



Growth at Any Cost (Op-Ed)

Joseph Newhard, Ph.D. • Associate Professor of Economics

There are now 4,000 data centers operating nationwide with another 3,000 planned or under construction. Many local leaders around the country are lobbying to attract data centers to their own communities, alleging a boost to jobs, output, and tax revenue. Yet the ones already in existence have also made it impossible to hide their true costs: massive water and electricity usage, increased air temperature, pollution, and a loud humming sound that never goes away. They harken back to what William Blake called the “dark, satanic mills” of the industrial revolution, but instead of textiles, these mills produce malware, fraud, memes, clickbait, and deepfakes.

To control the heat generated by thousands of GPUs, cooling equipment produces a constant hum, and generator testing is even louder. Nearby residents report disturbed sleep and an inability to enjoy their homes and gardens. The GPUs and cooling units consume an unbelievable amount of electricity and deplete water wells, potentially driving up utility costs for neighbors. Cooling transfers heat outdoors, and this can increase the outside air temperature by several degrees. During blackouts, diesel generators emit nitrogen oxides and particulates. Large data centers can block views and decrease property values. They occupy large tracts of land, and the bulk of their job creation is in their construction. They enrich distant millionaires while producing no goods and services for local economies.

Despite this, in their meeting of May 15, 2025, the Johnson City Commission voted 5-0 to permit data centers and cryptocurrency mines within city limits with a mere 200’ setback over the objections of 25 citizens who spoke out unanimously in opposition. One such proposal is for a data center next to Buffalo Mtn on Innovation Dr. A temporary moratorium expires in December—after the election—while the city drafts regulations intended to mitigate the harmful effects of the data centers they seem intent on building here. Cities argue that data centers provide a large boost to tax revenue for schools, roads, and police while requiring relatively few government services in return, as if the only thing that matters in life is money. In reality, the net benefits are reduced by a web of tax abatements, tax credits, tax exemptions, preferential rates for electricity and water, publicly financed infrastructure, discounted land parcels, and expedited permitting.

A Bitcoin mine in Washington County near Limestone produced nonstop noise pollution for five years until it was finally shut down in 2026 due to a legal settlement. One resident told WJHL, “This was normally a quiet, peaceful place to be. You heard the birds. You heard farm sounds. Now we sound like we’re on a four-lane highway.” In Memphis, the noise from the SpaceXAI data center complex, which relies on gas-burning turbines, was recently described by one nearby resident as “like a form of torture.” Similar complaints and concerns about noise and pollution now number in the hundreds if not thousands. In response, nearby cities updated their zoning laws hoping to prevent experiencing the same misery. The NAACP also filed a lawsuit against SpaceXAI alleging air pollution.

In addition to air pollutants, stored fuel, construction runoff, and cooling system wastewater all present environmental and health risks and cleanup costs can be shifted onto taxpayers. A Meta data center contaminated the Cheyenne reclamation water supply with bacteria, and an Equinix data center in New Jersey leaked 5,000 gallons of diesel fuel into a river tributary. Recent diesel spills have also occurred in Texas, Iowa, Massachusetts, and Oregon and there have been significant construction runoff incidents in Georgia, Washington, and West Virginia.

Silicon Valley, which slurs middle America as “flyover country,” sees low-income Black neighborhoods and Appalachian Mountain hollers as dumping grounds for hazardous, loud, power-hungry facilities that they would never build in their own backyards. Economists who would measure data centers only by GDP and tax collections see no value in mountains and forests and rivers beyond their auction price. They would convert the Cherokee National Forest into timber and quartz if it maximized shareholder value. Our region, which severely lacks affordable workforce housing after years of mismanaged growth, must now contend with the government's compulsion to boost economic statistics at the expense of everything that makes life worth living here.

When unemployment is already low, to what end are data centers needed to “create jobs” when we must also import the workers to do those jobs? Growing the economy for its own sake doesn’t necessarily make our citizens better off and often makes them worse off in ways the GDP doesn’t capture: Dentist shortages, traffic congestion, air pollution, noise pollution, homelessness, crime, and the loss of our rural character as every last farm, field, and forest is bulldozed. The only growth that actually matters is of GDP per capita, but not at the cost of everything else. There is such a thing as an optimal city size and bigger is not always better. The politicians who prefer to live in a larger city have plenty of other options. Even they won’t be building data centers near their own neighborhoods.

