

Powering the AI

By John Mauldin | March 23, 2024



Long, Hard Road
Squaring the Circle
Dot Plot Thoughts
Look Who's Coming to SIC!
DC, Cape Town, and Europe?

If, like me, you're old enough to remember the 1990s internet bubble, today's AI excitement might be giving you flashbacks. The parallels are unmistakable.

Back then, a new technology was grabbing public interest and investment dollars because it promised to dramatically improve our lives. Most people didn't really understand how it worked but could see the benefits... and sometimes hear them (AOL: "You've Got Mail!").

The era ended badly for many investors, not because they were wrong but because they were *early*. Excitement outpaced reality. The technology was quite real and would, eventually, achieve what it promised and more.

(I have told this story in front of George Gilder and he didn't object, so I will tell it here. I was invited to speak at a conference sponsored by Peters and Company at Chateau Lake Louise near Calgary. This was late January and it was -40F when I stepped out of the airport and ice was literally forming inside the windows of the limo driving into the mountains, where it was oddly 20° warmer. George was also speaking and though we knew of each other, we had never met. I had a private dinner with him one night where he was moaning about the technology stocks he had thought were winners which had dropped significantly.

The price action clearly frustrated George because these were all great companies. And he was right; every company we talked about ended up being incredibly valuable and successful. But at the time, they were overpriced so we went into a discussion over valuation and pricing. I am sure I could have that same discussion today with many technology enthusiasts. There are many technology companies that are enthusiastically priced (to coin a term) but that doesn't mean they aren't great companies. Pricing? Another story for another day...

Artificial intelligence has even greater promise. I believe it's quite real, yet experience and research tells me achieving it on a world-changing scale will be tough. Energy is a top obstacle. AI systems consume truly gargantuan amounts of electricity. Where it will come from is not yet clear.

Today I want to give you some possibly eye-opening information on AI and energy. Then, since it's Fed week, I have a few reactions to the latest dot plots.

Long, Hard Road

Investors often get overexcited about new technologies because they don't consider the inevitable implementation problems. You can have a great idea, do the work to prove the concept, and still face a long, hard road to executing at scale and making real profits.

This isn't new. Building a great sailing vessel to reach the New World was one thing; getting it across the ocean was another. It's true in warfare as well; an army really does travel on its stomach. World War II was won not just in battle but by the farms and factories that supplied the Allied forces. This is why amateurs talk about tactics, but great generals study logistics.

All this applies to AI technology as well. Achieving the breakthroughs that now seem within reach requires both software and hardware. Someone has to *design* the hardware, *build* the hardware, and then *operate* the hardware. That last step will require stupendous amounts of energy.

You may have heard the term "LLM" or large language model. That's the technology behind ChatGPT and similar AI systems. LLMs read documents then calculate, letter by letter, word by word, how bits of information usually go together. This requires many millions of computations.

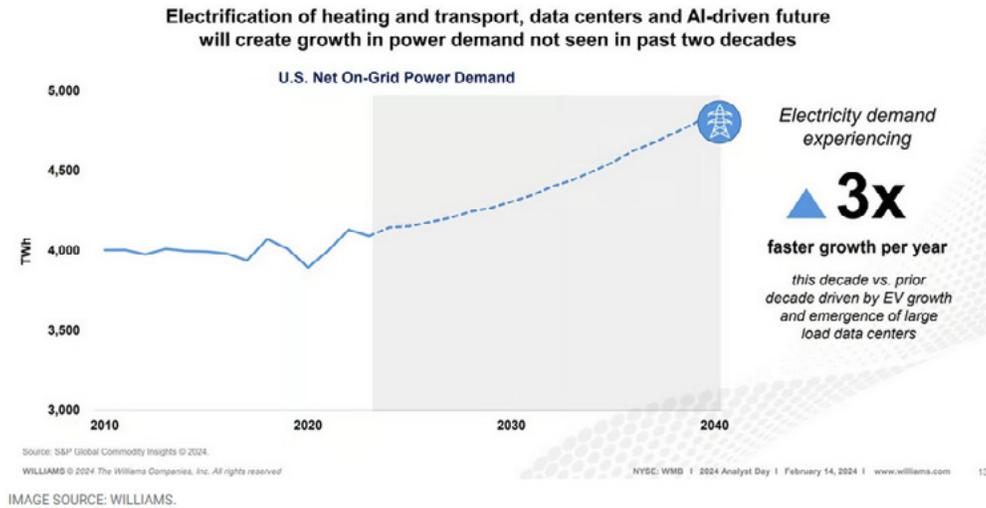
Think of it as building a sandcastle one grain of sand at a time. It's possible in theory but not in reality because you can't move fast enough to do it in one lifetime. The newest computer chips can... but *only* if they have a sufficient power source.

