1st Quarter Commentary

April 2025

Anyone who's read or listened to us recently—or any time in the past several years—pretty much knows the throughline: an extended inflationary period is legitimately in the cards, borne on long-building pressures that aren't going away. There's no need to walk that thoroughly plowed ground today—aside from noting that today is a bit different, because some inflation elements have begun to show themselves and catch at investors' attention. So this will be about how our portfolios are positioned for that scenario.

For many years, the great river of investment funds flowed in a certain direction, but not one we liked, so we waited. We were able to wait in companies with particular business models that afforded excellent financial compounding characteristics even in the absence of favorable conditions. Lately, the river is flowing back in our direction. Intelligent inactivity is a strategic practice we wrote about early on, and is something we try to do as often as possible – if that logic is possible.

And as diametrically unalike as our portfolios are to the stock market indexes and the extreme Information Technology weightings, they nevertheless encompass meaningful and salubrious exposure to the suddenly rampant new technology investment cycle: the explosive spending by every major IT company on data centers amid the mass adoption of artificial intelligence. Yet, without any significant changes on our part because, in truth, we didn't have to do much. AI's pertinence to our portfolios bears some explanation, though, so this first section will try to describe some of the reasons for the inevitability of the continued spending.

Section I: A.I. and Data Center Demand

Prefatory Remarks about AI and Data Centers

One hears about the enormous sums being spent on data centers. Meta recently said it will spend $60 to $65 billion on data center construction this year, more than 50% above 2024. The Boston Consulting Group estimates that the leading data center companies will spend $1.8 trillion in the U.S. between 2024 and 2030.[1] An astounding figure, but it's tough to judge how large it really is without a reference point.

The greatest relevant reference point must surely be the country's spending for World War II. For the human facet of its scope, the single best statistic might be this: in 1945, 40% of the male population in America, from ages 18 to 45, was in the military. An even greater number had served, but not all were still serving by war's end. Astounding, is it not?

In fiscal scope, between 1941 and 1945, the country spent $296 billion in then-current dollars on the war effort.[2] This amounted to 31% of 1940 GDP. Adjusted to reflect the impact of inflation since 1945, that's $5.1 trillion of spending today. Which means that the aforementioned data center capital expenditure projection is of the same order of magnitude as the cost of World War II, which involved a full, government-mandated mobilization of the economy, private factories and all. Alternatively, in 2024, the entire U.S. defense budget was $842 billion. These are the magnitudes. The build-out of extremely large data centers has the potential to be the greatest deployment of private investment capital in history.

---

[1] *https://www.bcg.com/publications/2025/breaking-barriers-data-center-growth*

[2] *Stephen Daggett, Cost of Major U.S. Wars: Congressional Research Service Report for Congress RS22926 (Foreign Affairs, Defense, and Trade Division, Congressional Research Service, The Library of Congress), 2008*

First Question

Not to ignore an obvious question, even if it seems self-evident, but what do the current generation of data centers do and why do they cost so much? As enormous buildings that house massive numbers of computer servers in vast ranks of racks, their most observable physical feature is the ungodly amounts of steel, concrete, copper, high-purity silicon and other resources that go into them. They also consume ungodly amounts of electricity. In fact, that is their defining feature, since data center size is quoted not in square feet or other real estate terminology, but in how much total power their processors draw.

How much power? An anecdotal example, from Meta's recent announcement, is that by the end of this year it will own 1.3 million processing units. Each NVIDIA-chip unit draws 700 watts, which is not much less than the 1,200-watt electricity usage of an average household.[3] (Households also contain a 700-watt appliance—the microwave—but that runs for a minute or two at a time, not 24 hours a day. Still, one can see how big a draw a microwave is on a home's capacity.) At a 100% utilization rate, Meta's 1.3 million chips would draw power at the rate of 7,972,000 megawatt-hours in a year, about the same as the state of Rhode Island, a touch less than Hawaii.[4]

A national scale picture is provided by the December 2024 report to Congress by the Lawrence Berkeley National Laboratory, which estimates the growth in U.S. data center electricity use from 2023 to 2028.[5]

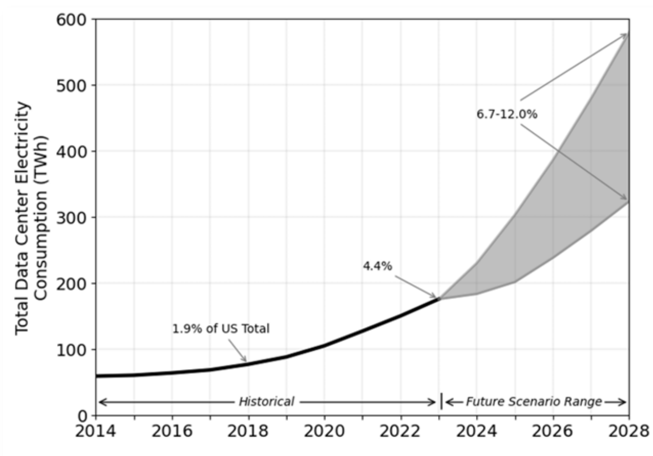

Figure ES-1. Total U.S. data center electricity use from 2014 through 2028.

Any number of investment firms and consultancies have made such projections, but Berkeley Lab has the property of not being a commercial entity, and of being staffed by scientists, with all the attendant thoroughness one might expect. The study even includes a survey of other studies, the aforementioned Boston Consulting Group paper among them. Within the Berkeley Lab paper is a sub-review that sorts those other studies by their methodologies: whether they employ a bottom-up approach, top-down, or an extrapolation of existing trends, and it reviews the advantages and limitations of each methodology. Berkeley even reviews the limitations of its own methodologies.

Berkeley Lab's conclusory projections, shown in the accompanying chart, makes allowances for everything from the range of NVIDIA chip models shipped per year, to the server types and server utilization rates. There is even attention paid to the energy used for cooling the data centers, with the latter being dependent upon the types of systems employed, such as evaporative cooling, and the weather conditions of regions with data center clusters. The point is that—however valid or invalid those projections turn out to be (the

[3] https://www.eia.gov/energyexplained/use-of-energy/electricity-use-in-homes.php

[4] https://www.eia.gov/electricity/state/

[5] https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf

history of new-technology growth estimates is rife with miscalculation)—there is much training-driven conservatism in these projections.

The low end of the range for data center electricity needs—*as a proportion of total U.S. consumption,* which in any case is already rising—is a 50% increase between 2023 and 2028, from 4.4% to 6.7%. The high end of the range is a near tripling, from 4.4% to 12.0% of U.S. demand. Given that total U.S. electric power production has been essentially flat for a decade, and given the aged and undermaintained U.S. electric grid, the question of how that can be achieved is another discussion.

Pre-First Question

The next question—perhaps it should have been the first—is ***why*** does AI require so much power?

What we don't know about technology basics...

The Complete Peanuts: 1957 to 1958 by Charles M. Schulz (Fantagraphics Books, 2005.)
Originally published July 9, 1957.

Ignorance Transparency Disclosure: The editor claims no education, training or expertise in computer programming, computational biology, neurology, biochemistry or other science disciplines referred to herein, and considers this disclosure sufficient to exculpate any related misstatements or other displays of ignorance. Editor merely endeavors, like many of us, to read—perchance to grasp—some basic principles as they might pertain to informed investing.

The Electric Power Needs of Large Language Models— the AI we Know and Love

For a ChatGPT query, the power requirements are really about the amount of data that must sorted to provide a response—which is the quantity of data that exists on the internet. It's as if, for an open-book essay test in a high-school History class, a particular student is unable to synthesize or draw a conclusion from the book assigned. Instead, the student has the unusual ability to teleport to the library, then super-fast read and verbatim-recall every word of every book and periodical in it, and from that select and piece together the 1,000-odd words for the standard five-paragraph essay.

But even that is more high-functioning than ChatGPT. Because piecing together to ChatGPT doesn't mean what it means to a person. ChatGPT can't understand context on its own: "The king was tyrannical and ruled the populace with an iron fist" is no more intrinsic a choice than "The king was tyrannical and tickled the populace with ice cream." ChatGPT uses statistical probability to estimate, from all the instances in the billions of indexed internet web pages, that after the words "the king was," which were themselves pieced together this way, the word "tyrannical," has a certain statistical likelihood of coming next in the sequence. Along with "benevolent," "indifferent," and a long list of other adjectives. This happens for each word or even part of a word.

There are ingenious ordering and selection methods to determine which next word to use, inspired by the way that neurons in the brain promote or suppress electrical input signals that they might pass on to other neurons. For each word chosen, hundreds of billions of calculations take place. Full responses involve trillions of calculations, and that is largely what the electric power is for.

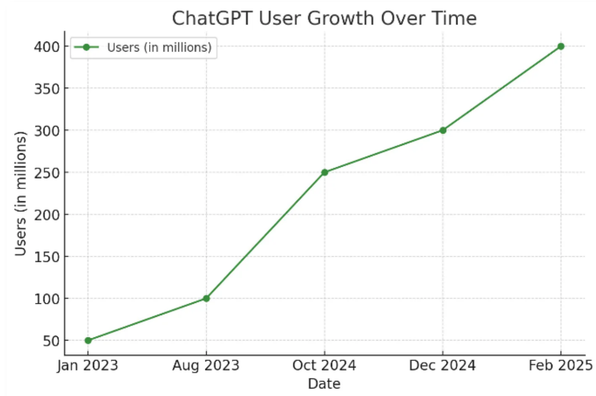

ChatGPT, which is one Large Language Model among many, is reported to now process over 1 billion queries daily, still only a modest fraction of the number of daily Google queries. The number of ChatGPT queries will continue to increase: in the two months from year-end to February, the number of active weekly users globally rose by one-third.[6]

Power usage is also a function of the magnitude of data available to sift through. The volume of data created, used or copied in the world in 2025 is expected to exceed 180 zettabytes. What's a zettabyte? As a number, it's a 1 followed by 21 zeros. In visual media terms, it's been cited as 250 quadrillion photos or 20 billion years of streaming Netflix video. In vacation terms, it's been cited to be as many grains of sand as on all the world's beaches. That's ***one*** zettabyte. The 180 zettabytes is almost triple the amount in 2020.

Only a small proportion of this is stored. In 2020, the installed base of storage capacity was 6.7 zettabytes, but was expected to increase at nearly a 20% annual rate for the five years up to 2025.[7] Those statistics are from IDC and Statista. A different firm, Cybersecurity Ventures, estimates that storage will be a far, far greater 100 zettabytes this year. The difference between the two figures is based on a higher estimate for the *proportion* of new data that is permanently stored, which the latter firm now puts at 50%, up from 25% in 2015.[8] As large as those differences seem, they are trivial for our purposes, as will be shown shortly.

