



Volume 49 - Number 3

Fall 2025

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 – *Is No News Really Good News?*

by David A. Jaffe, M.D.

The United States Congress ended the third quarter of 2025 by shutting down the Federal government after the Senate failed to pass a continuing resolution which would have funded operations through November 21. Financial markets shrugged.

Historically, government shutdowns have had little direct economic impact. Of note, however, most prior government shutdowns have been brief in duration. The risk that the current shutdown exceeds existing norms looms; the economic impact may exceed past experience as well.

A worrying feature of the current environment has been the interrupted collection and processing of important economic data disseminated by the Federal government and scrutinized by economists worldwide. One critical measure of the U.S. economy is the health of the jobs market, the report normally released on the first Friday of each month. Thanks to the government shut down – no news for September. Another victim of the shutdown has been the release of the September consumer price index (CPI), a measure updated monthly. Circulation of the September CPI, compiled by the Bureau of Labor Statistics (BLS), is promised by October 24th. This has received priority status because the information is critical to the annual cost-of-living adjustment (COLA) for social security payments and other government programs.

In addition to impaired timeliness, there has been deterioration in the usual depth of national data collection, a labor-intensive process handled by BLS field operatives. Due to earlier staff cuts by the Department of Government Efficiency (DOGE) operations, the volume of imputed pricing data (derived by statistical estimation) incorporated in CPI calculations, rather than direct measurement, has risen from 10% historically to a current level of about 33%. The consequence is diminished quality and damaged confidence in the data.

Beyond the September CPI, data collection remains suspended. Access to dependable economic data is more than just an academic concern. The Federal Reserve lowered short-term interest rates in September, the first cut since December 2024, helping to fuel healthy stock market gains during the quarter. While investors have worried that current tariff policies may spark inflation, the August employment statistics suggested a weakening job market, leaving the Federal Reserve facing a difficult conundrum. Rising prices in the face of a weak economy, a phenomenon dubbed

“stagflation,” was last seen in the 1970s following the oil embargo. The challenge is that policy intervention to manage inflation would likely depress economic growth and hiring, while efforts to boost the economy generally boost prices as well. Without reliable and timely economic data, the challenge for the Federal Reserve is heightened.

Despite current uncertainties, the stock market continued its relentless advance in the third quarter. Investor attitudes may be simplest captured by a bit of economic satire seen online recently, a poster of old friend Alfred E. Neuman and his trademark refrain: “What Me Worry?” (to our younger readers – “Google it”, or maybe better, ask AI!).

The Market Cap Weighted (MW) S&P 500 added over 8% in the third quarter alone, fueled largely by appreciation of large AI-related companies and enthusiasm regarding industry developments (addressed by Nathan Polackwich in the article which follows). The MW S&P 500 ended the quarter with 2025 year-to-date returns of 14.83%. The broadly diversified Equal Weighted (EW) S&P 500 rose 4.84% in the third quarter, posting a YTD gain of 9.90%. The PASI stock portfolio added 5.53% in the third quarter, showing a year-to-date gain of 9.46%.¹ All returns cited here include reinvested dividends.

There is no denying that this has been a challenging environment for conservative investors who believe in the value of broad diversification for risk management. We have not seen the current degree of sector concentration (technology and more specifically Artificial Intelligence or AI) since the “dot-com bubble” of the late 1990s. Stock valuations for related businesses reflect very high prices and carry the risk of an abrupt reversal if expectations are not met. At PASI we have worked to identify beneficiaries of the extraordinary ongoing and scheduled capital investment in AI which meet our requirements for financial strength and reasonable stock prices relative to growth potential. Additions in the last year include Oracle, Advanced Micro Devices, three semiconductor manufacturing equipment suppliers (Applied Materials, ASML, and LAM Research), and indirect beneficiary Jacobs Solutions (see Jeremy Goldberg’s article later in this newsletter). We also maintain investments in participants Alphabet (Google), Amazon (notably Amazon Web Services or AWS), and Microsoft (a substantial OpenAI partner and investor).

In the context of the AI goldrush, our selections largely provide “picks and shovels” to the prospectors laboring to foster the evolution of AI models. As Nathan discusses in the following article, whether these AI pioneers are able to channel their efforts into profitable businesses is an open question. We will be happy to support the infrastructure providers necessary for the journey and cheer the prospectors along the way.