Note also, the power isn't just for the microchips. Running so many calculations so quickly generates heat. The disposing of it requires cooling systems which consume yet more power (and sometimes massive amounts of water—another not-cheap resource). Yet more power goes into transporting data between users and servers, that may be thousands of miles apart.

You begin to see the challenge. AI generates enormous new energy demand on top of everything else. This is pure growth, not the replacement for something that will go away—at least not yet.

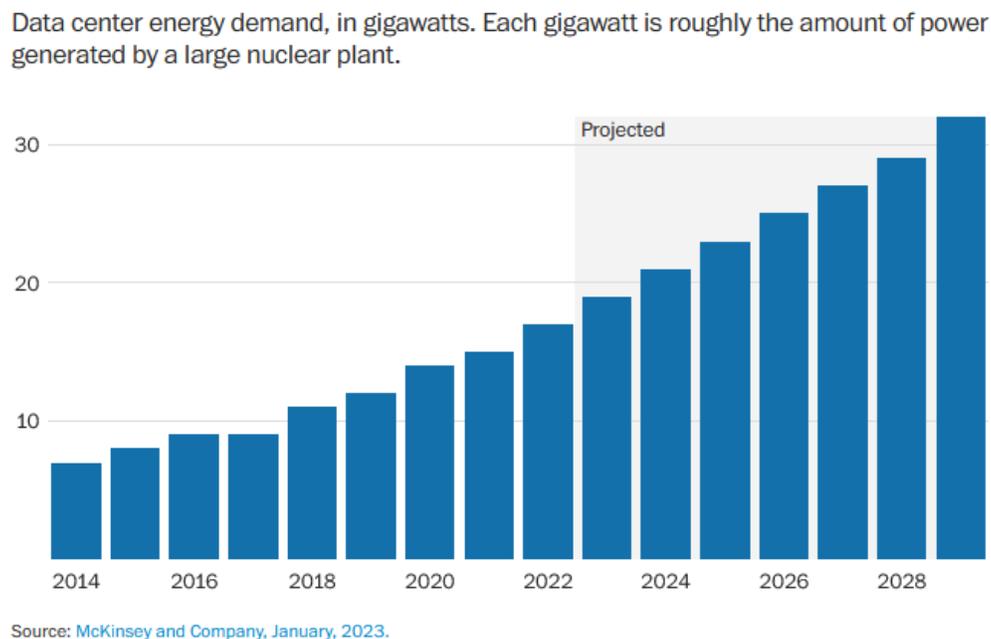
It's hard to get to Net Zero when we keep inventing new technologies that consume ever larger amounts of energy. The goalposts keep moving.

We are also shifting some energy demand from direct burning (gasoline, heating oil) to electric vehicles and heating systems. The amount consumed may not change much but it puts more strain on the electric grid. US electricity demand growth has *tripled* since a decade ago.



