Q2 2026

Winners, losers, and the law of large numbers

In one NFL season, we bet $20,000 to win $125.

Right after college, a couple of my best friends and I decided we'd get rich building a model that could beat the NFL betting markets.

We spent six months of the offseason building and refining a neural network* that could take inputs ranging from player injuries to how far away teams were flying and predict the outcomes of the games.

Now, in order to make money betting on NFL spreads and totals, after you take into account the sportsbook's take, you need to get it right roughly 52% of the time. Our model, after months of work, suggested we'd be right...54% of the time.

We each chipped in a couple hundred dollars week one, and each weekend we followed the model (probably ruined the fun of gambling for me), bet on nearly every game, and recycled our winnings into the next week. Some weeks we made tons of money (relatively), some weeks we lost even more; the swings were large for kids eating Domino's.

In the end, we placed over $20 thousand across hundreds of bets and won just over 53% of them. Good for $125 and about $0.75 per hour of my time.

More importantly, it was a lesson in the volatility of single bets or even single weeks, and the impossibility of evaluating whether one is right or not with too small a sample. It seems like the right time to revisit this lesson.

Looking at the market today, there appear to be a lot of winners and losers being crowned well before it's possible to know the truth, and probabilities have gone out the window.

To us this seems a mistake.

"The transaction we are announcing today represents a strategic proactive move to optimize our financial flexibility and maximize long-term shareholder value creation."

- Anat Ashkenazi
Alphabet CFO

** In this letter it may sound as if I'm incredibly bearish on AI, while in fact the opposite is true and has been since this story, well before ChatGPT.*

Smart CFOs want to make smart decisions.

On June 1st, Alphabet announced it was raising $80 billion in equity from investors, the largest raise in company history, to fund their AI plans. The raise included common shares, a private placement to Berkshire, and convertible preferred shares with a 6.25% dividend.

Meanwhile, earlier this year, the company issued roughly $32 billion in debt, with coupons ranging from 3.7% for 2029 debt to 6.125% for a *100-year* bond, with the reported order book north of $100 billion.

Unlike Oracle (which we spoke about in our year-end letter) it's clear Alphabet had a path to raise the capital it needed from the debt markets. And also unlike Oracle, it could do so at spreads that effectively priced it as riskless.

So, on the back of issuing billions of dollars of debt at the risk-free rate, the Alphabet team made the decision to tap the equity markets for incremental capital.

Anat Ashkenazi and Sundar Pichai aren't dumb.

The announcement rationale was summarized as a "… plan to fund our investments in a balanced way." In our view the rationale for issuing this much equity largely boils down to one of the below:

- Lack of access to debt markets, which we know was unlikely;
- They have concerns about the likelihood of payoff relative to the fixed nature of debt obligations;
- The decision that on a risk-adjusted basis, the cost of equity is lower than the cost of debt.

So, the real question is what makes a smart management team unwilling to pledge fixed obligations against its own flagship investment. We think there are two answers, and they compound.

***On average,* incremental hyperscaler ROIC is likely to be lower.**

The case for spending this much rests on the idea that the capital earns software-like returns. Oracle's management talks about return on invested capital "in the high 20s at a steady state," and the sell side models 25% to 30%.

Those are returns on the integrated platform: the software, the distribution, the ads layered on top. They are not returns on the concrete and the silicon underneath.

We can observe what the concrete and silicon actually return, because two large companies do almost nothing else. Equinix and Digital Realty own and operate data centers for a living. And their returns aren't software returns.