The Numbers Don’t Add Up

by Nathan Polackwich, CFA

PASI holding Oracle (ORCL) saw its stock surge 36% on September 10th, gaining \$255 billion in market capitalization after announcing its backlog (technically called “Remaining Performance Obligations”) jumped from \$138 billion to \$455 billion in a single quarter. Oracle also disclosed, however, that \$300 billion of the increase related to supplying computing power to a single customer – OpenAI – the company behind ChatGPT, and that the deal doesn’t begin until 2027.

Then on October 6th the stock of a second PASI company, Advanced Micro Devices (AMD), experienced a similar move, soaring 24% on the news that OpenAI would buy tens of billions of

¹ Please see our disclosures on page 10.

dollars' worth of AMD chips in the coming years. This agreement also gives OpenAI an opportunity to take up to a 10% stake in AMD contingent on the achievement of operational milestones.

While the gains in both stocks are certainly welcome, we're skeptical that OpenAI has anywhere near the financial capacity to fulfill its end of the contracts. These are just two of many examples of implausible numbers increasingly being touted by companies involved in the AI infrastructure buildout.

As it stands, OpenAI is estimated to generate about \$12 billion in revenue this year while burning over \$8 billion in cash. Losses are expected to more than double next year to \$17 billion and that's before the Oracle contract starts. At that point, OpenAI will have to raise an additional \$30-\$40 billion just for Oracle with that number increasing to perhaps \$100 billion by 2030. In addition to tens of billions of dollars in other financial commitments (AMD is just one of many), OpenAI also plans to spend *another* \$100 billion from 2026-2030 to rent servers from major cloud providers for backup computing "aimed at powering sudden or unforeseen AI research breakthroughs." Keep in mind OpenAI itself admits it has no path to profitability and sees its losses continuing to escalate in the coming years.

So where will it get the money? The entire U.S. venture capital industry only has about \$300 billion in dry powder capital and perhaps another \$600 billion globally. But OpenAI can't access most of this because a significant portion is earmarked for earlier stage investments (OpenAI is late stage), and most venture capital firms won't concentrate their bets on a single position. This is why even though its losses and financing requirements have so far been relatively mild in comparison to what's coming, OpenAI has already been forced to cobble together financing from multiple sources. These include not just venture capital but equity stakes from suppliers like Microsoft and NVIDIA, convertible (into equity) debt issuance, asset backed/leased financing for computing hardware, vendor financing for computing capacity, and special purpose vehicles (SPVs) funded by banks and insurers to finance data center construction. Scaling this patchwork capital to hundreds of billions or even a trillion dollars plus is well beyond the reach of private markets and would require U.S. government backing at a level no private company has ever received.

AI infrastructure costs have ballooned because the models lack the positive economies of scale traditionally enjoyed by the software industry. With software, once the code is written, there's basically zero marginal cost to add a new user and little processing effort for each command, as software just reuses pre-written instructions (writing the code was the hard part). AIs like ChatGPT, conversely, lack this dynamic and in a sense operate more like an electric utility. Specifically, as the number of users and workload increases, capacity (in the form of semiconductor chips, power, cooling, and data center buildouts) must scale roughly in line with demand. Because the AIs lack predefined answers, every question, even identical ones, force the model to run its servers and regenerate the results from scratch.

This means AI models consume massive computing power and thus also substantial electricity, partly to cool the data centers that house the servers. OpenAI's current commitments through 2030 alone imply a need for at least ten gigawatts of new electric capacity, enough to power approximately 8 million homes.

Given OpenAI's commitment to reach net-zero carbon, ten gigawatts of new electric capacity would effectively require building ten new nuclear reactors (data centers need continuous power, which intermittent sources like wind and solar can't provide). But with an estimated cost of at least \$10 billion per gigawatt, skilled labor and equipment shortages, regulatory resistance, and transmission bottlenecks, I put the odds near zero that anything close to this can be constructed and operational within the next five years. For instance, the only nuclear reactors built in the U.S. in the past three

decades (the recently completed Vogtle units 3 and 4 in Georgia) were originally expected to take 4-5 years and cost \$14 billion but ultimately required 14 years and more than \$35 billion.