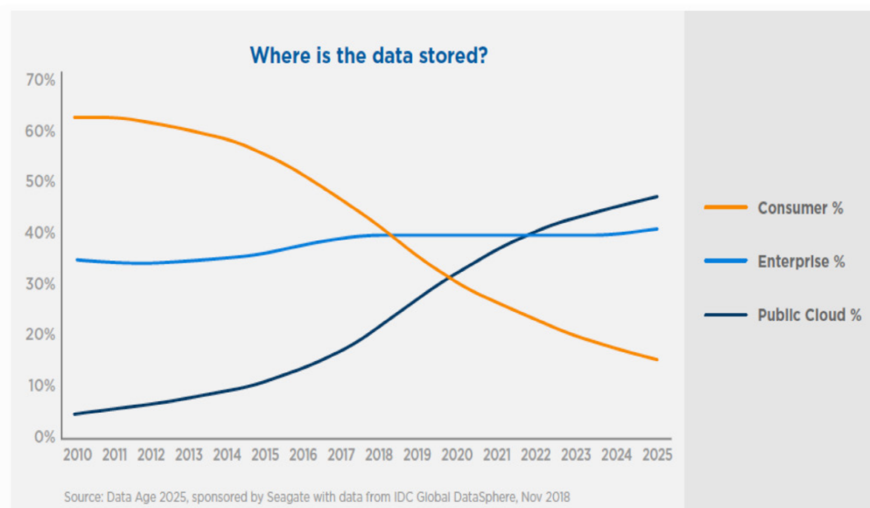

IDC on where data will be stored in the Datasphere
IDC REPORT FIGURE

[6] *https://nerdynav.com/chatgpt-statistics/*
[7] *Source: Petroc Taylor, Statista, 11/21/24.*
[8] *https://cybersecurityventures.com/the-world-will-store-200-zettabytes-of-data-by-2025/*

An important point is that the stored data ***also*** requires power: the servers must always be on. The IDC/Statista figure for annual data creation and usage is expected to more than double between 2025 and 2028. Less and less is being retained on home or office computers; even ordinary smart-phone-carrying civilians have tens of thousands of photos and videos in their archives. This is all emigrating to cloud storage, which means data centers.

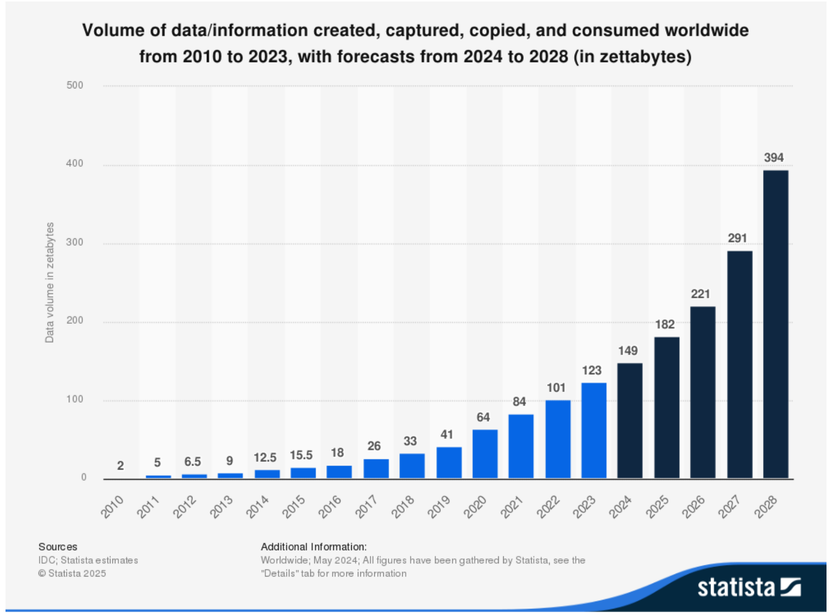

The explosion in energy consumption has everything to do with AI, because ChatGPT-type searches are variously estimated as requiring up to 10x the energy of a Google search. Image generation via AI has been estimated to require 5x more electric power than that. And video generation via the SORA text-to-video AI model released by OpenAI only this past December has been calculated to use 1,000x more energy than a ChatGPT search.

Pre-pre First Question

Even before the question about why AI requires so much electric power, perhaps it should have been asked why the world even needs it? Or needs it to a degree that justifies hundreds of billions of dollars of seemingly overnight capital raising and investment?

The only AI the public has interacted with is the Large Language Model type, like ChatGPT. It's easy to find lists of very large markets for it.

Among them: voice search and fake review detection (for e-commerce); proofreading and writing (for legal contracts, or for pretend investment research on financial news websites); data, intelligence, and cyber threat analysis and response (for commercial and military users); facial recognition, and gait and body language analysis (for surveillance and law enforcement); and back-office automation and portfolio analysis (for the banking sector). Other applications for AI include image and video generation, GPS navigation and autonomous driving, and robotics and industrial automation.

A comprehensive list would obviously encompass every sector and job description in the economy that lends itself to algorithmic processing—and to reworking patterns and procedures that can be committed to a database. Any one of those applications is a candidate for saving production time, improving success rates, or some other attribute of economic value to some entity. They are clearly very large dollar-value numbers, but perhaps difficult to quantify, and surely not every member of the workforce is thrilled about an AI substitute for their job

The Other AI—Large Math Models (or Large Quant Models, which some prefer)

But Large Language Models are not the only form of AI. There's another, with which the public is almost wholly unfamiliar, Large Math Models (or Large Quant Models, which some prefer). LMMs use the same GPU chips as the Large Language Models, but in a very different way.

Large Language Models are trained on external data; they can sort and pattern and rearrange that data, but they can't create their own independent data. Because of that, as much as they *can* do, there are important classes of tasks for which they are *not* effective, but a Large Math Models is. One such is new drug development.

A Sidebar on HI (Human Intelligence)

Why is generative AI so energy costly? Because it's not actually *intelligence*. It doesn't "know" things, but has to go outside itself to the entire library of human written and visual data to piece responses together through trillions of calculations. While there are things it has already learned and can re-access with less processing, global accumulated data always increases and must still be accessed.

Human intelligence is also energy-intensive in one very narrow sense: the brain comprises only about 2% of our body weight, but consumes about 20% of our resting energy, so that's 10x relative demand.

On the other hand, that 20% of an average resting metabolic rate of 1,300 Calories per day—before physical activity—translates to about 12 watts,. So however much memory the brain can hold and access, whatever the quantity or quality of output it can produce, inclusive of all the bodily functions that it regulates, it uses lot less energy than a refrigerator light bulb.

This means that the 170 million minds in the U.S. civilian workforce, whatever their aggregate creative and applied intellectual output might be, consume about 2,040 megawatts in total. That is 0.3% of the aggregate power capacity of all electric utilities in the U.S.[a] It would rank below the 50th-largest power plant in the U.S.[b]

Curiously, intense thinking or studying—such as for a exam or perhaps a quarterly investment review webinar—doesn't appear to involve any increase in brain energy consumption.

[a] https://www.eia.gov/electricity/annual/html/epa_04_02_a.html
[b] https://www.eia.gov/energyexplained/electricity/electricity-in-the-us-top-10.php

The Growing Problem—and the Emerging Solution—for New Drug Development

Not only is new drug discovery an enormous and growing market, it is primed for massive investment in LMM AI. To characterize the market opportunity in basic business terms, there is intense inherent demand both from sellers (the drug companies) and from buyers (everyone with a drug-treatable malady) for faster-to-market and less egregiously expensive treatments. And perhaps especially for the thousands of poorly treatable or as-yet untreatable conditions. And to make a more powerful triad, even demand from a government social policy and fiscal perspective.

New drug development expense has long been rising on the order of 13% annually, multiples of the general price index.[9] One factor behind this is the challenging process to discover novel molecules that can actually have therapeutic effects.

[9] *A Comprehensive Review of Discovery and Development of drugs discovered from 2020-2022* *https://www.researchgate.net/publication/376382939*

Once a new compound reaches the regulatory evaluation phase, approval time is up to 14 years.[10] Most of the failures occur in final Phase II and III clinical trials, which are designed to assess both efficacy and toxicity. These are the most expensive part of the process, because this is where randomized double-blind testing takes place and, so, where so much of the drug development costs are lost.

Because the third major element of drug development expense is that even once a molecule is promising enough to reach Phase I clinical trials, the overall failure rate for FDA approval is over 90%[11] (this from a 2020 study across 12,700 new compounds). For certain broad condition sets like oncology and neurology, the failure rates are about 95%.

The R&D cost of each of those failures, from a 2020 study, averaged over $2 billion and as high as $4.5 billion.[12] They might be higher by now. The price of a new drug must incorporate the aggregate cost of the nine out of 10 that failed. Otherwise, the expense would overwhelm the profits of the successful ones and threaten the future of the entire pharmaceutical sector.

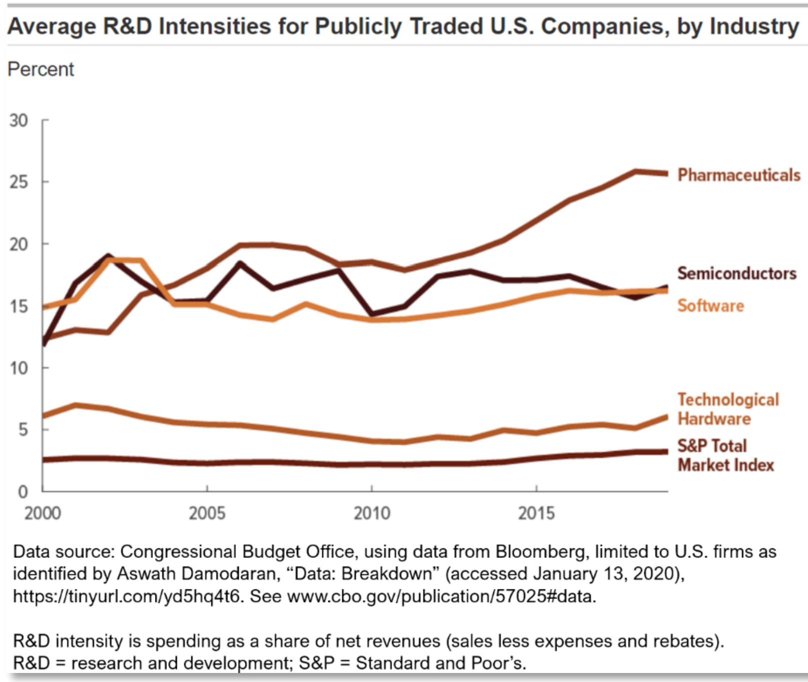

[10] https://pubmed.ncbi.nlm.nih.gov/38204591/
[11] https://go.bio.org/rs/490-EHZ-999/images/ClinicalDevelopmentSuccessRates2011_2020.pdf
[12] https://pmc.ncbi.nlm.nih.gov/articles/PMC8516790/

The financial motive to reduce the cost and time to market can be measured. Global pharmaceutical sales in 2024 were $1.5 trillion, over half of that coming from the U.S. In the past few decades, R&D spending, now at $306 billion, has claimed an ever-greater portion of drug company sales. It is no coincidence that investment returns on R&D have been declining.

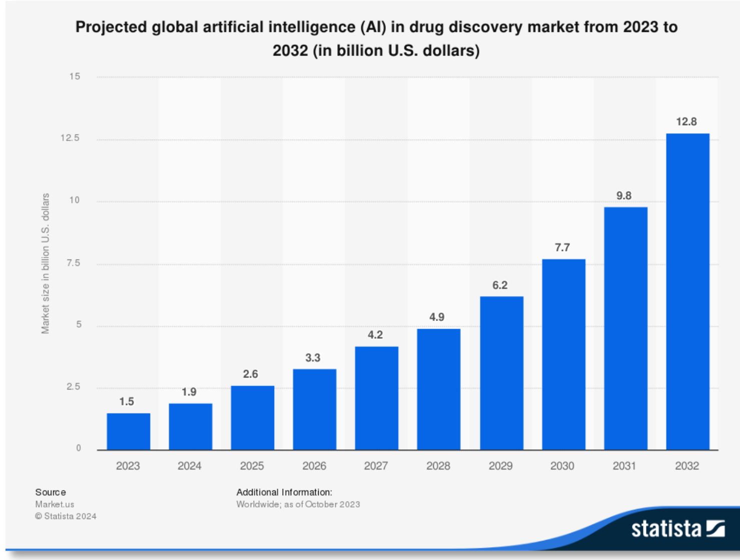

Drug discovery spending within R&D is expected to be $71 billion this year. In 2017, this was projected to grow at an 8.1% rate for the nine years to 2025.[13] As of October 2023, a nine-year projection for spending on AI *within* the drug discovery segment was for 27% annual growth. One will recall that both NVIDIA's H100 chip and ChatGPT were released less than a year earlier, in late 2022.

