

Memo 4: All-In, All-Out

August 2026



Vision & Mission

Our vision is that public markets are both beautiful and broken. Beautiful because they give us full access to an opportunity set that contains many of the world's strongest businesses. But broken because human behavior and incentives have caused most market participants to think short-term and avoid volatility. Mix these two together, and we're still left with beauty, as these dynamics allow us to pursue the work we love while having the ability to grow our capital with duration and liquidity.

Our mission is to show that there's incredible opportunity in public markets. We want to show what I've experienced firsthand: that growing our capital differently isn't about building small armies of analysts led by risk sheet-wielding decision-makers who feel compelled to maximize daily and monthly returns above all else. And to show instead that winning in public markets with longevity is increasingly about small teams or sole practitioners who are highly trained in the art of extreme simplicity that most can't stomach for very long.

Ethos


We are challengers. On the investing side, we challenge commonly held views around public market efficiency and volatility. We believe public markets have become less efficient over time. And that we might see more inefficiencies around larger, highly liquid companies with lots of "eyeballs" on them than smaller companies. For example, the 1-5-year share price charts of some of the largest companies on the planet like Netflix, Meta, NVIDIA, Oracle, UnitedHealth, and Taiwan Semiconductor do not suggest "fully understood businesses where all information has been priced in and prior consensus scenario probabilities have been amazingly accurate." Instead, the returns and share price movements of an increasing number of companies tell a story of a market that has become dominated by firms focused on short-term decision-making, volatility smoothing, or tracking benchmarks. And there are many other common views around investing in public markets and winning firm and team structure that we challenge today. Our ethos is that we should embrace the pursuit and hunt for greatness, and that in public markets, there needs to be far more seats at the table for those with the right ingredients to be exceptional investors.

2Q 2026 Thoughts: All-In, All-Out

2Q 2026 continued to be what we call an all-in, all-out market. With the rise of passive, quant / multi-strategy pods, and the AI revolution all happening at the same time, we have unprecedented market confusion meeting unprecedented market short-termism head-on. This creates a bomb of volatility, which we can see inside our own portfolio. It has become more common than not in recent months for >70% of the holdings in our portfolio to move plus or minus 2% daily. Over the last 12 months, >80% of the >200 companies we cover had *both* a 20%+ rise *and* fall within the period.

Exhibit 9. The index hid the violence underneath

Price swings inside our coverage universe, trailing 12 months ended June 30, 2026.



Source: Denmark Capital analysis of >200 public companies in its research universe, trailing 12 months ended June 30, 2026. Uses peak-to-trough and trough-to-peak trading ranges.

These dynamics, which we describe in our recent essay, [The Two Clocks](#), all drive towards more volatility rather than less, have increased both the frequency of the mispricing events we look for around the businesses we admire most and the speed with which market narratives can drive valuation multiples higher and lower than we would expect, within shorter periods of time than we would expect. It is a different type of market regime, and we believe the right investment philosophy, structure, temperament, and overall flexibility and discipline around position sizing presents a lasting opportunity to take advantage of its tendencies.


In our last memo, we discussed the two wars that were raging at the same time: the Moat War and the Iran War. While it remains to be seen whether the Iran debacle will resolve anytime soon, it seems to us that interested parties have shown a willingness to avoid a return to the full-on bombings that we saw earlier in the year, and Iran has stopped short of actions that would draw Israel back into the war as of now. We'll continue to watch with a particular focus on the inflationary effects, but our bigger focus is currently the Moat War. The Moat War is still being fought on a daily basis and could continue for so long as the slope of AI progress remains steep, which we think could be a while.

The second quarter featured a face-ripping rally in almost everything semis and AI infrastructure, including our holdings in these areas. Outside of semis and AI infrastructure though, it was difficult to make money almost anywhere else through June. It has been a narrow market through 2Q, with your exposure to semis and AI infra most likely dictating your quarterly and year-to-date absolute and relative performance. While the market continues to take its all-in, all-out approach, we have more of a pie expansion view. We are both positive on AI and the productivity it is and should continue unleashing, and on other types of business models that will no doubt be affected by AI and need to continue innovating but should have exciting growth opportunities ahead despite not being preeminent AI model builders or those directly downstream from them.