Some hope that small modular nuclear reactors (SMRs) might offer a solution, but that also seems unlikely. Instead of getting ten conventional reactors through the site planning and regulatory process, you'd now need to do it for dozens. Moreover, SMRs' cost per gigawatt has actually proven to be at least as high as conventional nuclear plants, which benefit from better economies of scale due to their larger size. In theory, mass production could eventually lower SMRs' construction and deployment costs, but with no such manufacturing infrastructure yet in place, they will remain expensive prototypes for the foreseeable future.

Unlike most industries in the tech sector where you see advances accelerate and productivity increase over time, the opposite is occurring with the cutting-edge AI models the big U.S. tech companies are trying to develop into Artificial General Intelligence (AGI) – an AI that can understand, learn, and apply knowledge with human-like flexibility. This quest is reminiscent of Einstein's Theory of Relativity – as mass nears lightspeed, energy demands approach infinity. Similarly, while early advances with the latest AI models like ChatGPT came relatively cheaply, staggering sums are now being spent to deliver increasingly modest improvements.

For example, ChatGPT-3.5 (released in 2022) cost about \$5 million to train and reached 44% accuracy; GPT-4 (2023) cost \$100 million for 86%; and GPT-5 (2025), at over \$1 billion, advanced performance only marginally to about 90–92%. Improvements have become increasingly difficult due to the technology's inherent limitations and a lack of fresh, high quality training data. We're now in the midst of the biggest investment bet in history with spending on AI infrastructure expected to reach \$375 billion this year and \$500 billion in 2026.

Meanwhile, Chinese AI company DeepSeek has gone the opposite direction, engineering its models to drastically cut costs. In November 2024 it released its V3 model, trained for just \$6 million yet with performance close to ChatGPT-4. Then in February 2025 DeepSeek released its R1 model, trained for just \$294,000 and capable of running at about 2% of ChatGPT's cost, with nearly comparable performance.

Most concerningly for the big tech companies hoping to ever cash in on their outsized AI spending, the DeepSeek model is open-source and freely replicable, which means any company can now create a powerful AI for practically nothing. In effect, DeepSeek has unleashed a deflationary force that's upended the trillion-dollar AI buildout narrative. Yet the major U.S. tech companies continue to push ahead with their exorbitant spending undeterred, chasing the dream of Artificial General Intelligence.

Text based models like ChatGPT and DeepSeek as well as image/video creation models like Dall-E and Midjourney are called *Generative* AIs. They represent a new class of artificial intelligence systems that learn patterns from existing data and generate new content (everything from legal documents to poems to cartoon images of politicians) that reflect and build upon those patterns.

The consensus among investors seems to be that Generative AI is the next great technological wave, analogous to the buildout of the Internet in the late 1990s and cloud computing from the mid-2000s through the 2010s. But the key difference between this wave and the two previous is that both the Internet and cloud solved immediate problems, dramatically lowering companies' costs while spawning new business models and even entire industries. Generative AI, conversely, still looks more like an expensive solution in search of a problem.

The Internet, of course, was an information revolution. There was suddenly no limit to how quickly people could communicate, and much of the world's knowledge became available to anyone with

the click of a mouse. Businesses could now reach their customers anywhere on the planet and fulfill orders instantly. Within a short time the Internet gave rise to all sorts of profitable new industries like ecommerce, search, social media, and many others.

The cloud revolution, while not quite as groundbreaking, did address another key challenge – companies were wasting significant resources building and maintaining their own IT systems. This infrastructure also had to be capable of handling peak demand (like Black Friday online retail spending faced by Amazon and other online retailers), which meant significant computing capacity sat idle much of the time.

In the early 2000s Amazon's management realized that its idle capacity problem wasn't unique and they could rent it to other companies on a flexible, ad-hoc basis. This insight led to the formation of Amazon Web Services (AWS) in 2006, which, like the Internet, immediately lowered companies' costs and fueled new businesses like Software as a Service (delivered over the cloud rather than installed and run on local hardware), the gig economy (e.g. Uber and DoorDash), mobile apps like Tik-Tok and Snapchat, and even the Generative AI startups like ChatGPT and Anthropic that are the focus of this article.