An insatiable appetite for taxes explains why government officials are forever fixated on growth no matter how crowded the streets get or how strained the infrastructure becomes. It is why Johnson City paid remote workers to move here during the worst housing shortage in history while homelessness doubled. It’s why Western governments see it as their role to cram as many humans as possible within their own borders even as housing costs go parabolic, wages stagnate, and fertility rates accordingly crash. To manage growth, the federal government proposes opening 45 million acres of national forests to logging and road construction. There is also the peculiar disorder among local governments that actively recruit business investment while routinely blocking the construction of housing for the new workers to live in.


Given the severe reduction in quality of life, no Tennessean should have to live next door to a data center against their will so that petite dictators on city councils can grow the breadth of their miniature empires. Existing data centers and their power plants should also be required to buy out homeowners who are negatively impacted by them. Their profits are a lie to the extent that they are permitted to operate without fully internalizing the social costs involving energy, noise, and pollution that they impose on the rest of us. If data centers are truly wealth-creating, they will still be viable and profitable after they fully compensate everyone who is harmed by them. Otherwise, the profits and taxes are an illusion being subsidized by the misfortune of surrounding communities.



Financial Market Overview: AI Bubble or Balloon?

William J. Trainor Jr. Ph.D. CFA, Professor of Finance and
Director, Center for the Study of Finance

The start of the AI boom can be traced back to the launch of ChatGPT 3.5 on November 30, 2022, with the “big bang” for AI stocks beginning May 24, 2023, when Nvidia announced a surge in earnings beyond all expectations driven by AI demand (Lauricella and Schlueter, 2026).



There have been several surges and drawdowns and pronouncements that the AI bubble is bursting but so far, every drawdown has recovered with AI stocks reaching new highs. These new highs are almost always associated with one company or another announcing blowout revenues and earnings which spill over into every other AI related company. Currently, most of stocks examined below are down 25% or more from the previous 2-year highs with many of these highs reached in late June of 2026.

The balloon analogy is based on the idea that prices are inflated but reflect real underlying demand as industries that surround AI are broad based including everything from grid power, real estate, construction, plumbing, electrical, air conditioning, and even insurance, (See Eaglesham, 9/5/2026 for the insurance idea). Data Centers are currently in the news and “love them or hate them,” the value of some of these installations are in the tens of billions with up to 2,000 more planned for construction. On another front, Constellation Energy is working on restarting Three-Mile Island for Microsoft. This is not speculation and it shows demand related directly to AI is real and profitable, at least for those providing the infrastructure.

For investors, there are several ways to profit from this, some obvious while others, if you can think of it first, may lead to extraordinary gains. Below are tables that show stocks directly and indirectly profiting from the AI infrastructure build. For reference, the S&P 500 is up 62% over the period examined and is currently near an all time high.

Core Infrastructure: (data covers Jan. 2, 2024 to Sept. 14, 2026, S&P 500 up 62% over this period)

Ticker	Company	Jan 2, 2024 Price	Maximum Price during period	Sep 14, 2026 Price	% Change Jan24-Sep26	About
AMAT	Applied Materials	\$158.66	\$722.23	\$424.21	167.4%	Semiconductor equipment
AMD	Advanced Micro Devices	\$147.41	\$580.91	\$493.41	234.7%	GPU competitor to Nvidia
ASML	ASML Holding	\$740.23	\$1,986.87	\$1,575.15	112.8%	Monopoly lithography machines
AVGO	Broadcom	\$108.75	\$480.81	\$344.72	217.0%	Custom AI chips & networking
INTC	Intel	\$49.59	\$140.94	\$97.19	96.0%	Lost AI chip market share
KLAC	KLA Corporation	\$56.92	\$301.37	\$169.09	197.1%	Semiconductor process control
LRCX	Lam Research	\$76.69	\$433.33	\$273.49	256.6%	Semiconductor equipment
MU	Micron Technology	\$84.65	\$1,213.37	\$924.03	991.6%	AI memory chips (HBM)
NVDA	Nvidia	\$49.52	\$235.74	\$212.10	328.3%	World's first \$4T company Jul 2025
TSM	Taiwan Semiconductor	\$100.66	\$477.57	\$418.01	315.2%	Manufactures chips for Nvidia/Apple

The semiconductor industry has been on a tear not just recently, but for the last 10 years. MU especially benefited from the memory surge tied to Nvidia demand over the last two years.