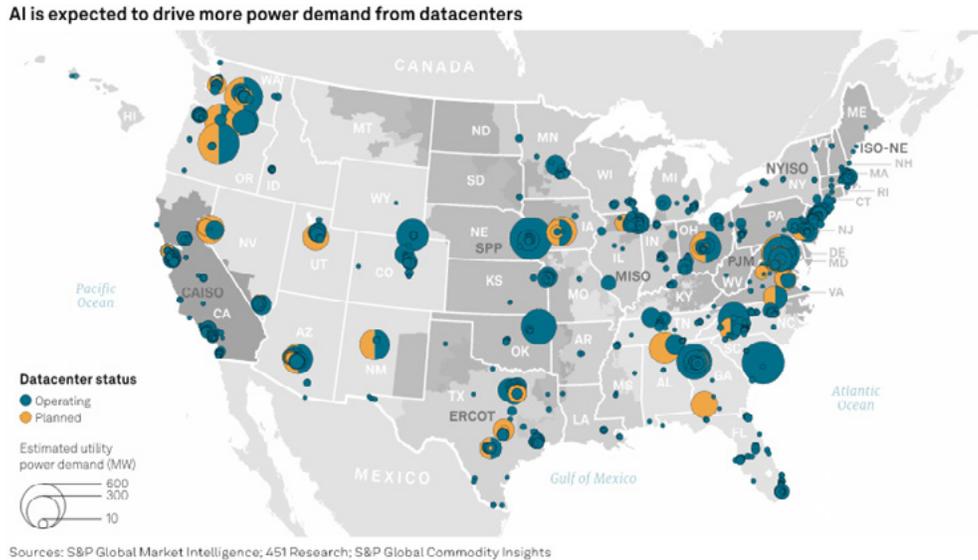
Source: Motley Fool

AI and crypto aren't causing all this demand growth but they are substantial parts of it. This chart shows past and projected data center energy usage, compiled by McKinsey last year.



Source: *The Washington Post*

Here is a map from S&P showing currently operating (blue) and planned (yellow) US data centers. The size of the circles represents their estimated power demand.



Source: Standard & Poor's

Many think AI will revolutionize, well, almost everything. I agree, but we will have to find ways to meet the new energy challenges that go with it.

The picture below is me and good friend Frank Holmes in one of his bitcoin mining data centers in Iceland (where cooling is no problem). This was the center of three aisles in one of several massive buildings. And the AI data centers dwarf this.



Source: John Mauldin

If you search for “how much energy will AI need,” you’ll find many of the articles refer to a Dutch analyst, Alex de Vries. He was an early crypto skeptic and began tracking crypto mining’s vast energy requirements.

Then in 2022, de Vries looked into AI’s energy demand. Since few of the companies doing AI would talk about their energy usage, he simply counted the number of chips Nvidia said it was shipping, how much energy they use, and then extrapolated. He estimates by 2027, powering those chips will take well over 1% of total world energy usage—not counting the cooling and other related demand.

This is from [a recent *New Yorker* piece](#) on the subject:

“There’s a fundamental mismatch between this technology and environmental sustainability,’ de Vries said. Recently, the world’s most prominent A.I. cheerleader, Sam Altman, the C.E.O. of OpenAI, voiced similar concerns, albeit with a different spin. ‘I think we still don’t appreciate the energy needs of this technology,’ Altman said at a public appearance in Davos. He didn’t see how these needs could be met, he went on, ‘without a breakthrough.’ He added, ‘We need fusion or... radically cheaper solar plus storage, or something, at massive scale—a scale that no one is really planning for.’”

In the US, data centers now account for about 4% of electricity consumption, and that figure is expected to climb to 6% by 2026.