Here is what we will remember most about 2Q 2026:

- **Semis.** The Philadelphia Semiconductor Index (SOX) (contains the 30 largest eligible U.S.-listed semiconductor companies) rose nearly 70% in April and May combined and nearly 90% for the full second

quarter. This was the strongest quarter since the index began in 1993 (previous record of ~68% in Q1 2000), and April alone was the SOX's second-best month ever, behind the ~50% gain in February 2000. Within the index, several stocks nearly doubled or more within the period including Arm, Astera Labs, SK Hynix, and Micron. The broader market rally was remarkably concentrated, with the significant majority of the appreciation in the MSCI ACWI and S&P 500 during this period coming from semis and AI infrastructure. This chart from Nomura shows the 10 stocks that drove ~70% of the 28-session S&P 500 rally from March 30th to May 8th.



- **The second tier outperformed the top tier.** Even within semis, an interesting dynamic was how the #2 or #3-type players outperformed the companies most commonly viewed as the clearest leaders and highest quality operators. Intel outperformed Taiwan Semi. AMD outperformed NVIDIA. Marvell outperformed Broadcom. When a market goes all-in on a theme, the fastest money tends to hunt for whatever hasn't moved yet and / or the companies showing the strongest absolute growth, and the perceived laggards can become the highest-octane way to express the theme. We understand the trade and believe some of these challengers are, in fact, becoming much stronger competitively while expanding their serviceable markets. We'd simply note that the multi-year winners in this industry have historically been determined by process leadership, customer captivity, and execution, and those factors have a way of reasserting themselves once the catch-up trade runs its course.



- **The left-behind.** Despite the euphoria within AI infrastructure, multiple companies and sub-sectors sat in correction or bear market territory, spanning much of enterprise software, ecommerce, information services, healthcare, and consumer-facing businesses.
- **Frontier lab product releases continued at a fast pace.** The starkest example came on April 17th, when Anthropic released Claude Design and sent Figma's stock crashing. Adding intrigue, there were reports that Figma had been working closely with Anthropic and was caught completely off guard by the direct manner in which Anthropic chose to compete. Further AI-related product development and model improvement should be expected for some time, and with it, continued market uncertainty and volatility as investors handicap who might be next.
- **Incumbents responded with product development efforts of their own,** like Salesforce's April launch of Headless 360, which makes everything on its platform accessible as an API, MCP tool, or CLI command. Increasingly, companies across industries are realizing that both revenue growth acceleration and meaningful operating margin improvement will require token consumption. This is causing more companies across more industries to focus even more intensely on how to build, and become more efficient, through more effective token usage. A key question is which incumbents will be able to innovate and develop AI-focused product that customers actually want versus failing to reach even "good enough" status.
- **The frontier labs faced their steepest competition yet from China's open source, open weight models.** These models from companies like DeepSeek, Z.ai, Moonshot, MiniMax, Alibaba, Tencent, Meituan, and Xiaomi have improved at a fast clip with some reaching the frontier for various types of workflows. More enterprises are rallying for more open source, open weight solutions, not only for the potential cost savings but for the ability to better control their models and their data. To us, this always comes back to innovation. If the frontier labs are still early on the innovation curve, premium tokens should remain premium for some time. If innovation stalls, frontier lab token pricing should come careening lower along with it and the frontier labs would have a problem. Note there are reasonable ways for both frontier labs and open source to win.

Even if frontier models are best for only ~20% of work with the rest going to open source, that 20% could be more than enough for the frontier labs particularly if the pie continues to grow exponentially.