Cloud and Software: (data covers Jan. 2, 2024 to Sept. 14, 2026, S&P 500 up 62% over this period)

Ticker	Company	Jan 2, 2024 Price	Maximum Price during period	Sep 14, 2026 Price	% Change Jan24-Sept. 26	About
MSFT	Microsoft	\$319.53	\$542.07	\$505.41	57.6%	OpenAI partner, Azure AI
GOOGL	Alphabet	\$130.44	\$402.62	\$349.39	161.2%	Gemini AI, Google Cloud
AMZN	Amazon	\$129.33	\$284.02	\$253.54	87.4%	AWS cloud AI infrastructure
META	Meta Platforms	\$295.73	\$790.00	\$665.60	122.1%	Llama open-source models
ORCL	Oracle	\$109.43	\$328.33	\$144.79	28.3%	Cloud AI infrastructure
CRM	Salesforce	\$208.61	\$367.87	\$259.43	21.8%	AI in enterprise software
NOW	ServiceNow	\$109.91	\$234.08	\$142.35	24.7%	AI workflow automation
PLTR	Palantir	\$14.00	\$207.18	\$173.31	1075.8%	AI data analytics
DELL	Dell Technologies	\$68.44	\$567.29	\$534.28	675.0%	AI servers & infrastructure

Four of the “Magnificent 7” stocks are in the group above and they have led the stock market for several years now.

Data Center Construction and Data Center Reits (sample of each only):

Ticker	Company	Jan 2, 2024 Price	Maximum Price during period	Sep 14, 2026 Price	% Change Jan. 24-Sept. 26	About
CAT	Caterpillar	\$272.03	\$1,064.90	\$784.00	181.4%	Building data centers
DE	John Deere	\$387.64	\$698.37	\$681.46	69.4%	Land clearing & construction
VRT	Vertiv Holdings	\$36.32	\$376.23	\$237.39	539.5%	Data center cooling & power
ETN	Eaton Corporation	\$210.65	\$459.96	\$393.16	83.2%	Electrical equipment for data centers
FIX	Comfort Systems USA	\$170.51	\$2,066.51	\$1,582.80	781.9%	Electrical & Plumbing
EQIX	Equinix	\$733.51	\$1,115.94	\$998.72	31.8%	Data center REIT
DLR	Digital Realty	\$123.94	\$203.91	\$179.14	39.9%	Data center REIT

Two of the stocks above (both data center REITS) have not outperformed the S&P 500. Other areas and stocks you may not have directly thought of are listed below:

Power & Utilities, Nuclear, Insurance (sample of each only):

Ticker	Company	Jan 2, 2024 Price	Maximum Price during period	Sep 14, 2026 price	% Change Jan. 24-Sept. 26	About
VST	Vistra Energy	\$33.45	\$217.92	\$140.73	335.2%	Power producer with nuclear assets
NRG	NRG Energy	\$37.89	\$184.03	\$108.47	178.9%	Power generation
CCJ	Cameco	\$38.77	\$134.09	\$93.26	136.1%	Largest Western uranium producer
UEC	Uranium Energy Corp	\$4.93	\$20.14	\$10.26	102.4%	Uranium producer
ALL	Allstate	\$113.85	\$275.11	\$258.48	126.4%	AI in underwriting & claims

There are other major areas such as Networking (CSCO & ANET), and I have ignored the AI companies themselves that will soon be making their initial public offerings. If you can think of another angle, someone else probably already has as well. The best way to find out is to look at the stock price over the last couple of years, if it has already soared, you have been beaten to it.