5Y Average UAFRS ROA

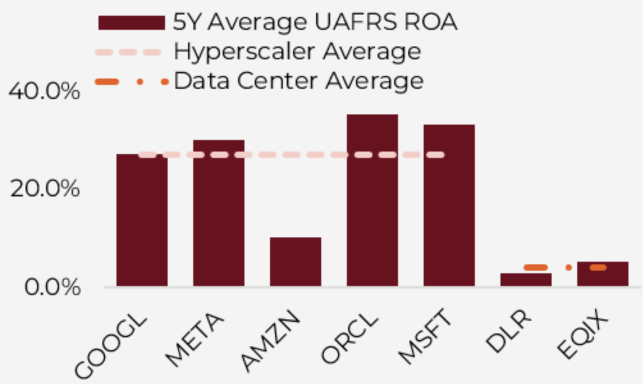

Source: Valens Research; Monas Investments

The above graphic highlights the last five years of ROIC for the hyperscalers versus the data center businesses they are building. Now, not all the investment is going into the construction and hardware, but it's a majority (estimates are for around 75% into new AI infrastructure).

The development math looks better on a cash-on-cost basis, but the all-in return on the full capital base, especially if funded by equity issuances at software multiples, appears to land in the low single digits.

Meanwhile, assets are short-lived, and prices are going the wrong direction.

Microsoft's CFO Amy Hood mapped out their spend as two-thirds GPUs and CPUs, with the remaining one-third going toward the rest.

The challenge is the useful life of these chips is roughly five years (six for the more generous estimates). If every five years we need to spend billions (or trillions) to replace and upgrade, the ROIC math becomes much less attractive. This isn't the same R&D investment that has me paying for Microsoft Word over 40 years after it was created.

For the chips to generate expected return on investment, we would expect to see significant pricing power and competitive advantage within these frontier-model companies. Unfortunately, to date it appears the opposite is the case.

OpenAI recently announced it was deliberating significant pricing cuts on API tokens and AI services, and although the frontier models continue to get more expensive over time, legacy models (which are plenty sufficient for most work today) continue to get cheaper.

GPT-4 Family Input Price Per Token

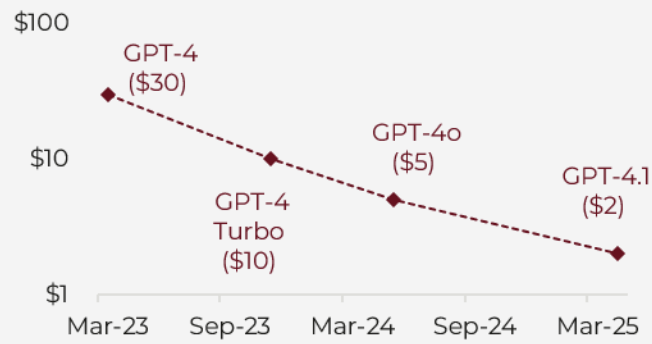

Source: TechCrunch; OpenAI; Monas Investments

As a regular user of these tools, it is already clear that smart corporate finance capital allocators will be figuring out ways to route as much work to as cheap a model as suffices, and that in most scenarios, these tools are commodities.

There will be a winner, but plenty of losers.

The spend makes sense under one condition: this is a winner-take-most market and you are the one who wins it. And they are spending like it. In 2026, capital expenditure is forecast at 70-110% of operating cash flow at Alphabet, Meta, Amazon, and Microsoft, and north of 150% at Oracle. These firms are investing at a rate that suggests an existential bet, not an IRR-based allocation exercise.

This looks less like an infrastructure buildout than a venture bet: hope one firm gets it right, builds the model that matters, embeds it into the products and workflows people actually run on, and earns a long-term return on everything it poured into the build.

It will work out for one of these companies. It also, like venture, will be a bet where most of the capital deployed does not come back.

Speaking of winners and losers…

This quarter saw the launch of the final boss battle between SpaceX/Elon and retirement savers.

On June 12th, SpaceX (SPCX) went public, priced at $135 per share, and closing at ~$161 that day. Closing at a valuation north of $2 trillion, it was the largest IPO ever.

Plenty has already been written on the merits of the company as an investment:

Damodaran weighed in.

Our friends at Valens wrote a three-part series.

Chamath wrote a piece. Believe it or not, you have to pay for it.