Large Math Model AI for New Drug Development, More Specifically

AI can save enormous amounts of time and money for new drug discovery. A great challenge has been to identify and develop molecules or proteins that bind to or fit into very specific locations of a certain target protein or gene, often in very complex structures, to activate or deactivate them. This is what drugs do.

This entails understanding how proteins change, or fold, their very elaborate shapes from an initial state into their final configuration. Just about every bodily function is based upon or regulated by proteins. Proteins can only perform their biological functions once they have folded. Diseases like Alzheimer's and cystic fibrosis are related to mis-folding proteins, so this is an area of intense study. Merely identifying a therapeutically suitable and specifically constructed protein, though, is hideously complex.


ANDROLOGY

REVIEW ARTICLE | Open Access

Folding, misfolding, and regulation of intracellular traffic of G protein-coupled receptors involved in the hypothalamic–pituitary–gonadal axis

Alfredo Ulloa-Aguirre, Ross C. Anderson, Teresa Zariñán, Rubén Gutiérrez-Sagal, Eduardo Jardón-Valadez, Claire L. Newton

First published: 07 March 2025 | https://doi.org/10.1111/andr.70018

For instance, the number of shapes that even a modest-sized protein can take, relative to others near it, actually exceeds the zetta level of 10^{21} that was discussed earlier about the zettabyte volumes of internet data.[14] The below ball-and-stick graphics below show two of practically innumerable theoretical three-dimensional configurations a small molecule might take relative to its receptor protein.[15] It's seen as the changed pattern of connected atoms relative to the fixed-background pattern of the target, only a small

---

[13] BIS Research, Sept. 28, 2017, in Statista https://www.statista.com/statistics/765535/drug-discovery-market-worldwide-by-segment-globally/
[14] https://magazine.hms.harvard.edu/articles/did-ai-solve-protein-folding-problem
[15] https://www.cheminformania.com/never-use-re-docking-for-estimation-of-docking-accuracy/

portion of which is shown. The factors that impact how and why a specific protein folds with respect to a target occur at the quantum mechanical level, where the outermost electron shell or wave function of atoms dictates how bonds might form.

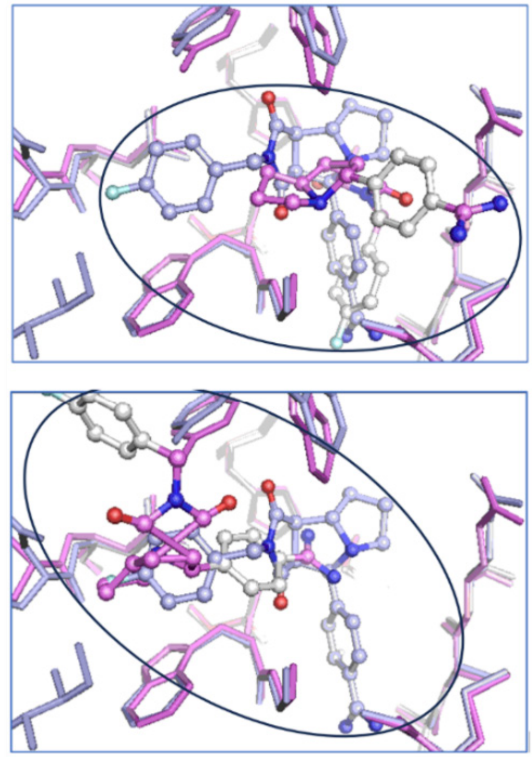

Two configurations out of more than billions?

Generative AI-based efforts, beginning many years before the NVIDIA H100 chip, have greatly aided research by identifying hundreds of millions of potential protein structures. That allowed researchers to model a compound "in silico," as they say, without actually having to manually synthesize some limited number of variations of it. However, this hasn't resulted in much progress in disease treatment. Although many more molecules are being brought to, and succeeding at, the pre-clinical Phase I level—where testing is done on a small group of healthy people to try to determine the maximum safe dose—the failure rates at Phases II and III aren't much different.

That's because of the intrinsic limitation of the Large Language Model: its pattern recognition capabilities are dependent upon on huge volumes of *existing* data. But what if there is no data? As in, how will a new molecule interact with a specific pathogen or pancreatic cancer cell protein? How will the immune system or the cancer cell respond? A Large Language Model can't be asked to create new data—or if it is asked, it might "hallucinate" it. But that is what Large Math Models can do.

LMMs aren't—as IT companies like to refer to them—artificial "intelligence" any more than Large Language Models are. People in medical research tend to use more accurate descriptors: high performance computing, machine learning, or computational biology. LMMs use the same NVIDIA GPU chip as ChatGPT, but rather than searching an external database, they create their own data through highly iterative modelling.

This can start by making a digital copy of an existing molecule and a virtual representation of a particular receptor or tissue, and running simulations of their interactions. The LMM makes incremental changes in response to results or non-results, then runs more simulations, creating new data all along the way. When a pharmaceutically promising structure is identified, biochemists can synthesize and test it. Then it can be refined and re-modeled again, and so forth.

In fact, because the data sets generated are so massive, generative AI models are necessary in order to discern patterns in the data. They are trained on the LMM-created data to accelerate the process and make better inferences about toxicity or drug intervention outcomes. This appears to be able to meaningfully improve success rates at the Phase II and III clinical trial level.

For a sense of the massive computing power that must be brought to bear, It should be understood just how iterative and data-intensive this is. For instance, the data to be simulated must incorporate four dimensions: not just the three dimensional structure of all the proteins involved in the interaction, but also their changing shapes over time as they fold. The shapes themselves can be extremely complex and can contain thousands of atoms in different groupings, each with their own structures and folding dynamics.

As astounding as it might seem, the time it takes a protein to reach its unique folded state can be as little as 4 millionths of a second—a few hundred thousand times per second. Enzymes, which catalyze reactions in other molecules, fold even faster. How much data is collected, nano-second by nano-second, if one person were to be scanned for just a few hours? Trillions upon trillions of simulations are made, and using huge amounts of data.

The Permanent Data Storage Load: An Example

There is no doubt that the emerging use cases for AI will cause storage needs to climb drastically. Staying with intelligence and neurology:

- Just this month, a report was released by The MICrONS Project. This is a project of IARPA, the Intelligence Advanced Research Projects Activity program, of the Office of the Director of National Intelligence. IARPA has an interest in improving pattern recognition systems. One effort toward this is reverse-engineering the way the brain is able to learn patterns from very modest amounts of information, and to thereafter recognize them even when significantly distorted or degraded.

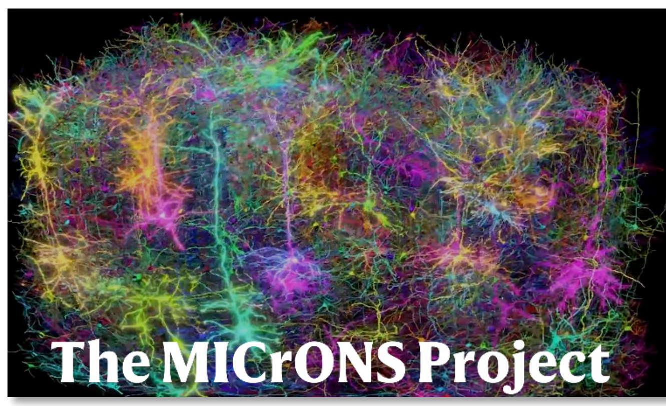

- MICrONS achieved an apparent first: the complete and fully detailed visual and digital reconstruction of just less than a cubic millimeter of the visual cortex of a mouse's brain. For reference, the thickness of a dime is 1.35 millimeters.
- That cubic millimeter included more than 200,000 cells, 75,000 neurons, and 523 million synaptic connections amongst those neurons.
- There are many other fascinating things to say about this, but pertinent to this discussion is that the raw dataset for this project amounted to 2 petabytes. Numerically, that's a 1 followed by 15 zeros. In experiential terms, it has been put, in ascending data density order, at 500 billion pages of standard printed text, 2,000 years' worth of MP3 audio, or 13 years of high-definition video content.

  Just for this single, 1-cubic-millimeter study. Which is not, in any case, over; it's just the first step.

Summary

This high-performance computing section was all a way of supporting, with more accessible and relatable detail, the idea that the growth in data centers and their resource requirements will continue, because it is economically based demand:

- *The early evidence is that high-performance computing applied to drug discovery is about an 18-month process rather than the norm of five to 10 years.*
- *It also appears to substantially "de-risk" new drugs before they get to the extremely high-failure-rate, high-expense clinical trial stage, resulting in much higher FDA approval rates. Both the discovery and the testing phase should experience dramatic benefits in terms of cost and years-to-market.*
- LMMs require gigawatt-scale data centers, too. Will people want the data center economy and the implications of all the energy consumption? Some of the purposes to which the electric power for LLMs is put, say for surveillance or replacing large swathes of white collar jobs, could be socially controversial. But will people say no to a technology that promises to vastly improve drug discovery and cure previously untreated diseases? No, they will say yes.
- Moreover, rather than costing more—which is the norm—high performance computing promises to dramatically reduce the cost of billion-dollar-class drugs. Right there is the powerful economic dynamic of buyer demand meeting seller demand for the very same product.
- Also, drug discovery is only one application of high-performance computing. The same iterative modelling methods for creating new compounds is being used for chemicals (such as for improved batteries) and materials science and engineering (lighter yet stronger composites for, say, jet wings or wind turbine blades). The various forms and applications appear to be transformative.

Section II: A Fork in the Data Center Road—Two Antipodean Approaches to Portfolio Positioning

Let's say you've concluded that high performance computing and data center construction of one-for-the-history-books scope are here to stay and on the way. It will impact growth, profits, and valuations in of all sorts of industries.

How do you decide where to invest?

One Direct Approach–the GE Vernova Example

A direct approach might be to buy a manufacturer of gas turbines for electric power generation, like GE Vernova. It already has about a four-year backlog. In 2022, it had orders for 30 of its 500-odd MW units—the world's most powerful—and is increasing capacity to deliver up to 80 next year. Only a modest portion of the backlog is related to data center demand; that bounty is yet to come. It also has an order for one of the first small modular nuclear reactors in the U.S. The company is number 89 in the S&P 500. A possibility, but it is a single-company risk. (There's a more fundamental problem, too. *See the text box, below.)*

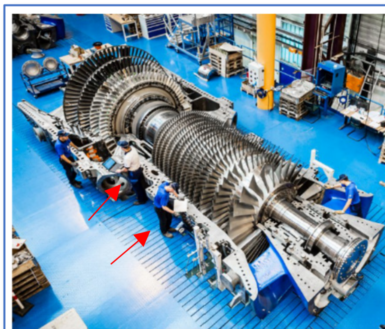

With two available models—the 9HA.01 gas turbine at 448 MW and the 9HA.02 turbine at 571 MW—customers can easily select the right model and capacity to help meet their generation needs. (GE Vernova)

(The Paradox of) The Indexation Approach

The default choice is the **indexation approach.** The logic in favor—aside from the assertions about its efficiency and safety versus individual security selection—is that it's already chock-full of AI exposure. Health Care, at 11%, is the third largest S&P 500 sector, and we already know how pharmaceutical companies can benefit.