However, all this may still be in the early stages and most of these stocks are currently in a drawdown or depending on your point of view, their stock bubble burst. Long-term, I wouldn't underestimate the changes AI will bring both to our lives and stock portfolios, assuming it doesn't take over...

References:

Eaglesham, Jean. Sept. 5, 2026. The next big opportunity in data centers; Insuring them. Wall Street Journal, <https://www.wsj.com/finance/the-next-big-opportunity-in-data-centers-insuring-them-f133f3cb>
Lauricella, T. & Schlueter, R. May 22, 2026. 3 years of the AI stock market boom in charts. Morningstar, <https://www.morningstar.com/stocks/three-years-ai-stock-market-boom-charts>



What Data Centers Actually Sell

Richard P. Gregory, Ph.D. • Professor of Finance

Most arguments about data centers are arguments about physical things. Megawatts. Cooling water. Acres of farmland. How long it takes to get a transformer. Whether the local grid can handle one more substation. These are real limits, and they deserve the attention they get. But they all describe what goes into the building. They say very little about what comes out.

What comes out is computation, which is not a commodity the way electricity is. Computation is different, because it is an input to decisions. Its value to a buyer is not the processor time itself. It is the improvement in whatever decision that processor time supports. Once you take that seriously, the market for computation turns out to have a flaw that never shows up on the power bill. And building more data centers will not fix it.

Start with what a buyer is really purchasing. A shipping company renting solver time to plan tomorrow's delivery routes is not buying processor-hours. It is buying a route plan that misses fewer delivery windows. The processor time matters only because it moves that outcome. But the link between the two is not simple. It is not a straight line, and it is not the same for every buyer. It depends on the structure of the problem being solved — on how hard it is to get the next bit of quality out of the next bit of effort. That is, the problems firms and people face can be unique, but the computational resources to address those problems may not be. Some problems improve fast and then flatten out. Others resist for a long time and then crack open all at once. Some never get solved adequately. The buyer's demand for computing power runs through this link, and the link is a technical fact about the problem, not a matter of taste.

Now look at what goes into the contract. Almost none of that link is written down. What gets written down instead is a guarantee. The vendor promises a response time, or an error tolerance, or a deadline the job will finish by. Guarantees are cautious by design. They have to be. A guarantee that gets tested is a guarantee that failed. Everyone signing these contracts knows that the promised level and the typical result are two different things, and that the distance between them is usually large. That distance, by itself, is not the problem. It is understood, discounted, and priced in.

The problem is one step past that. When a firm buys another block of capacity, two things improve at once. The contractual guarantee gets tighter. And the firm's actual results get better. Those two improvements happen at different speeds. More importantly, the relationship between those speeds is not the same for every buyer. For one firm, extra capacity mostly buys a tighter promise, because its real performance was already near its ceiling. For another firm, the promise barely moves, but the actual routing quality improves a lot. Both firms see the same posted price. Both do the math correctly. And the resulting allocation of capacity is still wrong.

It is wrong because one price is being asked to do two jobs it cannot do at the same time. Buyers weigh the price against the value of the thing they can write into a contract, which is the guarantee. But what society gets depends on what actually happens, which is the result. If the ratio between those two were the same for everyone, the difference would wash out and nothing would be lost. It is not the same for everyone. So computational capacity flows toward the firms whose promises improve fastest, not toward the firms whose outcomes improve fastest. Nobody cheated. There is no monopoly here, no hidden information in the usual sense, no pollution. The market clears. The outcome is still not the one a benevolent planner with the same power and the same chips would have chosen.

There is one more layer. A buyer usually cannot measure how far off their own answer was. To know how far you missed, you would need the right answer, and if you had that you would not have needed the computational resources. Faced with that blind spot, a sensible firm buys against the worst case it can point to. It is protecting itself against an error it cannot size. That is reasonable behavior. But it means a large share of the demand for computing power is really demand for insurance against unknown error, not demand for accuracy. Billions of dollars of concrete and copper are being committed against demand curves shaped by what buyers cannot see about their own work.