Our take is simple. Since the market bottom in 2009, by our calculation there have been 72 S&P 500 companies that at one time traded above 30x Forward Sales (SPCX is at ~50x and *100x trailing sales*), and 258 instances of one of these companies trading at that multiple at the end of a quarter. Across this sample, just 35% outperformed the S&P over the next 3 years, averaging 5% of annual underperformance.

The company's stated goals for future markets include manufacturing facilities in space, energy production on Mars, and asteroid mining.

Their defined TAM for Starlink is "every global household".

Their TAM for AI includes estimates for the entirety of the global "digital economy".

These things require stewardship and capital.

Meanwhile, due to the dual share class, even with ~46% ownership of SPCX, Elon retains ~85% of the voting power, and his track record of respecting minority shareholders is wanting.

Source: CNBC

Needless to say, there is a lot of hope embedded in the price today; the poster child for an early-crowned-winner of which we learned to be wary.

Fortunately for SpaceX investors, thesis matters less than investment flows.

For most of index-fund history a company had to season before it could join a major benchmark. It would list, trade for months, build liquidity, get through its lock-up expirations, and only then become eligible. The seasoning rule existed precisely so that index investors, who must buy whatever the benchmark tells them to, were not forced into a thin, freshly minted stock in the chaotic window right after it debuts.

Well, Nasdaq rewrote its methodology effective May 1, 2026: a new listing among the 40 largest can join the Nasdaq-100 after 15 trading days rather than roughly three months, and it scrapped its 10% minimum-float requirement, replacing it with a rule that lets a low-float stock be weighted at up to three times its actual float.

FTSE Russell followed on May 26th, allowing a large enough IPO into the Russell indexes after the fifth trading day instead of at the next quarterly reconstitution, and companies with less than 5% float can qualify if lock-up expirations are expected to lift them above minimum within 12 months.

For now, the impact is smaller than the headline suggests, because index funds weight a company by the shares actually available to trade, not its $2 trillion sticker. With only about 5% of the company floating, SpaceX enters as a rounding error: roughly a 0.1% weight in the CRSP Total Market Index and around 0.7% in the Nasdaq-100.

Even so, back of the envelope math suggests an estimated $22 to $27 billion of forced "passive" buying against roughly $85 billion of stock actually floating. And as lockups expire, and float rises, there is a natural forced buyer of shares as insiders exit.

None of this is a conspiracy. The case for the rule changes is reasonable: an index is supposed to reflect the market, and a $2 trillion company is now one of the largest things in it. Leaving it out for a year or more arguably makes the benchmark less representative, not more.

The problem is not the logic, it is the buyer. The flows we are describing are price-insensitive by design. The funds on the other side are 401(k)s, target-date funds, and the model-portfolio ETFs of people who have never heard of an investable weight factor, who will now sell little slices of everything they own, to buy stock in an unprofitable, 100x-sales stock on a schedule.

Maybe Warsh comes to the rescue.

We've written in our last few letters about the four legs of the market. In the year-end:

"For 2026 to look like 2025…we likely need to see four different legs of this bet pay off:

- *Growth must stay high*
- *Inflation must remain muted*
- *Rates must fall*
- *Broader geopolitical risk must cool"*

Well, mixed bag so far, but 2026 looks pretty similar to 2025.

Source: S&P CapitalIQ

Meanwhile, inflation has been elevated (largely due to energy prices), but expectations are down significantly over the last month. The market expects the energy spike to be short-lived.