- **Non-fundamental decision-making was on full display.** Marvell was up 20%+ in overnight trading during the second quarter because Jensen Huang came on stage at Computex and yelled to the crowd, “[Marvell is the] next trillion-dollar company!” Then, in early July, Brookfield announced that it was expanding its partnership with Bloom Energy from \$5 billion to \$25 billion, a big jump, sending Bloom Energy shares up 15% after the market close. But Brookfield had already said over a month earlier that it would be expanding this partnership “by multiples,” so this wasn’t new information and should’ve already been priced-in. In fact, because we had assumed the partnership would go from \$5 billion to \$30 billion following the earlier announcement, our numbers actually went *down* following this announcement. Not surprisingly, the stock came all the way back down within 48 hours.
- **Google announced a historic \$80 billion equity financing.** A company with one of the strongest balance sheets on the planet concluded that it should raise equity to help finance its infrastructure buildout plans. To us, this says a lot about the scale Google believes is required to compete from here across TPUs, data centers, power, and frontier model development, and about how seriously it is taking a new class of large and well capitalized players like SpaceX and Anthropic that are also raising around \$80 billion each (and soon OpenAI, it seems). When the most well capitalized incumbents and the most well capitalized challengers are all raising the same enormous sums, we take it as another signal (among many others) that this buildout continues to hit new levels.
- **SpaceX went public at a >\$2 trillion market cap and, even under very bullish scenarios, 20x 2028 sales.** A large part of SpaceX reaching \$100 billion in 2028 revenue depends on a new AI infrastructure leasing business, anchored by customers like Anthropic and Google. SpaceX has shown an ability to go from planning to powered shells and live servers very quickly, can monetize that scarce compute externally, and retain the option to use that compute internally if its own models become competitive with Anthropic and OpenAI. But the revenue is hard to value when these contracts are short term and cancellable by either party. Putting valuation aside, SpaceX is a clear leader across several markets that should become very large. Still, going public at this valuation, with this much still “on the come,” makes it hard. It is also a useful case study on founder premiums, private versus public market attitudes, and inconsistency across AI infrastructure valuations. For example, many participating investors who celebrated SpaceX’s jump from its last private valuation mark of ~\$800 billion in December 2025 to its post-IPO valuation of >\$2 trillion are negative on publicly-traded neocloud and AI infrastructure companies that are also monetizing similar compute scarcity, or on semis and AI infra valuations more broadly. On the ways in which value can be ascribed to founders, the premium attached to Elon at Tesla and SpaceX versus the discount that we believe is attached to other highly successful founders like Mark Zuckerberg is fascinating.
- **A single day captured it all.** On June 2nd, a day that was mostly flat for the broader market, we had more than 15 positions (the significant majority) up or down 2% or more. This included one company up 10% on news that it might be rolling out agents within its apps, one company down 10% on news that its CFO is being replaced, and multiple companies up or down 3-6% on no news at all.

What Our Work Supports: AI Pilled, But Still a Balanced View in an Unbalanced Market


We've flagged throughout 2026 that our ideas have been falling into three categories: (1) companies fighting powerful narratives related to the new technological breakthrough (GenAI and agents) that we do not currently subscribe to, (2) companies that we believe are both far removed from technological disruption and could actually benefit from it, and (3) companies leading and shaping the technological breakthrough where consensus views might be positive but we believe remain too linear and not positive enough. We've then shared views on: enterprise software, where we concur with many that AI disruption risk is real, but unlike many, we see a few diamonds that we believe should perform well with solid execution despite being priced for sudden deaths; ecommerce, where we see much lower disruption risk for the marketplaces we admire; and AI infra, where we believe we are still early overall, have identified almost too many companies that we believe to be unusually competitively advantaged and benefiting from expanding product portfolios and customer captivity, and need to contend with the general market optimism for these companies and their inherent cyclicity during the inevitable air pockets and slowdowns.