> **A note on GE Vernova:** *One must be a little careful here to not mistake a consumer of critical resources for a provider.*
>
> Many would leap to the conclusion that GE Vernova is a provider. Rather, as a manufacturer, the company must purchase steel and electric power—and a lot of it—in order, paradoxically, to supply turbines and generators that produce electric power. The Vernova 9Ha turbine weighs 9 tons.

Financial Services, the second largest S&P 500 sector, at 15%, is an obvious beneficiary of institutional-scale AI services. Banks' largest non-interest operating expense is compensation. JPMorgan pays $51 ***billion*** to its 317,000 employees. AI could bring meaningful efficiencies.

Balance sheet management is even more important. With a $4 trillion balance sheet, JPMorgan has every sort of asset—from wholesale credit, retail loans, credit cards and mortgages to real estate. They're distributed across industries, credit quality and geographies. There's currency and commodities trading; the associated hedges and derivatives; varying debt leverage on all these asset groupings and of course on the individual loans and positions within them; and on and on. Buried within all this, aside from efficiencies that might be realized, is a lot of risk. The problem for the bank is, "How much, and where is it buried?"

However many tens or hundreds of thousands of positions it has, they are too numerous and complex to be understood on a loan-by-loan or security-by-security basis. They can only be assessed on a statistical basis—precisely what generative AI can do exceedingly well.

The default asset allocation decision is to participate via the major indexes. The IT companies are already there, and so are their primary AI-beneficiary customers.

There's a different way to look at this gameboard. The associations just described mean that **the IT, Financial Services, and Health Care sectors share a functional covariance. Together, they account for 66% of the S&P 500.**[16] Seen in this light, the ascendance of AI services concentrates even more technology risk in the index.

[16] Including, as always, Amazon, Meta/Facebook and Alphabet/Google, which are among the largest AI companies, but are officially in the Consumer Discretionary and Communications sectors.

Until now, the biggest risk was IT's 40% weight. Put less dramatically, this covariance reduces the index's diversification resilience.

Thinking about these interdependencies, the high-margin IT companies have begun to load their formerly asset-light balance sheets with vast amounts of physical resources to build their cloud/data center services. Their commercial customers, like Financial and Pharmaceutical sector companies, will buy vast quantities of AI services. In a sense, they are all consumers of AI related resources. Shareholders, in turn, will be consumers of those companies' financial returns, which will partly be a function of the supply and pricing of the physical resources. *Everyone's heading to the same square on the gameboard.*

This raises a big-decision question: **is it better to invest in the data center companies, which consume the resources required to build and operate them, and in the consumers of the data center AI services? Or is better to invest in the providers of the resources required to build and operate the data centers?**

Here's a look at the gameboard.

Observation: The additional electric power needed for the data center buildout is startlingly large. Among the primary commodity inputs without which there won't be electric power, are these two of the three that were reviewed in the *4th Quarter Commentary*:

- **Natural gas**. Not only will it be the fuel of choice, it already is. U.S. electricity generation in 2023 is no higher than in 2005, but natural gas usage for power generation more than doubled, while coal's use declined two-thirds.
- **Water.** There is no other choice. That's for steam to move the turbines in thermal power plants (gas, coal and nuclear), and for cooling the data centers and the power plants.
- Followed by steel (which is 98% iron), copper, and other hard commodities.

It is beyond question that data center expansion on the scale envisioned will impact commodity demand and pricing.

Observation: There are vastly more investment dollars allocated to resource consumers—the AI and data center companies and their customers—than there are basic resource providers.

While the AI/data-center consumers are 66% of the S&P 500, the resource providers comprise less than 4% of the index. Really, hardly more than 1%.

U.S. Electricity Generation by Source

Source (in bill. kWh)	2005	2023
Coal	2,013	675
Natural gas	**761**	**1,802**
Nuclear	782	775
Renewables	358	894
Petroleum & other	142	32
Total	**4,056**	**4,178**

A brief aside on steel reveals some of the unavoidable feedback loops in an economy.

Coal use for electric power declined by two thirds in the past two decades, and some of that was replaced with wind turbines; renewable electric power production rose 150%. However, a 5 MW offshore wind turbine requires up to 900 tons of steel, which requires up to 2,000 MW of electric power to produce,* including metallurgical coal.

Data centers also use large quantities of steel in the construction of the campus buildings, the server racks, security caging, airflow containment, and other infrastructure equipment. Last April, Nucor, the steel maker, purchased a data center infrastructure equipment manufacturer, and has created a new subsidiary known as Nucor Data Systems as a way to participate in data-center steel demand growth.

*https://ieefa.org/sites/default/files/2022-06/steel-fact-sheet.pdf

- *Metals mining:* 0.25%
- *Energy:* 3.3%. However, this is weighted toward oil. Based on its contribution to energy company production volumes, natural gas would have an index weight of 1%.
- *Land:* 0%. Unless one applies the proportion of Texas Pacific Land Corp. market value that is attributed to land, but that figure does not fully exist on the balance sheet. In any case, it would have to be some fraction of TPL's 0.05% index weight.
- *Water:* 0%. Unless one applies the proportion of TPL revenues from water, in which case the water weight in the index would be 0.02% (rounded up from 0.018%).
- *Metallurgical coal*, for steel production: 0%
- *Data center REITs*: a very iffy entry, only because it looks good to try. REITs are an asset intensive and debt-leveraged business structure. Because REITs must distribute 90% of net income, they typically resort to equity issuance to fund expansion, so they are generally poor *per-share* earnings and book value compounders relative to their total growth rates. For what it's worth, there are two data center REITs in the S&P 500, and they total 0.27%.

Realty Income Corp (O)	Annualized Change 25 Yrs. to 2025
Revenue	17.3%
Net income	12.9%
Total dividends paid, dollars	16.4%
Shares outstanding	11.7%
Revenue/share	5.0%
Net income/share	1.0%
Dividends per share	4.1%

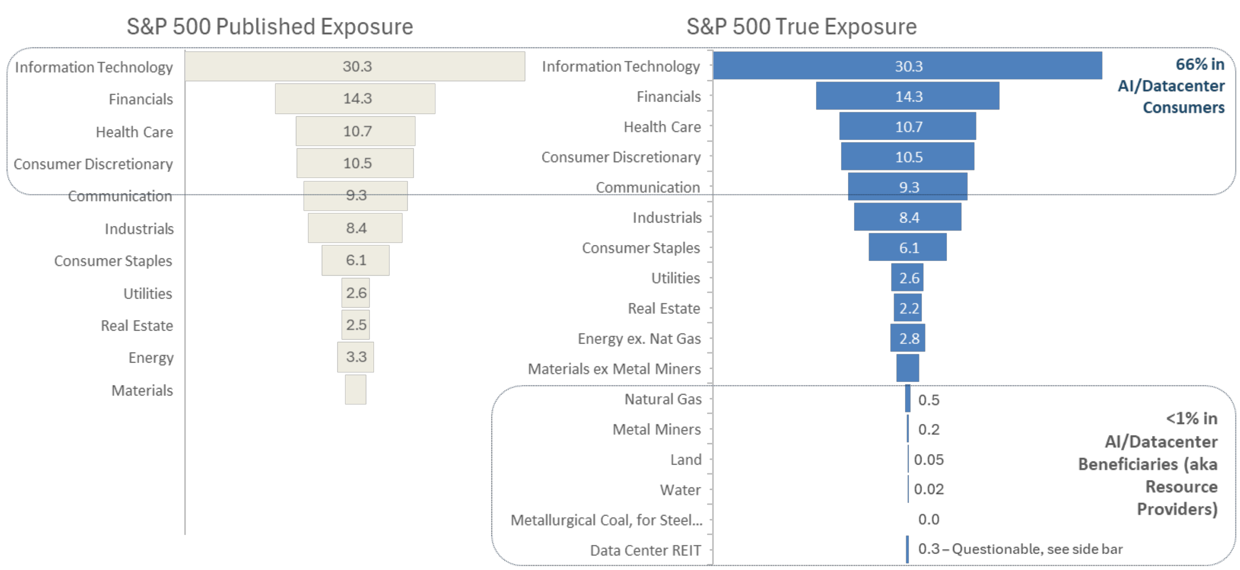

For a broader picture beyond the S&P 500 Index, the total market value of equity ETFs is 7.8 trillion,[17] and they number 2,859. For those intrigued by more bespoke offerings, etfdb.com lists 79 different industries. Of those, there are 12 industries, with 82 funds that sound as if they could be supplier beneficiaries to AI-data center demand growth. Their total AUM is $98 billion,[17] which is a 1.3% aggregate share of ETF dollars.

[17] *Source: https://etfdb.com/etfs/, as of April 28, 2025*

"Sound as if" because—well, because we can't help ourselves—we do have to look under the hood. Two examples will illustrate the challenge of getting resource supplier exposure via indexation.

	# of ETFs	AUM ($MM)
Broad Energy	6	$1,919
Broad Materials	14	$10,012
Commodity Producers	1	$0.001
Copper Miners	1	$2,192
Gold Miners	14	$23,982
Hard Assets Producers	1	$146
Metals & Mining	4	$2,415
Natural Gas	1	$307
Natural Resources	7	$8,968
Oil & Gas Exploration & Production	21	$40,288
Silver Miners	5	$2,926
Water	7	$4,872
Total Data Center Suppliers	**82**	**$98,027**
Total Equity ETFs	**2,859**	**$7,815,806**
% Market Share		**1.3%**

The largest of these sectors, Oil & Gas Exploration & Production, about 40% of the total AUM, is a little problematic, since only about one-third or less of oil company production is natural gas. Very little oil is used in electricity production.

The largest of the Broad Materials ETFs (Materials Select Sector SPDR Fund, XLB, $5 billion of AUM) has only two metals producers. They're 10% of the fund's value. The rest of the holdings are mostly manufacturers, such as of chemicals, paints, packaging, paper, industrial gases and so forth. XLB's security selection rule set is simply to replicate the S&P 500 materials sector.

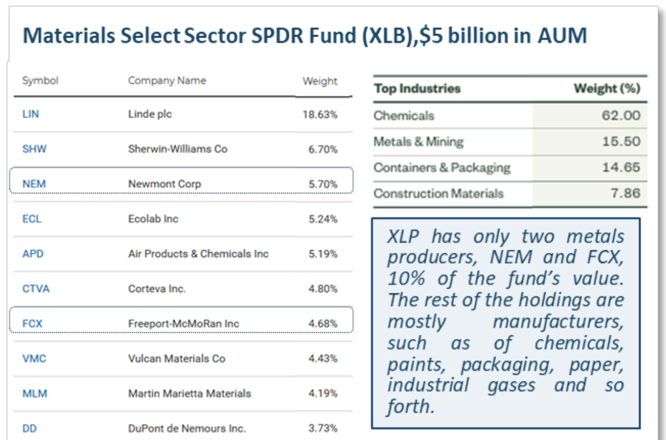
Materials Select Sector SPDR Fund (XLB),$5 billion in AUM

Symbol	Company Name	Weight
LIN	Linde plc	18.63%
SHW	Sherwin-Williams Co	6.70%
NEM	Newmont Corp	5.70%
ECL	Ecolab Inc	5.24%
APD	Air Products & Chemicals Inc	5.19%
CTVA	Corteva Inc.	4.80%
FCX	Freeport-McMoRan Inc	4.68%
VMC	Vulcan Materials Co	4.43%
MLM	Martin Marietta Materials	4.19%
DD	DuPont de Nemours Inc.	3.73%

Top Industries	Weight (%)
Chemicals	62.00
Metals & Mining	15.50
Containers & Packaging	14.65
Construction Materials	7.86

XLP has only two metals producers, NEM and FCX, 10% of the fund's value. The rest of the holdings are mostly manufacturers, such as of chemicals, paints, packaging, paper, industrial gases and so forth.

