have a position in securities mentioned in this newsletter and may execute transactions that may no longer be consistent with this presentation's conclusions. Reference to investment performance of the PASI composite stock portfolio is made gross of expenses. For formal performance disclosure with net returns please contact our office.

Performance Disclosure

To obtain a detailed analysis of Professional Advisory Services, Inc.'s (PASI) historical performance, inclusive of gross and net results from our balanced accounts and performance data for our segregated asset classes, please contact our office at 800-847-7274. It is important to note that PASI performance data presented in this newsletter is stated before the deduction of fees and in the context of each article. For a clearer understanding of the impact of fees, please refer to the following disclosures including a hypothetical example based on the maximum PASI investment management fee.

The **PASI Stock Portfolio** includes the reinvestment of dividends; and is reduced by brokerage commissions but is gross of Professional Advisory Services, Inc. fee, which is described in Part II of form ADV, available upon request. Our fee is a maximum of 1% and decreases based on assets under management. As an example of fee impact, over a ten-year period, $100,000 invested in stocks growing at 8% per year would increase at the end of ten years to $205,419 net of 1% fee versus $220,804 gross return.

PASI Stock Portfolio Benchmark: The ***S&P 500 Index (Market-Cap-Weighted)*** is an unmanaged index of the 500 leading publicly traded common stocks in the U.S., including reinvestment of dividends. This index is weighted according to the market capitalization of each participating company. As a result, companies with larger market capitalizations exert greater influence on the index's overall return, reflecting their proportionate size to the overall market.

Other Indices: The ***S&P 500 Equal Weight Index (Equal-Weight)*** is an unmanaged index of the 500 leading publicly traded common stocks in the U.S., including reinvestment of dividends. Designed to be size-neutral, it assigns equal weight to each participating company, irrespective of their market capitalization. This approach equally captures the influence of each company on the index's overall return relative to its individual performance, providing a balanced reflection of the collective market activity.