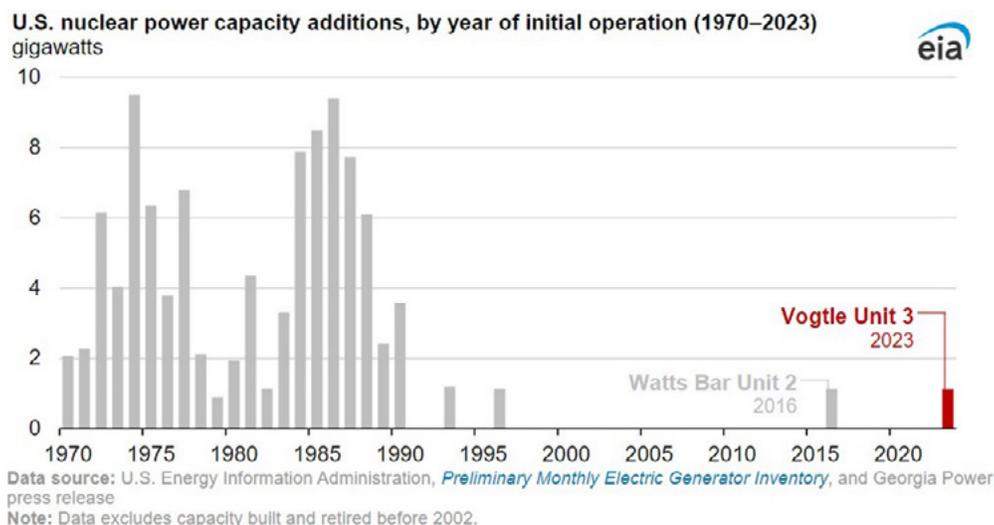
Squaring the Circle

The excitement surrounding AI is comparable in some ways to the excitement surrounding renewable energy. They might seem like a match made in heaven: boundless AI electricity demand met by clean, plentiful solar and wind power. A virtuous circle for everyone.

These renewable energy sources are indeed growing quickly. They are an important part of the mix, but they have limitations. The main one is “intermittency,” the fact that they produce electricity only when the sun shines and the wind blows. Data centers operate 24/7 and need generally constant amounts of electricity around the clock. They can’t just shut down because the winds are calm.

People are trying to square this circle. Google is running an experimental [geothermal power plant](#) in Nevada to help run two of its data centers. Others are building giant battery farms to capture excess solar and wind production for later use. This helps but not nearly enough.

As of right now, the most consistent around-the-clock clean electricity generation is nuclear power. The tech industry sees this and is getting more interested. That’s great because our nuclear power base is aging rapidly and not being replaced.



Source: [EIA](#)

Microsoft is reportedly investing in small [modular reactor technology](#). These could be located adjacent to data centers and provide the kind of steady supply that other renewable sources can't. Amazon recently moved in that direction, [buying a data center](#) next to an existing nuclear power plant in Pennsylvania.

I'm a huge nuclear power advocate, especially if we can cut through some of the bureaucracy that makes it so expensive and time-consuming to build. (See that lone 2023 addition in the chart above? That was the expansion of an existing Georgia plant which cost more than \$30 billion.)

Even better and hopefully less costly nuclear technologies like thorium reactors are coming. But none of them are coming as fast as AI power demand is growing.

All this leads to an obvious conclusion many people, particularly environmentalists, don't want to admit. We aren't going to "phase out" fossil fuels anytime soon. They are and will remain a large, critical part of global energy supplies. Nothing else has the same combination of abundant supply, well-developed infrastructure, and ability to keep producing in any weather conditions. The International Energy Agency (IEA) increases its demand forecast seemingly every year, sometimes twice a year.

If you are worried about the climate, you will likely say this isn't good. What are the alternatives, though? The world doesn't have sufficient clean energy to supply demand, nor do we have good ways to reduce demand in the amounts needed. A deep, multi-year global depression might do it but at high cost.

All this energy use is from current versions projected forward of AI. As we get newer and more data-hungry versions, and more companies emerge offering AI, and even more companies use it, that energy use is going to compound in ways that we don't understand now and we certainly aren't planning for.

This makes the aspirational dream of Net Zero even further away. We will be needing oil and gas and coal for much longer than many hope. Which is why I am still bullish on energy. And nuclear.

Much like monetary policy, this is one of those things where we don't have ideal choices. All the alternatives are bad in different ways. We can only choose the least-bad one.

But speaking of monetary policy...

Dot Plot Thoughts

We had a lot of central bank news this week. First, the Bank of Japan raised its policy rate for the first time since (checks notes) 2007! They hiked from -0.1% to 0%, thus (I hope) ending the era of negative interest rates. It was a baby step but still hugely significant. The BOJ also ended its yield curve control strategy which had been keeping long-term government bonds at 0% and will also stop buying ETFs and REITs.