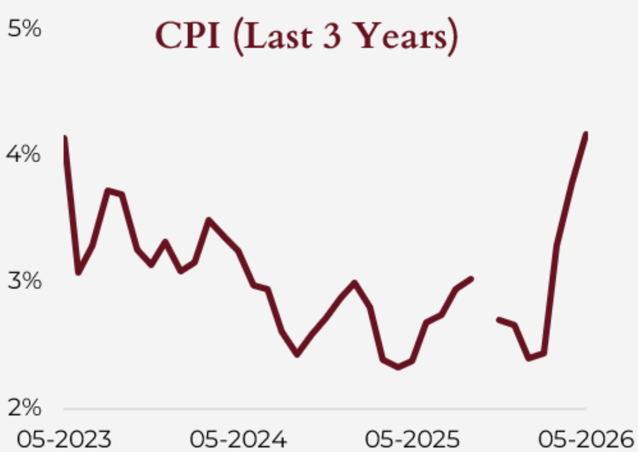

Source: FRED

Source: FRED

Maybe an open (ish) Strait of Hormuz means we're back to business as usual.


Iran Attacks Cargo Ship, Testing Trump's Deal to Reopen Strait

Attack takes place hours after Iranian warning to ships not to use routes that the regime hadn't sanctioned

By Shelby Holliday and Rebecca Feng
Updated June 25, 2026 7:40 pm ET

Source: Wall Street Journal

And yet, short-term rates have risen significantly over the same time.

Source: FRED

Meanwhile, expectations have flipped from cuts this year, to a potential hike.

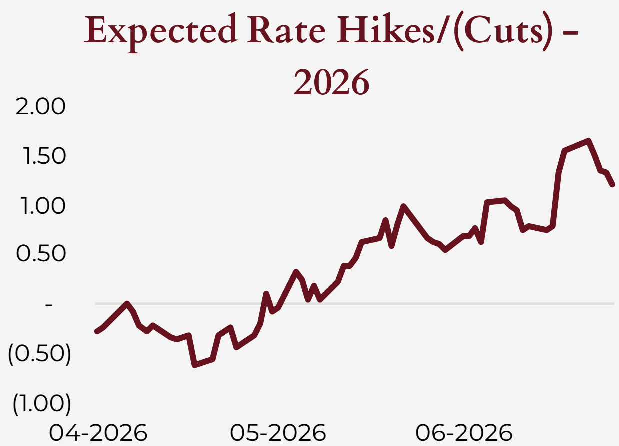

Source: CME Group, Monas Investments

Kevin Warsh took his seat as Fed Chair on May 22nd, with what would seem to be a mandate to get rates lower. This administration's plans to bring down rates is well documented, and he was picked to be the man to do so.

However, there is clearly uncertainty about his ability, if not willingness, to bring rates down.

He has said all the right things.

Warsh spent the last year building the case for himself. In a Wall Street Journal op-ed in November he called AI "a significant disinflationary force," drew the parallel to the late-1990s productivity boom, and argued the Fed had room to take rates lower to support households and smaller businesses. He has spent years arguing the Fed talks too much, that the dot plot and forward guidance trap the committee into defending stale forecasts long after the data has moved.

He is also about as wired into this administration as a Fed chair gets. He and Treasury Secretary Bessent go back two decades. Both learned under Stanley Druckenmiller, and Bessent now has him to breakfast every week, the same standing slot he kept with Powell. Bessent pushed the confirmation through a Senate blockade to get him the chair.

The Treasury wants lower rates to fund the deficit. The President wants them for every reason he has ever given. The chair wrote the intellectual case for them.

The table is set. Warsh to the rescue.

Except it isn't just up to Warsh.

Three things sit between wanting to cut and getting it done.

First, the committee around him just got more hawkish.

The Fed's annual rotation seated four new regional presidents as voters in January, and three of the incomers (Dallas's Lorie Logan, Cleveland's Beth Hammack, and Minneapolis's Neel Kashkari) are among the most inflation-focused people on the committee. They replaced a more dovish 2025 group (and will in turn largely be replaced in 2027).