The largest water sector fund, the Invesco Water Resources ETF (PHO), has $2 billion of AUM. Among the industry sector representations within the fund are Machinery, Building Products, and Chemicals. They each probably have something to do with water, but they're not actual water exposure.

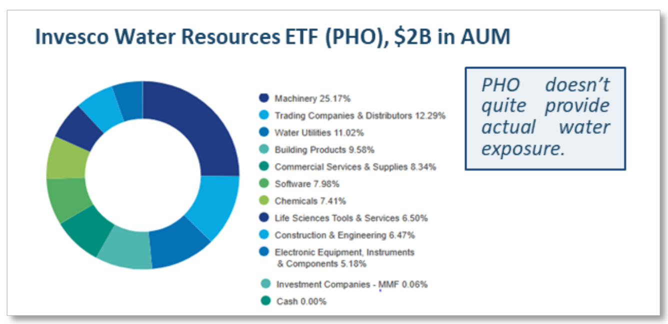

The closest match is the 11% in Water Utilities. Water utilities are probably not getting you to the data center-growth promised land. They are regulated and can't simply sell to the highest bidder. The ones with the most water volume are probably located in population centers. The evidence to date is that municipalities will not tolerate a large-scale data center claiming their water. Plus, the water must be proximate to plentiful natural gas supply and suitably sizable acreage.

In any case, a regulated water utility is a different business than pure free-market water provision, such as by a land company with subsurface water rights, like TPL. It can also be a water pipeline company that directly provides water delivery and takeaway services, like LandBridge offers, or Aris Water Solutions.

A pending supply/demand imbalance that extreme, especially if not yet manifest, is something investors should really pay attention to.

Turning Asset Allocation on its Head

Indexation and large-scale investing follow a logical approach: if an industry is a large part of the economy, it's wise to own it in a proportionate fashion. You partake of the long-term economic growth, magnified through the operating leverage and profits of the representative companies. Otherwise, you're underrepresented in the growth of the economy. Contrarily, you wouldn't invest much or at all in marginal sectors, which might not fare at all well even when the overall economy is robust.

But can those dynamics work the other way 'round? Can something that's not well-represented in the indexes or economic statistics nevertheless be critical to the economy and to those large sectors and companies? Here's a recent relatable example.

Because of the avian flu, egg prices just about doubled between year-end and March. But since they have only a 0.17% weight in the CPI, their impact on the published inflation figures was negligible. Besides, the 100% egg price increase was substantially offset by the roughly 30% price decline of potatoes and pasta, which together have a 40% greater weight. The Consumer Price Index registered pretty low inflation: in the three months from December to March, it rose by a nice, normal 0.63%, or 2.6% if annualized. Nevertheless, egg buyers did in fact experience meaningful—call it "localized"—inflation.

If apple prices had doubled, consumers might readily have substituted them with pears, oranges or grapes. Eggs, though, are not easily substituted, and egg prices became a cause of national political concern. Yet the supply declined by only 10% during those few months.[18] That's what happens when there is a supply constraint for a commodity with fairly inelastic demand; it can happen with any commodity.

A Note on Historical S&P 500 Sector Norms

The absence of commodity and inflation-beneficiary companies in the S&P 500 is not *the* norm, only a recent and non-permanent norm.

That absence is the result of decades of increasing financialization of the securities markets during a long period of artificially low interest rates. Those rates were enabled by a long disinflationary period, that disinflation being partly enabled by a sustained supply surge of hard commodities from China and the Soviet Union. Those exports accompanied reversals, in the 1980s, of the then extant restrictions on international capital flows, which relaxation in turn allowed the development of novel financial instruments to hedge such flows and trade, and which in turn engendered greater and greater volumes.

A lot of long term geopolitical, regulatory and monetary policy support went into the current historically extreme market structure. It is not an inherent norm; it only exists within the context of forces that support it.

Investors were once predominantly hard-asset oriented; future circumstances might dictate they again turn their attention there.

A lesson for anyone wishing to hedge against or profit from inflation—they're really the same thing—is to concentrate on assets and commodities for which there is no immediately available substitute, most especially in a period of increasing demand. That is a circumstance tailored for localized inflation.

[18] *https://www.ams.usda.gov/mnreports/ams_3725.pdf*

As such, there's probably general agreement—humorous asides aside—that natural gas and electricity are far more important to the economy than coffee and eggs; coffee prices also doubled this past year. And, like eggs, we already saw the negligible weightings of such hard commodities in the equity indexes. They don't even approximate their GDP weightings. Energy is reported as 8% of GDP. [19] Water isn't reported in GDP and it's not in the CPI, either. But natural gas and water are a lot more important to the AI/data center industry than even coffee and eggs are to us breakfast eaters, even if we do consider coffee to be a critical asset.

This is an object lesson—in diametric opposition to today's market structure orthodoxy and received knowledge—that *the serious performance advantage can now reside not in the large-scale investing model, which is done at the $100-billion-plus scale,[20] but in **small-scope investing**. And for the AI/data-center era, in **limiting-factor assets**.*

"Cornering" the Market with Small-Scope Investing

Attempting to corner the market in a security or commodity, by securing control of a large portion of the supply, is illegal and closely monitored by regulators. The idea is to force the price up by artificially limiting the supply. An artificially boosted price might create a situation with forced buyers who had sold short, typically on a leveraged basis, such as in the commodities futures markets. The short-sellers had to purchase the item in question to close out their position. That was the "short squeeze."

But what if cornering a market could somehow be replicated legally and naturally? If it involved no market manipulation in the legal sense of the term? Weirdly, large-scale index investing has itself constructed the conditions for a number of cornered markets, and which are open to the general public to employ.

It might be described this way. Indexed Equity ETFs and mutual funds amount to about $14 trillion.[21] This does not include in-house indexing by large institutions like endowments and defined benefit pension plans. If none of those pension fund assets are in the publicly traded ETF and mutual figures, they could be another $12 trillion, [22] or even greater by some calculations, but not wanting to double count, the figure will be used provisionally.

As ETF promoters seek ways to provide exposure to the data-center-growth phenomenon, their attention will at some point turn to the enabling commodities, and to companies that are associated with those commodities. If asset allocators' attention translates, even if only to the marginal degree of **a single percentage point** of portfolio allocation, applied to this rough estimate of indexed equities of somewhere between $14 trillion and $26 trillion, that **would be $140 billion to $260 billion of buying interest**. Moreover, for this scenario, there's no need to restrict the potential interest to indexed equity assets—actively managed portfolios should ultimately have a keen interest in this theme as well. In any case, there simply is not enough market capitalization of this relative handful of companies to go around. Those equities would, in that sense, be overwhelmed, which would be expressed in their share prices.

---

[19] https://www.statista.com/statistics/1451878/share-gdp-oil-and-gas-production-select-countries-globally/
[20] That's GE Vernova, 0.2% S&P 500 weight, $100 billion market value.
[21] https://www.ici.org/research/stats/combined_active_index, as of March 31, 2025.
[22] https://www.pionline.com/exchange-traded-funds/worldwide-index-assets-18-vanguard-stays-top-and-worries-about-concentration. Data as of June 2024.

Being pre-positioned in such companies could be very much like participating in a cornered market, albeit of indexation's making.

Section III: Portfolio Insights—The Use of Predictive (as opposed to descriptive) Attributes, and A Few Data-Center-Expansion Beneficiaries

Predictive Attributes and Time Arbitrage (AKA Equity Yield Curve)

Indexes, and index and stock screeners, are based upon descriptive attributes. These are statistical measures, starting with the GICS sector and industry codes[23] into which companies are classified, followed by their stock market capitalizations, and trading liquidity. There is also dividend yield, price volatility, whether companies have experienced rapid earnings growth, have high or low valuations, and so forth. The statistics can describe how a constituent company has performed, and there are published earnings estimates. But none of these inform how well a company might do in the future. They're not predictive.

We used to write a lot about predictive attributes—if you believe in that sort of thing—which are features that predispose a business or security to do well in the future. The most straightforward example is probably the dormant asset. A company might own an asset that produces no revenue or earnings. The company is reasonably priced relative to its reported earnings, which means the share price includes no value for this asset. It might be a piece of land acquired so long ago at so low a cost that it doesn't even qualify as a line item on the balance sheet. It's invisible.

HORIZON KINETICS
PREDICTIVE ATTRIBUTE SERIES

Note: The below selections represent sample research reports as of the listed publication dates. There have been no edits made to these research reports since they were published. As such, the data herein should not be relied upon without first reviewing more recent figures.

TABLE OF CONTENTS

Nevertheless, the land could be located on a block of a city business district that is undergoing a redevelopment. That lot might be worth a great deal, maybe as much as the company's operations, but its potential earnings or sale value won't be visible until the company takes some action. This is not a made-up example. A dormant asset can be predictive of a superior return. It's also, really, a function of another predictive attribute.

This other predictive attribute is not as discrete or easily demonstrated, but is universal in that it can apply to just about any security. That is the Equity Yield Curve. Some of our analysts, reflecting a newer generation of schooling, prefer the term Time Arbitrage. It is relevant to the three relatively recent additions to some of our strategies that will be reviewed shortly.

[23] *Global Industry Classification Standard methodology*

In brief, the equity yield curve is a way of showing just how much return investors require from a stock relative to how long they believe they must wait for that return. It's upward sloping, just the way a bond yield curve usually is: The farther the maturity date to recoup the principal, the higher the yield that investors want in return. Below is a 20th Anniversary depiction of a more classic yield curve than has been lately extant.[24]

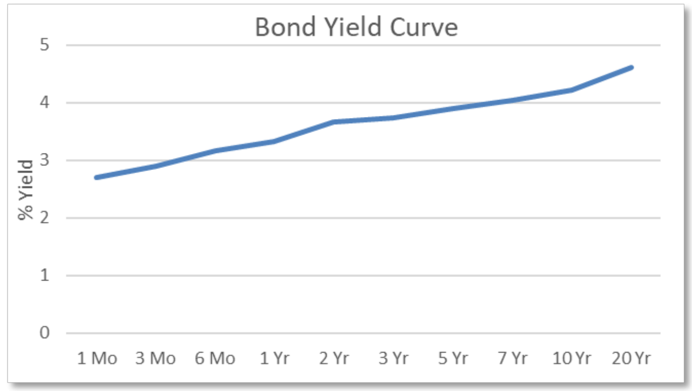

The equity yield curve is way, way steeper, though. Unfortunately, it's very rare that you can actually plot one. The problem is that you can't have reasonable confidence in what the price of a stock should be one year, two years, three years in the future; valuation is highly subjective. You say NVIDIA will grow at 40% for 10 years and maintain a P/E of 40. Someone else might say 20% and a final P/E of 12. You'll both have cogent arguments, and reasonable minds may differ.