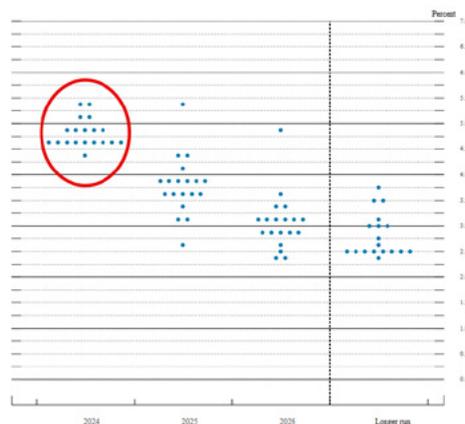
The Japanese authorities had little choice in this matter, as inflation is finally taking off there. It's nothing like we face here, but wages and some consumer prices are rising in ways unseen for many years.

Meanwhile in Europe, the Swiss National Bank became the first central bank to *cut* rates in this cycle. The European Central Bank, Bank of England, and Bank of Canada are also hinting they will cut soon. A turning point seems to be at hand, but the timing is still highly uncertain.

Back in the US, this week's Federal Open Market Committee meeting brought no policy changes but a wealth of signals to interpret. The last quarterly set of "dot plots" in December showed the median committee member expected the fed funds rate to be 4.6% at the end of 2024. That would be 75 basis points less than the current level, or three quarter-point cuts. The big question this week was whether this projection would change. It didn't, staying at 4.6%.

For release at 2:00 p.m., EDT, March 20, 2024

Figure 2. FOMC participants' assessments of appropriate monetary policy: Midpoint of target range or target level for the federal funds rate



Source: FOMC

Each dot in that circled cluster represents one member’s forecast. While the median stayed at 4.6%, the distribution did change a bit. As of December, three members thought year-end rates would be over 5%. Now it’s four. And back then, one member had their dot below 4%. Now the lowest is 4.5%, though we don’t know if this is the same member. Note that half the members see only two rate cuts or less. Two see no cuts and two see only one.

Even though Powell was “dovish” at his press conference, the committee is still cautious. The dots’ distribution shows the committee as a whole getting a bit more unified in the hawkish direction. That’s not surprising, given the latest inflation and other economic data.

The FOMC members make “economic projections” as well. They are enlightening not for any forecasting value, but because they tell us what the committee is thinking. Comparing it with prior forecasts also shows us the *evolution* of their thinking.

Here is the table summarizing their projections ([click here](#) for a larger image). The blue shaded rows are the latest guess and the white row underneath each one shows what they said in December.

Variable	Median ¹				Central Tendency ²				Range ³		
	2024	2025	2026	Longer run	2024	2025	2026	Longer run	2024	2025	2026
Change in real GDP	2.1	2.0	2.0	1.8	2.0–2.4	1.9–2.3	1.8–2.1	1.7–2.0	1.3–2.7	1.7–2.5	1.7–2.5
December projection	1.4	1.8	1.9	1.8	1.2–1.7	1.5–2.0	1.8–2.0	1.7–2.0	0.8–2.5	1.4–2.5	1.6–2.5
Unemployment rate	4.0	4.1	4.0	4.1	3.9–4.1	3.9–4.2	3.9–4.3	3.8–4.3	3.8–4.5	3.7–4.3	3.7–4.3
December projection	4.1	4.1	4.1	4.1	4.0–4.2	4.0–4.2	3.9–4.3	3.8–4.3	3.9–4.5	3.8–4.7	3.8–4.7
PCE inflation	2.4	2.2	2.0	2.0	2.3–2.7	2.1–2.2	2.0–2.1	2.0	2.2–2.9	2.0–2.5	2.0–2.3
December projection	2.4	2.1	2.0	2.0	2.2–2.5	2.0–2.2	2.0	2.0	2.1–2.7	2.0–2.5	2.0–2.3
Core PCE Inflation ⁴	2.6	2.2	2.0		2.5–2.8	2.1–2.3	2.0–2.1		2.4–3.0	2.0–2.6	2.0–2.3
December projection	2.4	2.2	2.0		2.4–2.7	2.0–2.2	2.0–2.1		2.3–3.0	2.0–2.6	2.0–2.3
Memo: Projected appropriate policy path											
Federal funds rate	4.6	3.9	3.1	2.6	4.6–5.1	3.4–4.1	2.6–3.4	2.5–3.1	4.4–5.4	2.6–5.4	2.4–4.9
December projection	4.6	3.6	2.9	2.5	4.4–4.9	3.1–3.9	2.5–3.1	2.5–3.0	3.9–5.4	2.4–5.4	2.4–4.9