Name	Position	Dove/Hawk Scale	2026 Voter	2027 Voter	2028 Voter	Board of Gov. / (Vice) Chair term ending *	Appointing President
Christopher J. Waller	Board of Governors		Yes	Yes	Yes	Jan30	Donald Trump
Michelle W. Bowman	Vice Chair Supervision		Yes	Yes	Yes	Jan34 / Jun29	Donald Trump
John C. Williams	FRB of New York		Yes	Yes	Yes		
Kevin Warsh	Chair		Yes	Yes	Yes	Jan40 / May30	Donald Trump
Philip N. Jefferson	Vice Chair		Yes	Yes	Yes	Jan36 / Sep27	Joe Biden
Lisa D. Cook	Board of Governors		Yes	Yes	Yes	Jan38	Joe Biden
Anna L. Paulson	FRB of Philadelphia		Yes	No	No		
Jerome H. Powell**	Board of Governors		Yes	Yes	Yes	Jan28 / May26	Barack Obama
Michael S. Barr	Board of Governors		Yes	Yes	Yes	Jan32	Joe Biden
Neel Kashkari	FRB of Minneapolis		Yes	No	No		
Lorie K. Logan	FRB of Dallas		Yes	No	No		
Beth M. Hammack	FRB of Cleveland		Yes	No	Yes		
Mary C. Daly	FRB of San Francisco		No	Yes	No		
Susan M. Collins	FRB of Boston		No	No	Yes		
Thomas I. Barkin	FRB of Richmond		No	Yes	No		
Raphael W. Bostic	FRB of Atlanta		No	Yes	No		
Alberto G. Musalem	FRB of St. Louis		No	No	Yes		
Austan D. Goolsbee	FRB of Chicago		No	Yes	No		
Jeffrey R. Schmid	FRB of Kansas City		No	No	Yes		

Source: BBVA

Now, a chair has never been outvoted by his committee, but it is hard to believe cuts are a given, especially with the hawks going out with statements such as:

"Inflation is taking too long to return to the FOMC's 2 percent target…[and] appears to be trending toward the mid 2's—not all the way back to 2 percent." (Logan, June 3)

"For today, it's reasonable to keep rates steady given the uncertainties around the economic outlook. But if recent trends continue, it may soon be appropriate to act." (Hammack, June 2)

"In March, I had penciled in one rate cut by the end of the year. In June, I've changed that to one rate hike by the end of the year." (Kashkari, June 26)

Second, the disinflation story has yet to show up in the numbers. In fact, companies have been rehiring after waves of AI-supported layoffs over the last couple of years, and in prior letters, we've shown the productivity relative to prior innovation cycles – the data has not changed.

Third, even a cut at the front doesn't fix the rest of the curve.

Even in a scenario where he is able to push cuts, the middle and tail of the yield curve will be much more unwieldy.

The middle of the curve is being pushed the other way by the wave of hyperscaler issuance we described earlier in this letter, hundreds of billions in new bonds competing for the same money.

Estimated Cumulative Debt Issuance ($B)

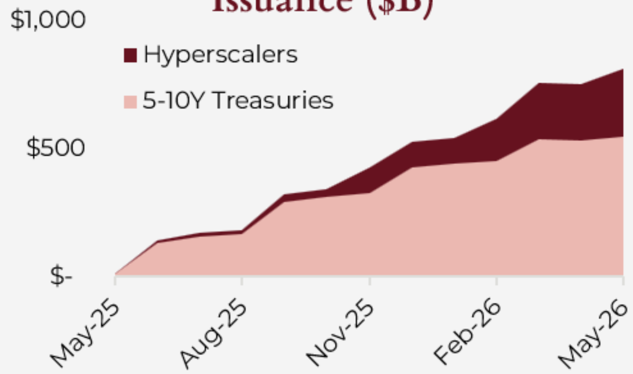

Source: S&P CapitalIQ; SIFMA; Monas Investments

The long end is being pushed by the geopolitical risk that still won't resolve, with the term premium now positive for the most consecutive months in over ten years.