So why is Generative AI different? There's no doubt the technology is useful for tasks like summarizing information, creating new text and images, translation, and computer coding. But despite hundreds of billions in spending and years of development, so far no successful new companies have emerged from the technology. In fact, the entire industry only produces about \$50-\$60 billion in revenue and most is either unprofitable or inflated. For example, Microsoft claims \$13 billion in AI revenue, but \$10 billion is attributable to OpenAI using Microsoft's cloud at cost (no margin – and in return Microsoft gets a share of OpenAI's nonexistent profits).

The most innovative new Generative AI businesses are the core model developers, particularly OpenAI and Anthropic. But as noted earlier, both earn only modest revenue while producing billions in losses. And unlike the Internet and cloud, other AI startups have struggled to build profitable companies on top of these core models, as their intellectual property is easily replicable (by OpenAI or Anthropic), and their margins are at the mercy of the model providers who themselves are losing money. I do expect attempts to build businesses using the more economical open-source models like DeepSeek, but so far that hasn't happened perhaps suggesting the technology itself, no matter how cheap, isn't easily turned into a functional, salable product.

If Generative AI fails to spur the formation of profitable new industries, the economic justification for the huge infrastructure spending will rest solely on existing companies' use of the technology to cut costs and improve productivity. But so far they've had little success. A recent report from MIT Media Lab found that “despite \$30-\$40 billion in enterprise investment into GenAI,...95% of organizations are getting zero return... Most fail due to brittle workflows [the models break down when faced with the messiness of real-world business operations], lack of contextual learning, and misalignment with day-to-day operations.”

A Harvard/Stanford research collaboration (as reported in the Harvard Business Review) identified another issue: “Employees are using AI tools to create low-effort, passable looking work that ends up creating more work for their coworkers... While some employees are using this ability to polish good work, others use it to create content that is actually unhelpful, incomplete, or missing crucial context about the project at hand. [This] shifts the burden of the work downstream, requiring the receiver to interpret, correct, or redo the work.”

The Achilles Heel of Generative AIs is that they often make mistakes, and this flaw may be insurmountable because it stems from the inherent nature of the technology. Specifically, the models

are *probabilistic*. Large Language Models (LLMs – Generative AIs that produce text) like ChatGPT, for instance, work by predicting the next word in a sequence based on statistical probability. This is why an LLM can produce different answers to the same question and will always remain vulnerable to error. Worse, if asked something it doesn't know, an LLM will just "hallucinate" an answer that sounds good. Since the models don't reason or understand the text they generate, they don't know when they don't know something.

Conversely, traditional software programs like Microsoft Excel are *deterministic* and follow strict sets of rules with a predictable, repeatable, and verifiable process. Excel, for instance, will calculate correctly 100% of the time. In contrast, Generative AIs still struggle with something as basic as building and computing numbers within tables. The result is that despite astonishingly powerful semiconductor chips, enormous data centers, and massive energy consumption, the AIs aren't close to being better than existing software for most applications. This is especially true when considering their cost.

The technology's fundamental limitations mean that Generative AIs are unlikely to take over companies' IT systems, supplant traditional software, or replace most human workers any time soon, if ever. They can't be used for mission critical applications like healthcare, accounting, legal compliance, engineering calculations, and many other areas. Even in industries where they can be useful such as advertising, content generation (e.g. blog posts, news articles, product descriptions), and digital image design, humans and/or traditional software will still have to check and refine the AIs' output.

The dream is that Generative AI will one day become a sort of "everything machine" capable of completing any task. But it now appears it will end up like my ability at sports – a jack of all trades but master of none. It can regurgitate the ideas of others, but because it doesn't reason or understand the words it uses, it can't produce new thoughts of its own. So ChatGPT excels at, for instance, writing boilerplate emails but not an original, well-reasoned essay. Similarly, it can explain how a car works but can't invent a unique design for a better one.

So how is this likely to shake out from an investment perspective? I'm skeptical that many of the pie-in-the-sky projections for hundreds of billions or even trillions in spending for future data center and electric capacity construction will ever materialize. While Generative AI will succeed as a technology and likely become ubiquitous in daily life, what it does best – generating text, summarizing, editing, transcribing, and translating – doesn't require the massively expensive cutting-edge models like ChatGPT or the latest NVIDIA chips.

From a return on investment standpoint, it's more practical to train AI models for narrow use cases. There's no reason, for instance, that a travel planning AI needs to train on the complete works of Shakespeare or the physics of satellites. Open-source models like DeepSeek running on inexpensive chips like ASICs (Application-Specific Integrated Circuits – chips custom-designed to perform specific functions), can handle these more limited tasks far more efficiently and cost-effectively. In fact, Google already relies heavily on its own ASICs for its AI models.