Today's market structure has created the perfect recipe for stocks across all of these categories to swing wildly, often on a daily basis. And now the breakneck pace of innovation, where AI carries potential benefits for some companies and disruption risks for others, has handed a market that has never been more focused on daily price movements an endless supply of fresh narratives to push. These narratives essentially boil down to whether a company is an AI winner or an AI loser, with many companies like Meta and Google swinging all the way from one side to the other throughout the year. And within these categories, there are almost daily debates over whether the ebullience assigned to the AI winners or the fear assigned to the AI losers has gone too far or has not gone far enough.

While this isn't always the case, we've continued to identify a fairly balanced set of ideas. First, why are we still long-term positive on AI? AI hit another major inflection point early this year when OpenClaw was released, followed by Claude Cowork and other agent harnesses. This marked the arrival, much sooner than expected, of agent harnesses that connect to frontier models and allow tasks to be handled over longer periods of time. The METR chart below shows how AI agents have been quickly improving their abilities to handle more difficult tasks.

GPT-5 has a 50%-time-horizon of about 2 hr 15 min (95% CI: 1 - 4.5 hours)
Task length (at 50% success rate)

 METR



Vertical axis measures the length of a task, based on how long it takes a human expert, that an AI agent can complete autonomously with a 50% probability of success.

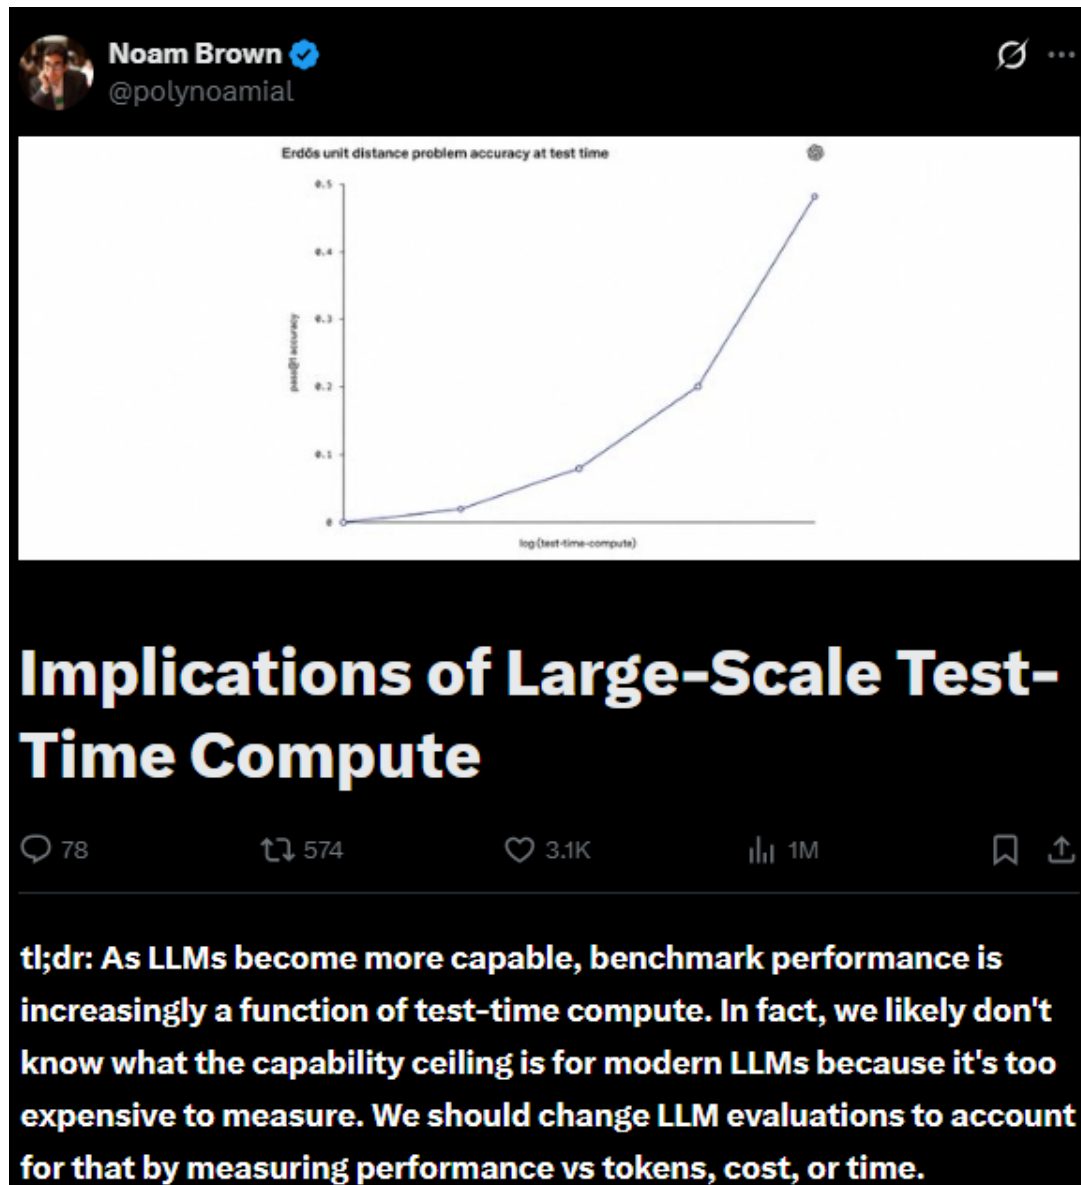
Task completion can unlock big productivity gains for most people on the planet across their personal and professional lives. Productivity gains should, in turn, translate to far more people paying for tokens both on their own and through their employers versus the low single-digit percentage of the U.S. population that we believe is using agents daily today. As agent adoption rates expand, and as agents continually improve from here, we believe AI-dedicated compute can go from under 50 gigawatts (GW) today to hundreds of GWs over the next several years. This has implications for the semis supply chain and others supporting AI infrastructure, ranging from chip designers, to chip manufacturers, to equipment providers to the chip manufacturers, to power producers, to those providing equipment and services to the power producers. While the temptation is to spread your capital far and wide across booming end markets, our goal is to stay focused on: (1) participating meaningfully in the best positioned companies that are leading and shaping these end markets and are still underappreciated in our view, (2) identifying certain companies that are earlier in their journeys but are showing clear signs of becoming durable leaders, and (3) avoiding the types of third tier and overly niche companies that look great while the boom continues but aren't strong enough in our view to make us want to continue to back them when we hit a more boring phase of this compute and infrastructure buildout.