Term Premium

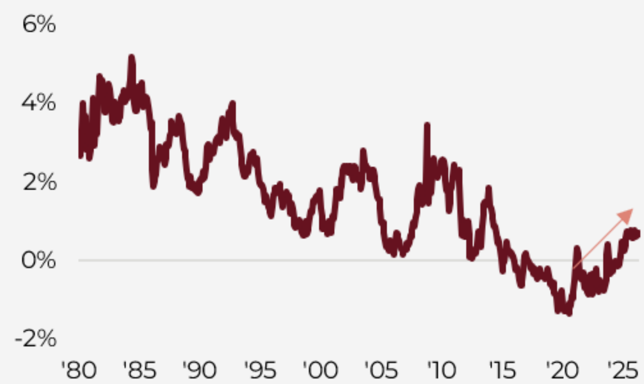

Source: New York Fed

And while credit cards, savings yields, floating debt are all wired to transmit policy changes quickly, a significant portion of the economy (mortgages and CRE financing, corporate capex funding, pension and insurance balance sheets etc.) is still tied to that longer-term debt, and transmission might not be as quick as desired.

So where does this leave investors?

Mixed signals, eye-watering valuations, and continued uncertainty around rates and where that leaves the economy. Fortunately, the investing world is massive beyond US-based 60/40 portfolios, and for every dollar capitalizing a consensus idea, less consensus ideas grow more attractive.

Here are two we're thinking about as we enter the back half of the year.

70% of the South Korean stock market is the same AI trade, and the other 30% is getting more attractive.

We calculate 54% of the MSCI South Korea Index (by weight) can be defined as explicitly the AI trade (Samsung, SK Hynix etc.) while 16% is adjacent (power and grid, cooling, components etc.), and just 30% is unrelated.

Over the last year, the cap-weighted AI theme is up over 400%, while the non-AI names are up a still strong, but significantly lower ~30%.

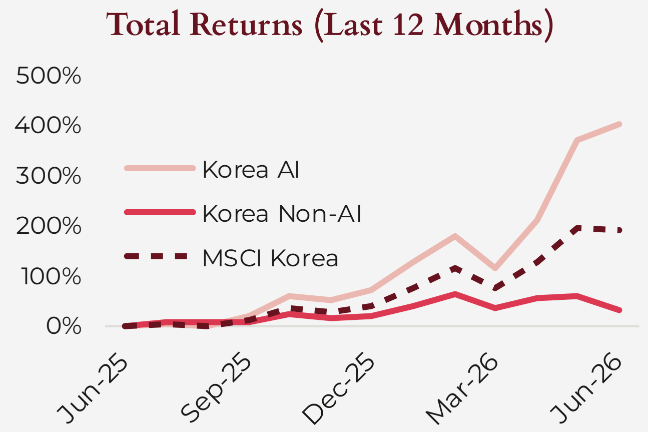

Source: S&P CapitalIQ; Monas Investments

Much like Japan, Korea has had a shareholder governance revolution in the last few years, as investors have sought to unlock shareholder value captured inside a universe of chaebols and state-owned organizations.

That value is no longer present at the top of the index, with just 3 of 19 of the "AI" names trading below 1.0x book (and one of those three is trading at 0.997x). However, in the rest of the universe, that opportunity is still widely apparent.

Source: S&P CapitalIQ; Monas Investments

Unlike Japan, this governance uplift is still in incredibly early phases.

First, there have been a number of positive changes over the last year.

July 2025's Commercial Act amendment expanded directors' fiduciary duty from "the company" to its shareholders explicitly, mandated fair treatment of all shareholders, and raised the independent-director minimum from a quarter to a third of the board. A controlling family that engineers a transaction at minority expense now faces a statutory duty it can be sued under.

With the second amendment, effective this year, cumulative voting was introduced for board votes. Simply, rather than each board seat being a separate vote, shareholders get votes equal to shares times the number of open seats, and can pile all those into one candidate. This lets minority investors concentrate their votes, and generally get more say on at least one director.