This should slow down the reflex that says the answer to any strain in this market is to build more. More capacity loosens a limit on quantity. It does nothing about a mismatch in rates. A bigger fleet handed out by the same price signal repeats the same misallocation on a larger scale, and now that misallocation is poured into twenty-year assets that are expensive to move. It is also worth saying that this problem does not sit inside the more familiar ones. Even if every standard distortion in this market were fixed tomorrow, the differences between firms in how promises track results would still leave capacity in the wrong hands. Two separate problems. Two separate tools needed.

What would actually help is writing contracts on rates of accuracy improvement instead of levels of accuracy. The instrument has to target how fast quality improves per unit of capacity, not how much slack sits inside a promise. That sounds impossible, since the error itself cannot be seen. It is less hopeless than it looks. The leftover quantities a solver produces while it works — the numbers it is trying to drive toward zero — can be measured. Under reasonable assumptions, they fence in the error from both sides. A contract written on those measurements gets closer to pricing the right thing than any contract written on a certified tolerance ever will.

One last point about who wins. The buyers who do best under the current setup are the ones with enough instruments to know how they are really performing: large, sophisticated operations with in-house expertise and good monitoring. Everyone else buys the promise. As computing becomes a bigger share of costs in industries that have nothing to do with technology, that advantage compounds quietly. It will not appear in any tally of the megawatts.

The scarce resource is not just the data center. It never was.



Can AI Beat the Simple Equal-Weighted Investment Strategy?

Son Wilson, Ph.D. • Assistant Professor of Finance

In 1952, Harry Markowitz revolutionized portfolio management by introducing Mean-Variance Optimization (MVO), offering a mathematically rigorous framework for asset allocation. Yet, decades later, financial economics still confronts the most enduring paradox first highlighted by DeMiguel, Garlappi, and Uppal (2009): sophisticated mathematical investment strategies rarely outperform the simplest allocation heuristic in out-of-sample testing. This equal-weighted strategy, known as naive diversification, or the 1/N rule, simply splits capital equally across N available assets.

The secret behind the 1/N rule's persistent success lies in parameter uncertainty. Traditional optimization algorithms rely heavily on historical estimates of expected returns and variance-covariance matrices. Because financial return data is inherently noisy, slight estimation errors in mean return vectors lead optimizers to assign extreme, unstable asset weights, generating excessive portfolio rebalancing and high trading turnover, which rapidly erode net returns once transaction costs are deducted. By imposing equal weighting without parameter estimation or return forecasting, the naive diversification avoids parameter error, overfitting, and excessive trading.

In a recent study conducted with my co-authors, we set out to revisit this foundational puzzle to ask whether modern artificial intelligence (AI) techniques, specifically machine learning (ML) and large language models (LLMs), could finally overcome these classical estimation hurdles and dethrone the equal-weighted benchmark. Evaluating 35 years of monthly U.S. market data spanning from January 1990 through December 2024 across ten industry portfolios, we compared 1/N against eight distinct allocation strategies. These included traditional strategies (such as the classical MVO portfolio, minimum variance portfolio, and Bayes-Stein portfolio) and AI-based strategies, including MVO strategies with ML techniques of elastic net regularization and moving average smoothing, the Hierarchical Risk Parity framework, and MVO strategies with LLM approaches (i.e., LLM-MVO and smoothed LLM-MVO). For the LLM-driven strategies, we prompted Claude 4.6 Sonnet to evaluate sector return prospects using five key macroeconomic indicators: real GDP growth (with a conservative four-month publication lag), CPI inflation, the unemployment rate, the federal funds rate, and three-month Treasury yields.

Our results revealed that, among the classical strategies and LM-based portfolios, the simple 1/N strategy was hard to beat, with the highest annualized Sharpe ratio (0.62) and extremely low monthly turnover (2.52%). However, LLM-powered strategies appeared to achieve a major breakthrough, with Sharpe ratios near 1.00 (0.96 for unconstrained LLM-MVO and 0.97 for smoothed LLM-MVO), dramatically exceeding 1/N's Sharpe ratio. Circular block bootstrap statistical tests confirmed that this gross outperformance was statistically significant at the 1 percent level. However, this apparent AI victory evaporated after we factored in two realistic conditions: implementation costs and training data contamination.