One of the more interesting use-cases for AI is video generation. But even this is subject to the same limitations (magnified) as text-based models like ChatGPT. Computational and energy demands are much greater, the outputs remain inconsistent, requiring substantial human editing, and open-source competition is growing. In addition, copyright infringement risks are particularly acute since most of the models are trained on movies, animations, and clips whose creators didn't consent to their use.

If functional Generative AI models can be trained and run relatively cheaply and no single company has a lock on the technology, it's hard to see much long-term pricing power or profit potential for the industry. That said, some differentiation will be possible for companies that can pair AI with proprietary data such as Booking Holdings' unique travel information or Google's YouTube content or by integrating AI into existing software platforms like Microsoft's Copilot in Office or Adobe's Firefly in Photoshop.

Still, current and projected levels of AI infrastructure spending only make sense if Generative AIs eventually achieve Artificial General Intelligence, which the technology appears to lack the fundamental capability of reaching. And if AGI never arrives, then today's trillion-dollar AI build-out looks less like the next Internet or cloud revolution and more like one of the costliest misallocations of capital in history.

Jacobs Solutions: *A Quiet Leader in Global Rebuild*

by Jeremy Goldberg, CFA, CFP®, MSF

Jacobs Solutions may not be a household name, but it plays a central role in how the modern world is being rebuilt. From water systems and pharmaceutical facilities to airports, electric grids, and data-center campuses, Jacobs operates behind the scenes in projects that define the next decade of physical and digital infrastructure. Once known as a traditional engineering and construction contractor, the company has reshaped itself into a science-based consulting and advisory firm focused on some of the world's most complex sustainability, critical-infrastructure, and advanced-manufacturing challenges.

The company's transformation has been deliberate. Historically, Jacobs managed large industrial and government projects that were often low margin and cyclical. Over the last decade, management refined its portfolio through a series of divestitures and acquisitions, exiting segments such as its Energy, Chemicals, and Resources business unit in 2019 and adding higher-growth, higher-margin operations including CH2M, a global leader in water infrastructure and semiconductor facility design, in 2017, and PA Consulting in 2021, where Jacobs holds a 65% ownership stake. In 2024, the company completed the spin-off of its defense-heavy Critical Mission Solutions and Cyber & Intelligence divisions, leaving a simpler, more stable company focused on design, advisory, and specialized engineering expertise.

Today, Jacobs reports through two operating segments: Infrastructure & Advanced Facilities (I&AF) and PA Consulting. I&AF encompasses water, environmental, life-sciences, advanced manufacturing, transportation, and energy systems – the core of Jacobs' long-term project pipeline. PA Consulting contributes higher-margin, advisory-focused revenue. Based in the U.K., it helps clients across energy, healthcare, consumer goods, financial services, and government rethink strategies, design new products, and improve execution. Its public-sector business, particularly in U.K. defense and energy, shows strong momentum and backlog growth. Together, the two segments create an integrated model in which Jacobs can guide projects from concept through design, construction, and ongoing operations, building long-term relationships and steadier margins across cycles.

Jacobs' I&AF portfolio highlights the breadth of the company's work: Its teams design next-generation data center campuses for some of the world's largest cloud providers, including a recent collaboration with NVIDIA to build "digital twin" models that let engineers test cooling, power, and design efficiency before construction begins. In life sciences, Jacobs designed and built vaccine manufacturing facilities for Pfizer and AstraZeneca during the COVID-19 pandemic and is now

applying that same expertise to the fast-growing GLP-1 (semaglutide) market, where pharmaceutical companies are racing to expand capacity. Its water business has become a standout as climate and contamination challenges grow. Jacobs is helping utilities in Florida upgrade treatment plants to remove PFAS, the “forever chemicals” that have become a national priority, and continues to lead desalination and flood-protection projects globally. In transportation and energy, Jacobs manages large-scale projects ranging from the Denver International Airport expansion to the Seattle and Phoenix light-rail systems, and is partnering with Xcel Energy on its \$45 billion effort to upgrade and expand the regional power grid across its eight-state service area.