Source: FOMC

Looking at the “median” section, we see the FOMC expects 2024 to end with:

- 2.1% GDP growth
- 4.0% unemployment
- 2.4% PCE headline inflation
- 2.6% PCE core inflation

Notably, their GDP projection made a big jump from 1.4% as of December while their unemployment and inflation projections didn't change much. That suggests the committee sees little or no chance of recession this year. Remember Q4 real GDP growth was at 3.2% in the second estimate (which will be updated next week). They don't foresee growth staying that high, but they also don't anticipate contraction.

At the same time, those economic projections are in the context of fed funds dropping by 75 points. Their collective belief seems to be, "We can cut rates three times and it will have almost no effect on inflation or unemployment."

When you put it that way, what's the reason to cut? Their own projections say it won't do anything beneficial.

The answer may be that the benefits aren't shown in this table. Lower rates would obviously help many borrowers, not least the US Treasury, along with many leveraged businesses and consumers. Lower rates might spur some construction activity which, if it expands housing supply, would help bring down inflation.

If the FOMC members really believe what their projections imply, then as of now at least half of them are still intent on three rate cuts by year end, though of course they will state it as assuming the data doesn't change, etc.

The election is a slight complication as they presumably won't want to be seen as influencing the outcome. But I'm not sure that will stop them; *not* cutting rates—after repeatedly signaling rate cuts were coming—could also be seen as interference.

Remember the Fed's #1 goal: Keep inflation headed downward toward their target, and certainly don't let it rise much. The projections show they think the goal (2% in core PCE) is reachable in 2026. But they also need to not keep policy so tight as to spark recession.

Can they have it both ways? Maybe. I'm one who said a "soft landing" would be tough to pull off. Two years later, it looks a lot more plausible.

Look Who's Coming to SIC!

We're at an historic crossroads of converging cycles—economic shifts, geopolitical instability, societal conflicts, and technological forces—reshaping markets against the backdrop of a critical US election. It's a perfect storm for turbulence and instability.

But despite the uncertainty, there is opportunity—SIC 2024 will highlight opportunities for growth, innovation, and resilience.

The theme of this year's Strategic Investment Conference is "***Into the Storm***." We can't avoid it. We simply must go into it and *through* it. That requires planning and, like the generals who focus on logistics, successful planning requires information. Which is what this year's SIC is all about.

Who better than Felix Zulauf to help us navigate heavy waters? Rational Optimist Lord Matt Ridley will give us a deep dive on innovation, including artificial intelligence. Seriously, I could listen to Matt for hours. His book [The Evolution of Everything](#) is one of my top five must-read books.

Leon Cooperman. Howard Marks. Who better to peer into an increasingly turbulent future? We just locked down an all-star panel of experts (whom you will not know but will want to) to tell us about dual-use technologies. Say what? Why? Because that's where the conflicts and opportunities lie. What might be a powerful technology for business can also have military uses. Who gets it first? Is there a way to control the spread?

[Click here](#) to see some of the others who will be presenting. And unlike any other year, we have kept 3–4 slots open for some VERY special guests we are working with. This 20th SIC will be my best ever. You really want to join us.

DC, Cape Town, and Europe?

I plan on being in DC in mid-April for another anti-aging drug presentation. Still a few years out, this is the most promising near-term drug I have seen in 20 years of research. It seems to be actually lengthening telomeres. We are trying to get something we can present at the SIC, along with the other latest health and biotech developments.

Shane and I will be in Cape Town for a speech in early June and then somewhere in Europe.

The letter is already too long, so I will hit the send button. Have a great week. It is starting to get warm and soon will be hot and humid here in Puerto Rico after a very mild winter. The Atlantic seems to be warming up, so maybe a more “active” hurricane season awaits? A little rain in paradise may be in our future. And don't forget to [follow me on X!](#)

Your down to 185 pounds analyst,



John Mauldin

subscribers@mauldineconomics.com

**Keep up with Mauldin
Economics on the go.**

Download the App



<http://www.mauldineconomics.com/members>

© 2024 Mauldin Economics. All Rights Reserved.