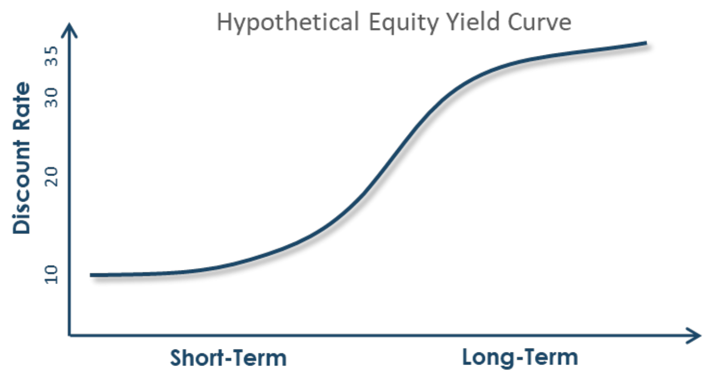

Once in a long while, though, an odd security shows up that does provide a fair degree of confidence about its future value on or about a certain date. It could be a pre-existing contract for a future buyout or a put feature that has a date and price or valuation methodology attached; it could be some form of arrearage or legal claim with sufficient enforceability to rely upon. With those two reasonably objective data points, an equity yield curve may be plotted.

Do that every time you locate such a security—even across different economic and yield environments, and vastly different businesses and asset types—and their yield curves are remarkably similar. In practice, once you approach the three-year range, the offered annualized return can be 35%.

This example from a past portfolio holding and associated research report is a pretty good one. It was for any or all of the nine series of preferred shares of Pacific Gas & Electric, which filed for bankruptcy in April 2001, as a result of the spreading financial crisis catalyzed by the Enron debacle. The company stopped paying the preferred dividends, and the preferreds fell to two-thirds of face value.

By May 2002, both the company and the California Public Utility Commission itself had put forth competing plans of reorganization. Although they differed in many respects, they were agreed on two issues: They were structured to achieve an investment grade rating for the company, and to result in no impairment to the preferred shareholders, including dividend arrearages. Accordingly, the substantive financial risk was

[24] *https://home.treasury.gov/resource-center/data-chart-center/interest-rates/TextView?type=daily_treasury_yield_curve&field_tdr_date_value=2005*

erased and a post-recovery value could be estimated with some confidence, whether it was to the face value of the preferreds or to the dividend yields at which non-troubled utility preferreds then traded.

From the May 2002 purchase, allowing two years for emergence from bankruptcy and the amount of arrearages that would be payable in May 2004, the final value would be, adding the par value and accumulated dividend arrearages, over 80% higher than the then current share price. Prospectively, in May 2002, the two-year annualized return was, roughly 35%:

Series	Price on May 21, 2002	Price on April 2, 2004	Arrearages	Annualized Return
6.0% PCGpfA	$ 18.00	$ 29.73	$ 4.875	30.8%
5.5% PCGpfB	$ 16.00	$ 28.26	$ 4.688	35.6%
5.0% PCGpfC	$ 14.50	$ 25.51	$4.0625	35.3%

Source: Bloomberg, Company reports

With every passing month after May 2002, one could check the trading price of the preferred, and calculate the annualized return to May 2004, to see what annualized return investors did in fact require for, first, a 23-month value realization, then a 22-month yield, a 21-month yield, and so on.

> "The stock market is a device for transferring money from the impatient to the patient."
> ***Warren Buffett, quoted in The New Zealand Herald***

Here's the thing, as the yield curve below will show. As strange as it sounds, the size of that wonderful ***prospective*** return—like a bull market in a bottle—is not that important to a time-constrained and relative-return-judged fund manager. What's important is WHEN? How soon can I have it?

The longer it will take to earn that return, If well beyond the standard institutional 12-month time horizon, the less and less value or utility that

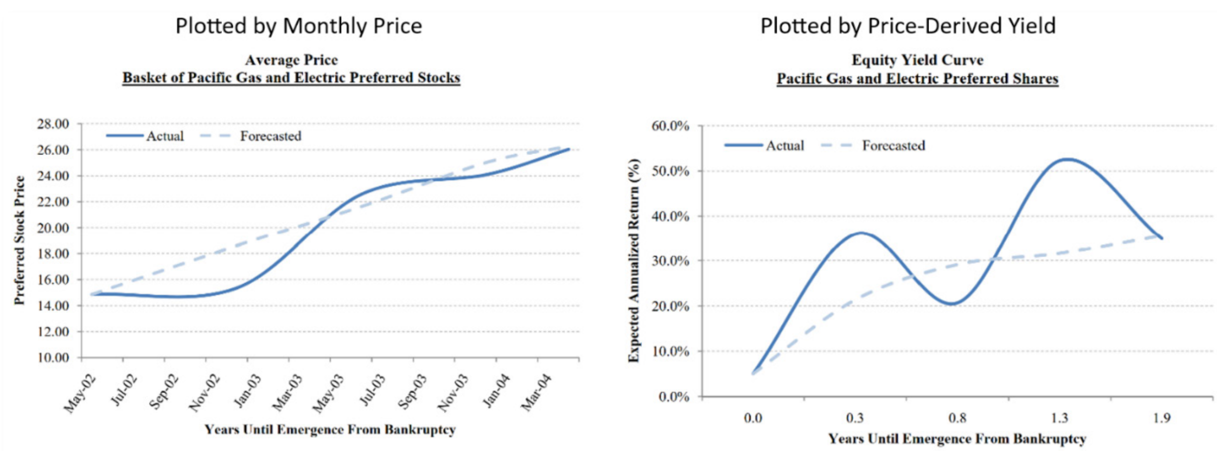

return has for them. Particularly if the timing is uncertain—in this case, it might have been as soon as 8 months, it might have been 2 ½ years. The institutional disinterest translates to a lower price/higher yield. But it interests ***us.*** Along with the high discount rate or yield, we'll take the interim time and performance risk.

Section IV: Case Studies—Some Data Center Expansion and Equity Yield Curve Portfolio Companies [25]

The following company synopses were written, on short notice *(see sidebar)*, by two of our research analyst/fund managers. Among their other benefits, each of these three companies is priced somewhere deep in the equity yield curve.

Two of them have positive exposure to limiting-factor resources necessary to the vast data-center-buildout demands.

Two of them—still only three companies—share very recognizable valuation and price pattern profiles with securities we've written about and owned before.

It was Carl Reiner, commenting about deliberately throwing an unexpected question at Mel Brooks during their 2,000 Year Old Man comedy sketches, that (paraphrasing heavily), it was a wonderful thing to corner a gifted mind into a panic and see what emerges.

San Juan Basin Royalty Trust, like Mesabi Trust, is an entity with no employees. Rather, it has a trustee with a contractually circumscribed compensation arrangement. Both Trusts merely pass their royalty income on to their unitholders as direct earnings from the underlying hard assets: natural gas and iron ore. Also, as with San Juan, Mesabi had announced a dividend suspension and just recently reinstated it, with the logical result in the fully-recovered share price. It was just about a two-year journey. They further have in common that the dividend suspension was (and is) confidently known to be a temporary condition with a reasonably predictable time horizon. A nice time arbitrage example, although we haven't plotted Mesabi's time arbitrage chart.

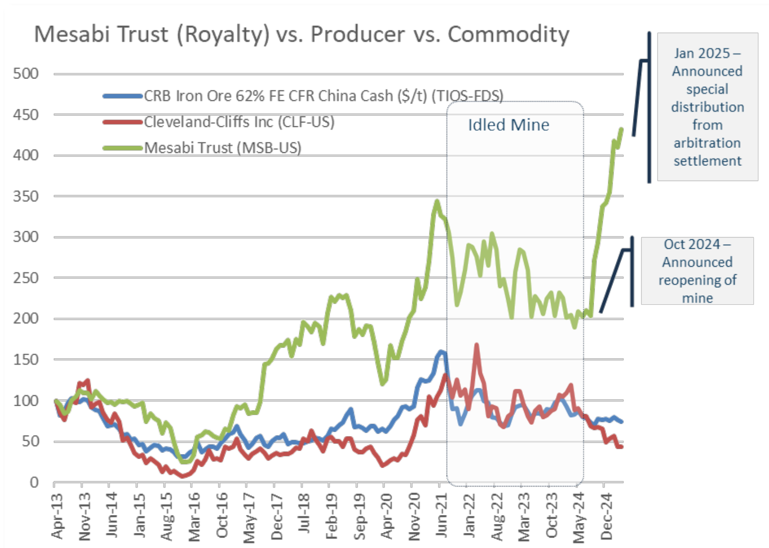

Hawaiian Electric Industries, as a utility that had to suspend its dividend in the face of bankruptcy-level liabilities, is very much in the mold of past utilities in similar circumstances as or shortly after their existential risk was erased by a regulatory and judicially approved recovery plan. The post-recovery value is readily estimated, as is the time frame, so it's just a matter of waiting. An interesting question about such a deeply discounted utility is whether, in a pure investment sense, it is a utility stock or something else entirely. Should developments unfold as expected, its highly asymmetric return profile from this point forward, as well as its financial progress, will look nothing like a utility stock.

---

[25] These selected companies are holdings in several funds and strategies managed by Horizon Kinetics Asset Management LLC. Not all investors will experience the same holdings, returns or weightings as the corresponding composite subject to constraints related to timing and implementation restrictions/considerations among certain group of accounts.

Aris Water Solutions (ARIS)

A conventional approach to sourcing investments: Identifying favorable conditions for capital appreciation. Reasonable as this approach may be, favorable conditions are often accompanied by robust valuations. There is little use in identifying favorable conditions if the market has fully priced in these fundamentals (and perhaps much more).

If we invert this approach and consider pervasive *unfavorable* conditions that might obscure an otherwise excellent value proposition, we are left with much more interesting opportunity sets for undervaluation. And we can examine Aris Water Solutions as a particularly timely security to discuss. There is no shortage of unfavorable conditions impeding the market from recognizing its value. However, we appear to be at an inflection point where the longer-term growth and profitability potential of the company is readily apparent.

Conventional investment approaches overlook Aris for several reasons. It is technically an oilfield service company (strike one), and it came public in 2021 as a small capitalization issuer (strike two) via several private equity market sponsors (strike three). It was an inauspicious IPO setup for the company, as the world was still grappling with energy demand concerns related to the pandemic, ESG considerations-dominated committee meetings, and profitless technology company valuations soaring with abundant capital flowing. It's a wonder that the investment bankers were even able to place Aris shares in the market.

We followed the initial offering from the sidelines, cautious of the leverage ratio, and had similar energy water infrastructure exposure through our preferred capital-light land holdings. Aris predictably traded in a volatile but mostly flat range for the next three years. Aris piqued our interest after we participated in the LandBridge IPO in 2024, and came to better appreciate the tremendous value of water infrastructure during our industry research.

Our firm is no stranger to the Permian Basin water industry, by virtue of our longstanding investment in TPL. Water is both a crucial input for fracking shale and a substantial well byproduct output. The input, or injected "source water," is very different from the formation "produced water" that comes out of the wells. The former is generally sourced from aquifers and injected into wells to stimulate oil production. This water must meet certain purity standards so it doesn't damage drilling casing and equipment or disrupt the natural flow of the hydrocarbon mixture.