The scale of work ahead is substantial. Jacobs ended its most recent quarter with a record backlog of approximately \$22.7 billion, up 14.3% from last year, and management estimates a medium-term opportunity pipeline exceeding \$100 billion, according to its August 2025 investor presentation.

Several secular forces underpin Jacobs’ growth story. Governments are investing heavily in infrastructure resilience, clean water, and power-grid modernization. Semiconductor and data center construction continue to surge from artificial intelligence (AI) demand and supply chain reshoring. Pharmaceutical capacity is expanding globally to meet the growing need for new therapies and drugs. Each of these areas plays directly into Jacobs’ expertise, giving it multi-year visibility that few industrial peers can match. Jacobs is increasingly differentiating itself with digital engineering tools. In particular, it uses simulation models and sustainability analytics to stress-test designs before construction begins, which helps reduce rework, accelerate permit approval, and lower energy use in built assets.

Every company in this space faces risks. Delays in government funding or private-sector project starts could affect near-term revenue, and large programs can always encounter cost inflation or permitting issues. Jacobs’ contract structure and client diversity, however, cushion these effects. The company has shifted more of its work toward reimbursable contracts, which now represent about 70% of total revenue, lowering the risk of cost overruns that often plague fixed-price projects. The business is diversified across end markets, with only about 9% of revenue coming from U.S. federal agencies, primarily the Department of Defense, and the rest spread across state, municipal, private, and international clients.

Through 2029, management expects 6.0-8.0% annual revenue growth and operating margins expanding from 12.8% to 16.0% (a 25.0% increase). Over the same period, Wall Street analysts estimate earnings will rise from \$5.28 to \$10.31, a 14.3% annualized growth rate. Since 2020, Jacobs has repurchased more than 12 million shares, reducing share count by nearly 10%, and still has \$1.37 billion remaining under its current buyback program. The stock trades at 22.7x forward earnings – a reasonable valuation for such a high-quality operator with momentum across all its businesses.

Jacobs is not a short-term play on infrastructure headlines, but a long-term participant in the rebuilding of essential systems that support water, energy, data, and health. It’s a quiet leader in the global rebuild that we’re happy to own.

BNY Mellon's Credit & Lending Solutions

by Jeremy Goldberg, CFA, CFP®, MSF

As custodian for most PASI clients, BNY Mellon offers a range of credit and lending solutions designed to provide liquidity without disrupting long-term investment strategies. Several of the most relevant options for qualifying clients are outlined below. PASI does not receive any financial benefit from the use of these services, and we encourage clients to explore similar offerings from other financial institutions.

- **Investment Credit Lines:** Revolving lines of credit secured by marketable securities, offering access to liquidity for personal or business needs while keeping portfolios fully invested. Clients must have at least \$1 million in non-retirement investment account assets to qualify.
- **Commercial Real Estate Lending:** Customized financing for the acquisition or refinancing of retail, multifamily, industrial, mixed-use, or limited-service hospitality and office properties. Loans feature terms of up to 10 years, amortization up to 25 years, and minimum loan sizes of \$5 million.
- **Mortgage Lending:** Financing available nationwide for primary, secondary, and investment properties. Options include adjustable- and fixed-rate products, interest-only or fully amortizing loans, 100% financing through pledged assets, and construction-to-permanent loans.
- **Construction Financing:** Single-closing, residential construction-to-permanent loans with interest-only payments during the build period, rate-lock options at origination, and flexibility to pledge eligible investment portfolios as collateral to preserve invested assets.

BNY also offers a **Loan Modification** program that allows eligible borrowers to adjust interest rates or product types without full re-underwriting or appraisal, providing an efficient alternative to refinancing for a nominal fee. Note: This program applies to Mortgage Lending and Construction Financing only.

To learn more about these credit and lending options or to review how they may fit your situation, please contact your Portfolio Manager.

Return of the PASI Blog!

Recent discussions disseminated among our investment team led us to believe that much of the content would be of interest to clients. We have thus rekindled work on the long dormant PASI Blog, found on our website at www.pa-services.com/blog. Articles include topics such as review of current economic conditions, financial planning issues, and investment strategies. Our goal is to post a new article biweekly. We hope you find the content interesting and informative. Follow-up questions to your Portfolio Manager are encouraged.

We plan to follow this alert with an e-mail to clients. We promise not to flood your inbox!

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.