Thoughts from the Frontline is a free weekly economic e-letter by best-selling author and renowned financial expert, John Mauldin. You can learn more and get your free subscription by visiting www.MauldinEconomics.com.

Any full reproduction of Thoughts from the Frontline is prohibited without express written permission. If you would like to quote brief portions only, please reference www.MauldinEconomics.com, keep all links within the portion being used fully active and intact, and include a link to www.mauldineconomics.com/important-disclosures. You can contact affiliates@mauldineconomics.com for more information about our content use policy.

To subscribe to John Mauldin's Mauldin Economics e-letter, please click here: <http://www.mauldineconomics.com/subscribe>

To change your email address, please click here: <http://www.mauldineconomics.com/change-address>

Thoughts From the Frontline and MauldinEconomics.com is not an offering for any investment. It represents only the opinions of John Mauldin and those that he interviews. Any views expressed are provided for information purposes only and should not be construed in any way as an offer, an endorsement, or inducement to invest and is not in any way a testimony of, or associated with, Mauldin's other firms. John Mauldin is the co-founder of Mauldin Economics, LLC. He also is the President and investment advisory representative of Mauldin Solutions, LLC, which is an investment advisory firm registered with multiple states, President and registered Principle of Mauldin Securities, LLC, a FINRA and SIPC, registered broker-dealer. Mauldin Securities LLC is registered with the NFA/CFTC, as an Introducing Broker (IB) and Commodity Trading Advisor (CTA).

This message may contain information that is confidential or privileged and is intended only for the individual or entity named above and does not constitute an offer for or advice about any alternative investment product. Such advice can only be made when accompanied by a prospectus or similar offering document. Past performance is not indicative of future performance. Please make sure to review important disclosures at the end of each article. Mauldin companies may have a marketing relationship with products and services mentioned in this letter for a fee.

PAST RESULTS ARE NOT INDICATIVE OF FUTURE RESULTS. THERE IS RISK OF LOSS AS WELL AS THE OPPORTUNITY FOR GAIN WHEN INVESTING IN MANAGED FUNDS. WHEN CONSIDERING ALTERNATIVE INVESTMENTS, INCLUDING HEDGE FUNDS, YOU SHOULD CONSIDER VARIOUS RISKS INCLUDING THE FACT THAT SOME PRODUCTS: OFTEN ENGAGE IN LEVERAGING AND OTHER SPECULATIVE INVESTMENT PRACTICES THAT MAY INCREASE THE RISK OF INVESTMENT LOSS, CAN BE ILLIQUID, ARE NOT REQUIRED TO PROVIDE PERIODIC PRICING OR VALUATION INFORMATION TO INVESTORS, MAY INVOLVE COMPLEX TAX STRUCTURES AND DELAYS IN DISTRIBUTING IMPORTANT TAX INFORMATION, ARE NOT SUBJECT TO THE SAME REGULATORY REQUIREMENTS AS MUTUAL FUNDS, OFTEN CHARGE HIGH FEES, AND IN MANY CASES THE UNDERLYING INVESTMENTS ARE NOT TRANSPARENT AND ARE KNOWN ONLY TO THE INVESTMENT MANAGER. Alternative investment performance can be volatile. An investor could lose all or a substantial amount of his or her investment. Often, alternative investment fund and account managers have total trading authority over their funds or accounts; the use of a single advisor applying generally similar trading programs could mean lack of diversification and, consequently, higher risk. There is often no secondary market for an investor's interest in alternative investments, and none is expected to develop. You are advised to discuss with your financial advisers your investment options and whether any investment is suitable for your specific needs prior to making any investments.

All material presented herein is believed to be reliable but we cannot attest to its accuracy. Opinions expressed in these reports may change without prior notice. John Mauldin and/or the staffs may or may not have investments in any funds cited above as well as economic interest. John Mauldin can be reached at 800-829-7273.