The produced water is a high-saline brine with various toxic compounds and high total dissolved solid content. Suffice it to say, this water is generally neither suitable for reinjection into wells, nor irrigation. In fact, it presents a major environmental liability, and must be transported from the well site—either to be cleaned for reuse in fracking, or disposed of via injection into saltwater disposal wells.

In the Delaware Basin, there are approximately four barrels of produced water generated for every barrel of oil equivalent. This "water cut" only increases as deeper shale formations are targeted and existing wells age. Thus, the basin requires at least four times the amount of infrastructure for water as for hydrocarbons. This is a large and growing business—but many oil and gas companies want nothing to do with it. Let's explore why.

Water management is an expense to oil and gas producers—the largest and fastest-growing part of lease operating expenses in the Delaware Basin—directly raising drilling break-even prices. The producers are focused on producing and, to a lesser extent, transporting their hydrocarbons, not managing the related bonded hydrogen and oxygen molecules. These companies generally don't want such activities on their balance sheets, as it would take capital away from drilling—and from shareholder returns. Firms increasingly appreciate that water is a highly specialized and technical business often better outsourced. In response, a concentrated group of specialty water infrastructure companies has emerged.

Aris was built from the ground up by Solaris Energy Capital to manage wastewater from Permian oil and gas wells. This generally involved permanent and semi-permanent pipelines to transport water from the wellhead to a disposal well. Most pipelines and "handling facilities" are operated via land leases, rights of way, and/or royalty injection agreements with the landowners. To provide context, TPL and LandBridge are amongst the largest landowner lessors of land to third-party water management companies.

The most prolific portion of the Delaware Basin spans the border between Texas and New Mexico. This is a critical aspect of water management due to differing regulations across state lines and varying private/state/federal lands. As a result, most of the Aris network is in New Mexico, and much of the economics are predicated on getting excess water into Texas, where it can be more easily disposed of. This setup presents an enormous competitive advantage to both the landowners (recipients of land lease and injection royalty revenues), but also incumbent operators with established infrastructure (such as Aris).

Aris is paid to take barrels of water away from the well head, but it doesn't have to simply dispose of it as waste. As a result, the company has developed a complementary recycling business that filters the water to standards acceptable for reuse in fracking. To the extent that this can be done economically, the company can earn a second revenue stream on the same barrel of water that would otherwise be disposed of.

Aris handled approximately 1.142 million barrels of water per day in the first quarter of 2022, at an average price of $0.68/barrel and an adjusted margin (on a cash operating basis) of $0.42/barrel, or 61.7%. The volume mix was approximately 70% produced water ($0.78/barrel) and 30% recycled water ($0.45/barrel).

In the subsequent three years (through the fourth quarter of 2024), the company's total volumes rose by 14% per year and revenue per barrel rose to $0.75. Aris has maintained an adjusted margin of $0.44/barrel. This growth has been a product of both organic volumes within its existing pipeline network as well as capital investment.

While the company operated at nearly a 60% adjusted margin per barrel in the fourth quarter of 2024, it's worth noting this was based on approximately 60% and 30% capacity utilization for, respectively, the produced water and recycled water systems. It would not be unreasonable to expect further margin expansion from higher capacity utilization. But this is only half the story.

The Permian Basin does not have infinite space for wastewater to be injected, and many of the most cost-effective areas near major oil and gas activity are rapidly approaching capacity. This is a function of pressure and porosity, where there has been seismic activity related to water injection, as well as interference with drilling activity.

As a result, regulators are limiting new injection permits, and water companies are developing long-haul pipes to remove water from the congested areas. This comes with a cost, and pricing leverage will belong to the landowners and incumbent infrastructure operators.

If we revisit the unfavorable conditions previously faced by Aris, which were plainly visible in the moment, we can present the longer-term, more economically significant conditions that were developing: secular water volume growth, price inflation, and barriers to competitive entry.

There is also nascent electric power and data center development in the Permian Basin—notably, being pursued by LandBridge. Thermal power generation (gas combined cycle) and data center liquid cooling are immensely water-intensive, and most of the Permian is a desert.

It remains to be seen what, if any, revenue exposure the incumbent companies like Aris will have to this water market. Clearly, though, this "option" is not priced into the stock, which trades at approximately 8.5x estimated pre-tax cash flow for next year.

-James Davolos

San Juan Basin Royalty Trust (SJT)

A casual energy market observer could be forgiven for not knowing the water dynamics in energy production, or the prevalence of oil and gas produced in New Mexico. Despite the common coupling of Texas with the prolific Permian Basin, the formation stretches well into Lea and Eddy counties in the southeast corner of New Mexico. And, between 400 and 500 miles northwest of these counties—mostly in New Mexico and along the borders of Colorado, Utah, and Arizona—lies the San Juan Basin. This is a largely forgotten gas field with methane primarily produced from a coalbed formation.

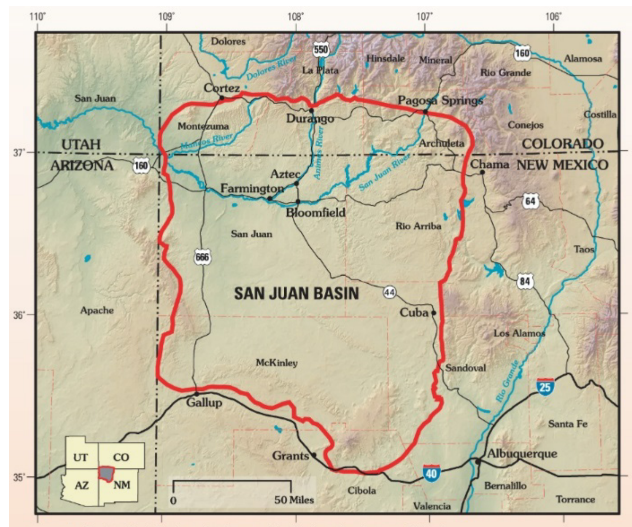

San Juan Basin Province of northwestern New Mexico and southwestern Colorado.

Speaking of unfavorable present and observable conditions for a stock, we present San Juan Basin Royalty Trust. As the name implies, the "company" is actually a trust that distributes a net profit interest (NPI) from certain properties in the area. It has no debt and no employees. It has a corporate Trustee, the compensation for which is limited by a schedule in the Trust Indenture. It begins at 1/20th of 1% of the first $100 million of annual gross revenues of the Trust, allows for hourly rates in excess of 300 hours a year, and establishes a minimum fee of $36,000 per year. In 2024, the Trustee fees totaled $120,108, although total general and administrative expenses were $2.1 million.

Natural gas may be appreciated by a growing minority as a cleaner fuel source compared to coal, but it is widely loathed by investors due to perceived oversupply (and attendant price volatility) in the United States.

The U.S. has an abundance of high-quality, low-cost natural gas reserves. This dynamic is exacerbated by a highly seasonal gas consumption pattern and limited storage/export potential relative to production levels. As a result, U.S. natural gas is a highly volatile market that has been responsible for many fortunes made and lost.

This dynamic, however, is rapidly changing due to enormous incremental gas power demand (data centers, reshoring of industry, etc.) and a doubling of LNG export capacity in the next 18 months.

While this will promote drilling at the margin, there are logistical pipeline constraints that make the market far less fluid than reserves data suggest.

These factors could quickly shift the market from structural surplus to deficit.

The thesis is not about gas pricing, but that contributes to unwarranted pessimism for the San Juan Basin Royalty Trust. The bigger headwind is that it currently distributes nothing to shareholders. This must be taken in the context of the Trust making a $0.41 *monthly distribution* per share in March of 2023. This annualized to approximately $4.10 per share, more than the stock price at the end of 2024 ($3.83). Welcome to the world of natural gas investing!

However, this is no ordinary dividend cut. It is the result of the operator (Hilcorp) increasing capital expenditures in 2024 to $36 million, up from the 2023 level of $4.4 million. This eightfold increase in spending will result in considerable production growth, but given the terms of the NPI, the investment must be recouped via earnings prior to trust holder distributions being made. Thus, there have been no distributions since April of 2024.

Hilcorp is privately held, so we must make informed estimates on the production profile going forward. However, based on disclosures from the Trust, we estimate close to a 70% production increase from the end of 2024 due to capital spending and increased drilling. Furthermore, it is reasonable to believe that the 2025 spending guidance of approximately $9 million can sustain this level of production.

Finally, if we assume comparable lease operating expenses and severance tax costs[26] going forward, we can estimate the date and amount of the distribution reinstatement (based on prevailing gas prices). For instance, if we assume that benchmark natural gas averages $3/mcf and apply a 10% discount to the local "basis," the run-rate production will result in a full paydown of the proportionate capital expenditure deficit in May/June of this year. Critically, this will be based on a monthly NPI to the Trust of approximately $3.6 million/month or $43 million/year, which equates to nearly a 17% yield based on the quarter-end share price.

This exercise delivers a dynamic emergent return scenario, particularly with incrementally higher gas prices, but quite favorable even at $3 gas prices. Of course, these hypothetically higher gas prices might prompt Hilcorp to add even more drilling, hence growing the net asset value further.

The Trust started the year with barely a $200 million market value, so we doubt many others are even running these figures. If by chance someone else is, few investors have the patience to wait for the inevitable resumption of the dividend.

While the base case scenario (conservative gas pricing and discount rate/distribution yield assumptions) delivers a return scenario that may rival anything else in the market today, there is also the "option" value of gas price related to data center and industrial demand.

Finally, it should be noted that the San Juan Basin is amongst the only regions with gas pipeline access to the Southern California market, which is prone to episodic disruptions to fuel supply, thus resulting in premium pricing for piped gas.

-James Davolos

[26] *A fixed percentage state tax levied on most natural resource extraction*

Hawaiian Electric Industries, Inc. (HE)

A direct example of the equity-yield curve in action is the utility company Hawaiian Electric. In a normal environment, the results of the business are relatively predictable. This is partly due to surrounding regulation, given the regional monopoly status utilities maintain. Both revenue and profits are governed by a performance-based regulation framework governed by the local public utility commission. This incorporates an annual revenue adjustment, and an earnings-sharing mechanism based on an established return on average common equity from regulatory capital. Any earnings that fall outside this range are shared between customers and utilities, so that both excess gains and diminished profits are limited.

In other words, growth is limited, but so is risk (absent extraordinary circumstances), so investors treat utilities more like yield alternatives. A majority of the earnings are paid out via dividends—though in comparison to bond yields, they offer potential for some annual growth; same goes for book value. Hawaiian Electric is not unique in this regard. In fact, as a utility, perhaps its most unique attribute was its ownership of the American Savings Bank, through a 90.1% stake was sold on the last day of 2024 in a cash transaction, so HE is no longer subject to capital regulation as a bank. Now HE is a pure utility more in line with peers, with the added complexity of securing energy for a remote island.

So what ties this investment to the equity yield curve? Hawaiian Electric hasn't paid a dividend since Q3 of 2023. Prior to that, dividends were paid on an uninterrupted basis since 1901. The shift came as HE faced extreme scrutiny and multiple lawsuits in the wake of the Maui fires that devastated the island in August 2023. HE chose to pause the dividend to shore up capital for any legal liabilities, and the stock dropped from $37.36 on August 8th to $9.66 on August 25th—a 74% decrease intramonth.

As of April 2025, there is far more clarity as to the company's legal responsibility. A decision by the Hawaii Supreme Court has allowed the settlement agreement proposed by HE to continue. The remaining portion of the settlement agreement requires HE to pay about $1.92 billion over the next four years. The payments will be broken out into $479 million increments, with the first one expected in Q4 of 2025.