That same amendment also capped a controlling family's vote at 3% when electing the audit committee, so in theory the watchdogs are not hand-picked by those being watched.

Finally, a new rule was introduced where dividends from "high-payout" companies are taxed separately at a lower 14-30% band, rather than 45% under the old regime, applying to companies paying out a significant portion of profits with a disclosed Value-Up plan. Wealthy families are now aligned to distribute the cash they generate, rather than hoard.

However, there remains significant work to be done, and investors are right to place some discount on whether more value can be generated.

Most recently, MSCI again declined to upgrade Korea beyond emerging market status, with continued concerns about investor access and governance.

Value-Up is a voluntary program, with tax incentives that are opt-in by design. Japan's re-rating came from the Tokyo Stock Exchange folding "improve capital efficiency or explain why not" into its listing rules. The Commercial Act gives Korea legal teeth Japan lacked, but the disclosures themselves remain optional, and adoption thins fast below the largest names.

Finally, in what we would argue is the biggest issue, the estate-tax misalignment is untouched. By our measure 44/58 of the "Non-AI" companies in the MSCI index are family or founder controlled, and inheritance tax reaches 50-60% for controlling shareholders, among the highest in the world. Aging controllers facing succession planning have billions of dollars of incentives to keep share prices *low* for the next generation, directly contrary to the rest of the shareholders.

While "duty" is now law, we know smart CFOs, and owners, make smart decisions. We remain positive on the outlook as the country becomes more shareholder friendly, but it is clear this is a long-term trade while AI is the immediate driver.

If rates remain higher for longer, negative duration assets remain attractive.

Floating rate debt, IO Strips, and Mortgage Servicing Rights are all beneficiaries of higher or rising rate environments.

Earlier this quarter, we met with a team who has built a Master Servicing firm, and is building a portfolio of Mortgage Servicing Rights that should benefit in this environment.

In terms of an imbalance, trade volumes in MSRs continue to run at roughly 2x the pre-COVID average, as banks and non-bank originators have had to sell to address capital pressures. Meanwhile, high rates keep prepayment speeds low, while resilience of homeowners in this economic environment has kept defaults favorable as well.

The result is unlevered yields on broker-syndicated deals of roughly 8%, down slightly but still comfortably ahead of most public fixed income solutions.

Meanwhile some of the floating rate solutions that were attractive last year and prior, particularly high-grade CLOs, are now trading at first-percentile tights, making them relatively less attractive. In terms of capital imbalance, the flows into the space have been significant, and driven at least in part by ETF introduction and model-portfolio passivity.

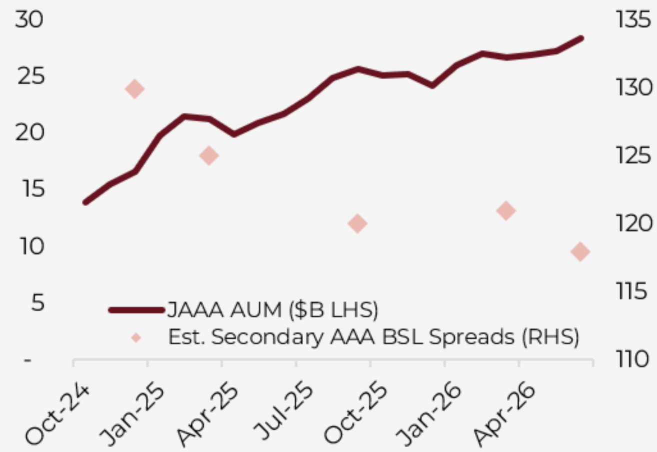

Source: Janus Henderson; S&P Capital IQ; JPMorgan; Monas Investments

It is possible this makes CLO equity tranches more attractive over the next several years, but in our view the return is more attractive elsewhere for similar economic risk. This letter is bearish enough, we can leave that discussion for another day.

Smart people want to make smart decisions.