First, LLM strategies required immense portfolio rebalancing, generating monthly turnover rates between 79% and 113%. When transaction costs were deducted, net risk-adjusted returns plummeted. At a standard institutional transaction cost of 50 basis points per trade, the net Sharpe ratios for LLM-MVO and smoothed LLM-MVO dropped to, respectively, 0.57 and 0.68, comparable to 1/N's highly cost-efficient net Sharpe ratio of 0.61.

Second, our research uncovered severe training data contamination within the LLM strategies. Our sub-period analysis showed that the significant gross outperformance of the LLM-driven portfolios was confined entirely to the 2010–2024 era covered by their training corpus, while the models did not perform statistically better than 1/N during the 1990s and the 2000s, signaling a form of data contamination. Because commercial LLMs are pre-trained on vast archives of historical financial commentary and news, providing calendar dates in prompts allowed LLM models to "memorize" past financial crises and market cycles rather than performing genuine forward-looking analysis. To isolate true macroeconomic reasoning, we executed a date-masking test by withholding formation dates and prompting the model solely with raw macroeconomic numbers. Without temporal dates, the Sharpe ratios of the LLM-powered portfolios fell by approximately 30% to 0.65 for LLM-MVO and 0.72 for smoothed LLM-MVO, rendering their performance advantage over 1/N statistically insignificant. Their net Sharpe ratios at 30 basis points also decreased sharply to 0.41 and 0.54, well below 1/N's 0.61.

Ultimately, our study demonstrates the persistence of simplicity: naive diversification remains the dominant benchmark in the AI era. For investment professionals evaluating LLM models, our work cautions that back-tested AI gains must undergo rigorous date-masking contamination checks before being attributed to genuine predictive capability.

REFERENCES

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. "Optimal versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?" *Review of Financial Studies* 22 (5): 1915-1953.

Markowitz, Harry. 1952. "Portfolio Selection." *Journal of Finance* 7 (1): 77-91.



Small Business Spotlight: Velvet & Moss

Joseph Newhard, Ph.D. • Associate Professor of Economics

Velvet & Moss Vintage Emporium

817 W. Walnut St. #9, JC

11:00 – 6:00 Tues to Sat

Those of us with an old soul and refined taste are looking for something beyond the typical thrift store discovery. Velvet & Moss Vintage Emporium in Johnson City is not a thrift store but rather a carefully curated collection of antiques, vintage clothing, and contemporary offerings with a timeless look. Opening last March near Italian Pizza Pub, the aesthetic ranges from art deco to the space age, from Biba to Boho.

Velvet & Moss is stocked floor to ceiling with dresses, shirts, jeans, and jackets. In Men’s Clothing I saw labels like Ralph Lauren and Woolrich. You can explore the collection of hats, coats, shoes, boots, sunglasses, handbags, jewelry, costumes, gloves, scarves, ties, belts, vests, vintage perfume bottles, luggage, vinyl records, posters, decorative items, and other oddities. One shelf displays a transistor radio; another displays a working uranium lamp. There is a shelf of old film cameras, a tennis racket, and a baseball glove. A Rolling Stones jacket hangs nearby. As you shop you’ll hear jazz, big band, bossa nova, and rock & roll: Glenn Miller, Oscar Peterson, the British Invasion.

Velvet & Moss owner Julie got her start as a young girl hanging out in her parents’ antique shop in Southern California. Inspired by Diana Rigg’s outfits on *The Avengers*, she used tracing paper to design dresses over *Betty and Veronica* comics. As a student of history, her knowledge of old films and music surpassed my own. Decades of collecting now culminate under one roof just a short walk from Backdoor Records, another shop for those of us who miss the analog world.

Shops like Velvet & Moss are as close as some of us will ever get to visiting Carnaby Street in Swinging London. A step into the past, it is a shopping destination for those of us who know that fashion and art and music were just better in the old days. If names like Twiggy, Jobim, Mary Quant, or Burt Bacharach resonate with you, you should check out Velvet and Moss which offers the kind of inventory that makes shopping a treasure hunt.