The dividend is off the table for the foreseeable future, and the company has committed to not raise rates in response to the settlement agreement.

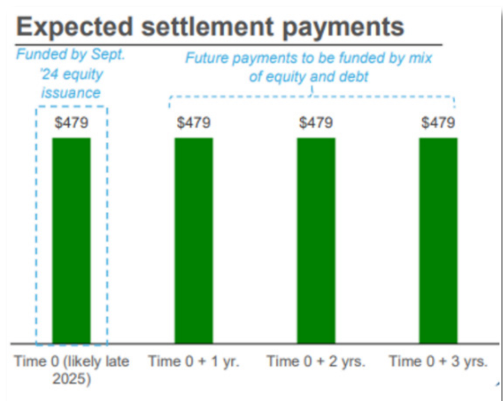

Source: Company

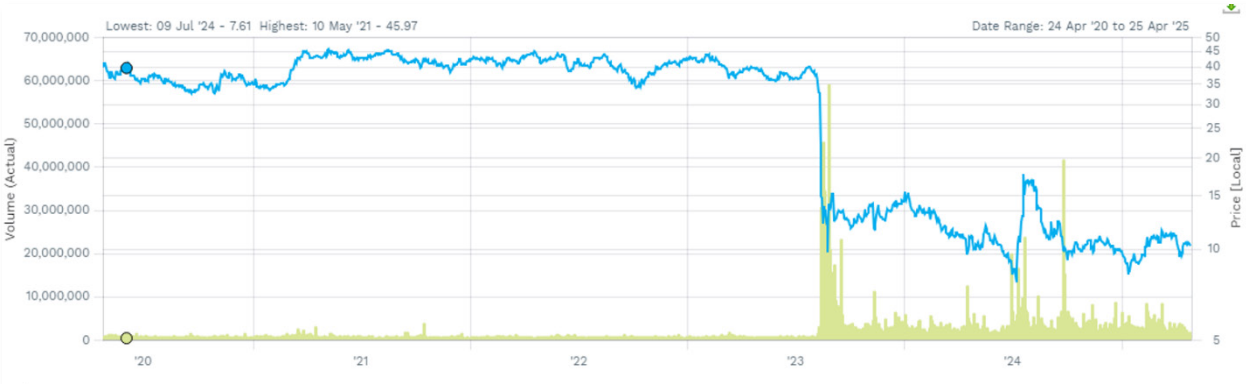

Source: FactSet

So we currently have a pure play utility company with the following attributes:

- No dividend yield
- Looming capital raise/debt issuance
- Limited ability to raise pricing or grow revenues
- Earnings-sharing agreement on hold (to avoid any lower-profitability impact with customers).

No wonder the stock trades at a small fraction of its pre-wildfire price. So why are we interested?

We have visibility into the normalized dividends from the business. We have been notified of the necessary capital raise (more on which below). We are already aware that most utilities do not have significant earnings growth, which means this could never be a solution. The suspension of the earnings sharing agreement has been noted to only be "temporary," and Q4 2024 GAAP income from the utility segment was not much different than in Q4 2022 (46.4 million vs 48.6 million). Management believes the company's current liquidity, along with some other measures taken, have alleviated the going-concern risk the company faced previously.

This is all information readily available for market participants. What the market lacks is patience.

Core income[27] from the electric utility segment was $180.7 million in 2024. The holding company portion detracted another $56.4 million, though some of these costs have likely already been mitigated due to interest/debt reduction and a simplification of the consolidated business. The core business can generate at least $150 million on an annual basis, likely more when considering the future tax benefits from prior losses.

HE needs to raise $1.44 billion for the wildfire liabilities (the first payment is already covered). To keep this simple, let's say 100% of the normalized earnings over the next four years are contributed. That equates to $600 million in earnings, which reduces the capital needs to $840 million. Subtract another $380 million, using the net proceeds from the recent sale of American Savings Bank, and that leaves $457 million for a potential equity raise. At the current price, this assumes shareholders could be diluted to about 80% ownership of the business currently, though this would be the most expensive form of financing available to the company, and therefore a conservative outlook.

The company announced that it has an additional $523MM available through an at-the-market

[27] *Excluding some non-recurring costs, such as for the Maui wildfire, and certain asset impairments*

equity program ($250 million), utility accounts receivable-backed credit facility ($239 million), and credit facility capacity ($34 million), so in reality there is a mix of liquidity options available to cover the remainder.

After all is said and done, how will this company look in four years, when the dust settles and dividend payouts seem likely? Apply some reasonable assumptions and you have visibility into a company completely overlooked due to short-term limitations.

Another way to view HE is just comparing valuations before and after the wildfires. Aside from potential dilution, is the company significantly different than it was prior to the settlement? If not, why would the valuation be so different post-settlement? The only difference is that the share price is one-quarter of the pre-crisis price or, put differently, the application of four years of a discount rate to achieve this valuation, which in the case of a regulated utility should not be very high.

-Brandon Colavita

Section V: Conclusion

The preceding dozens of pages would qualify as a long-form—maybe too-long long-form—discussion of just a few basic questions.

Is there a legitimate, long-term, high-order investment opportunity in AI and the data center **buildout**? There is plenty of mouthwatering news about it, but new technology is always an exciting show.

Without enough qualitative and contextual understanding, there can't be much confidence. Without much confidence, AI probably gets the marginal, take-a-flyer type investment treatment that won't matter much to a portfolio one way or the other.

But what makes it legitimate? How does it even work? What does it really do? What is the market for it, for the ultimate consumers of AI-based services? That requires a bit of exploration.

And there is the question of how to invest. Buy the technology maker, or choose a supplier or a distributor? From which industry? Different AI-proximate companies have different business models, different bargaining power. What do you really want to buy? Which is why it's easy to default to an index. But what's in the index? You can buy a water ETF that has no water in it, or a real estate ETF that has no data-center-ready land, just apartment buildings or storage facilities.

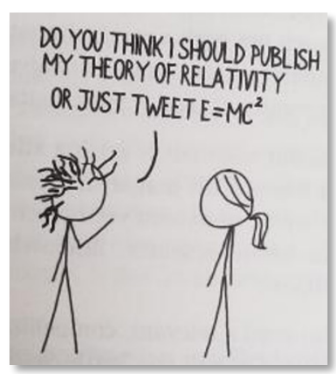

No doubt, this all could have been nicely condensed into several pages or slides, a page or two for each of the **first** four sections. **This** would have covered all the essential points and figures. Perhaps it would have sufficed. Somewhere, though, between too concise and too long, there is a border beyond which assertions and data no longer qualify as knowledge.

The best balance, then, is probably the mid-point answer to this question:

The communication cartoon above instigated an independent consultant I know, of a creative and anti-institutional mind-set, to forward these additional two. Apparently, he thought them applicable to these Commentaries:

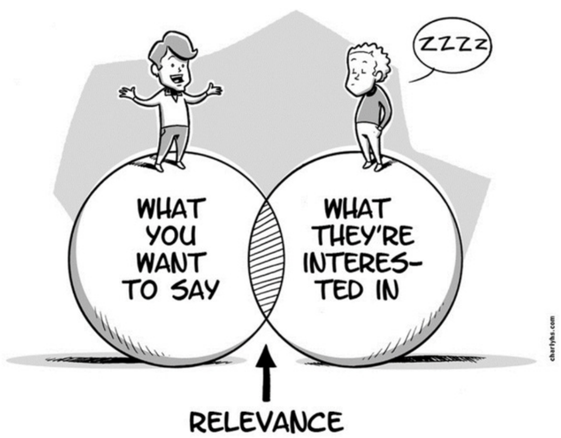

"In response to your request for better communications, I'm going to increase my lengthy explanations by 50%."

IMPORTANT RISK DISCLOSURES:

The charts in this material are for illustrative purposes only and are not indicative of what will occur in the future. In general, they are intended to show how investors view performance over differing time periods.

Past performance is not indicative of future results. The information contained herein is subject to explanation during a presentation.

Certain of the material herein is intended to portray the general nature of investor communications provided by Horizon Kinetics from time to time to existing clients. None of the investments or strategies referenced should be construed as investment advice and just because one investment is appropriate for one account does not necessarily mean it is appropriate for another. No investments should be made without the analysis of, among other things, an investor's specific investment objectives, which considers their overall portfolio and any income requirements. The accounts referenced herein pursue an unconstrained strategy—meaning they are not limited by capitalization, geographic region, or investment techniques. They generally primarily seek capital appreciation with a secondary objective of income.

Note that indices are unmanaged, and the figures shown herein do not reflect any investment management fee or transaction costs. Investors cannot directly invest in an index. References to market or composite indices or other measures of relative market performance (a "Benchmark") over a specific period are provided for your information only. Reference to a Benchmark may not reflect the manner in which a portfolio is constructed in relation to expected or achieved returns, portfolio guidelines, correlation, concentrations, volatility or tracking error targets, all of which are subject to change over time.

This material references cryptocurrencies, including bitcoin. Horizon Kinetics' subsidiaries manage products that seek to provide exposure to bitcoin and other cryptocurrencies. The value of bitcoins is determined by the supply of and demand for bitcoins in the global market for the trading of bitcoin, which consists of transactions on electronic bitcoin exchanges ("Bitcoin Exchanges"). Pricing on Bitcoin Exchanges and other venues can be volatile and can adversely affect the value of the bitcoin. Currently, there is relatively small use of bitcoin in the retail and commercial marketplace in comparison to the relatively large use of bitcoin by speculators, thus contributing to price volatility that could adversely affect a portfolio's direct or indirect investments in bitcoin. Bitcoin transactions are irrevocable, and stolen or incorrectly transferred bitcoins may be irretrievable. As a result, any incorrectly executed bitcoin transactions could adversely affect the value of a portfolio's direct or indirect investment in bitcoin. Only investors who can appreciate the risks associated with an investment should invest in cryptocurrencies or products that offer cryptocurrency exposure. As with all investments, investors should consult with their investment, legal and tax professionals before investing, as you may lose money.

The S&P 500 Index ("SPX") is a broad-based index widely considered as a proxy for overall market performance. It is the property of Standard & Poor's®.

This is not an offer to sell or a solicitation to invest. Opinions and estimates offered constitute the judgment of Horizon Kinetics LLC ("Horizon Kinetics") and are subject to change without notice, as are statements of financial market trends, which are based on current market conditions. Under no circumstances does the information contained within represent a recommendation to buy, hold or sell any security, and it should not be assumed that the securities transactions or holdings discussed were or will prove to be profitable.

Subsidiaries of Horizon Kinetics LLC manage separate accounts and pooled products that may hold certain of the individual securities mentioned herein. For more information on Horizon Kinetics, you may visit our website at www.horizonkinetics.com. The Core Value and Small Cap separate account strategies are managed by Horizon Asset Management LLC.

Murray Stahl is member of the Board of Directors of Texas Pacific Land Corporation ("TPL"), a large holding in certain client accounts and funds managed by Horizon Kinetics Asset Management LLC ("HKAM"). Officers, directors, and employees may also hold substantial amounts of TPL, both directly and indirectly, in their personal accounts. HKAM seeks to address potential conflicts of interest through the adoption of various policies and procedures, which include both electronic and physical safeguards. All personal and proprietary trading is also subject to HKAM's Code of Ethics and is monitored by the firm's Legal and Compliance Department.

Not all investors will experience the same holdings, returns or weightings as the corresponding composite. No part of the research analysts' compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views expressed by the research analysts in this report.