Alphabet issues equity when it is less expensive than debt. Korean families will focus on minimizing a massive inheritance tax. Corporate finance teams will allocate to the cheapest AI models that solve their problems.

The exception is the largest pool of capital in the market today, which doesn't understand it is even making a decision. The passive bid does not weigh price against value; it buys iterations of the same model portfolios, on the schedule the methodology sets, even if that is a 100x-sales rocket company fifteen trading days off its IPO. When the marginal buyer is price-insensitive by design, we are incentivized to figure out where that capital is not.

The themes we discussed today, and have written about over the last few letters, all share the same shape: scarce capital being outpaced by demand.

Late-stage shareholder liquidity, where 800-plus unicorns sit in an IPO waiting room (and seeing OpenAI delay their IPO makes this even more attractive).

Biotech, four years into a repricing the rest of the long-duration complex never had to suffer.

The lower middle market, where the spread between large and small deals stays wide and the succession wave keeps handing the patient operator more to buy.

And electrons and infrastructure, the power and physical plant underneath the AI build, where demand is inelastic, supply is genuinely scarce, and the owner gets paid regardless of which model ends up winning. The ROIC problem we discussed earlier is a price problem, not an asset problem: funding heavy, depreciating assets at a hyperscaler's equity valuation is what breaks the math, while there remain opportunities to access this theme at infrastructure-like valuations.

The newer ideas here, Korea outside the AI complex and the negative-duration book, are the same trade in different clothes: own what the crowd is structurally forced to ignore.

Having learned our lesson early on the volatility of single events, and the impossibility of claiming victory on small sample sizes, our goal is to identify where probabilities are on our side, and make enough bets for it to matter.

We are not smart enough to call the Fed, the Middle East, the winner of the model wars, or whether we will have data centers in space. But we do believe we are capable of finding asymmetry when most aren't looking.

Disclosures & Important Information

This material has been prepared by Monas Investments LLC ("Monas") for informational purposes only. Monas is not registered as an investment adviser with the U.S. Securities and Exchange Commission or any state securities authority. Nothing contained herein should be construed as personalized investment advice, an offer to sell, a solicitation to buy, or a recommendation regarding any security, investment product, or strategy. Any references to funds, securities, or managers are provided solely to illustrate research perspectives and analytical frameworks. Recipients should consult their own legal, tax, and financial professionals prior to making any investment decisions.

This communication is provided on a general basis and does not consider the specific investment objectives, financial situation, or needs of any particular individual or entity. Advisory or consulting services, where provided, are delivered pursuant to a separate written agreement and only to clients who have engaged Monas formally. Information contained herein is subject to change without notice.

Any performance figures or examples presented may be estimated, preliminary, unaudited, and subject to revision. Past performance is not indicative of, and does not guarantee, future results. All investments involve risk, including the potential loss of principal. Benchmarks and indices are shown for comparison purposes only; they may not be directly comparable to any strategy discussed and are unmanaged and not investable vehicles. Returns described may not reflect fees, expenses, or other factors that would reduce actual client returns.

Certain information has been obtained from third-party sources believed to be reliable, but Monas does not guarantee its accuracy or completeness. Opinions reflect the firm's judgment as of the date of publication and are subject to change without notice.

This document may contain forward-looking statements, projections, forecasts, opinions, or assumptions. These reflect current views as of the date indicated but involve known and unknown risks and uncertainties. Actual outcomes may differ materially from those expressed or implied. No guarantee can be made regarding future performance, market conditions, fund execution, liquidity, valuations, deployment pacing, sourcing success, exit environments, or macroeconomic trends.

Investments in private markets, alternative strategies, and illiquid assets involve heightened risk, including but not limited to valuation uncertainty, leverage, liquidity constraints, extended time horizons, and the potential for capital loss. Strategy examples and case studies are hypothetical and provided for illustrative purposes only.

By reviewing this document, you acknowledge and agree to these terms.