Despite how positive we are on AI's long-term trajectory, we continue to feel that the world is a big place, with many companies outside of the purest "AI winners" delivering unusual value to their core customers that is incredibly difficult for competitors, including the AI labs like Anthropic and OpenAI, to replicate. As a bonus, valuations can be unusually cheap for these types of businesses currently. OpenAI's experience of roaming into social media applications like Sora, only to realize that sustaining engagement was difficult while the required time and resources allowed its more focused competitor, Anthropic, to take a big lead in the very important and already lucrative enterprise workflow market, is instructive. So is the more recent phenomenon of U.S. labs running into strong Chinese competition at the frontier, where open source, open weight models with high-level capability are being made available to enterprises at low cost. We believe two things can be true at the same time: Anthropic and OpenAI can sustain incredible businesses within their areas of competence today, and there can be many other winners, including open source model providers and certain incumbents across markets like ecommerce, gaming, and enterprise software that the frontier labs won't be able to win themselves and that should prove very difficult for AI native new entrants to wrestle away from established companies that are doing all the right things to accelerate their respective businesses through more effective AI integration efforts.

In this market though, where almost every debate frames AI as zero-sum, this type of balanced view has mostly felt like running against 50 mph winds, at least thus far in 2026. On a daily basis, we can almost guarantee that 1/4 to 1/3 of our portfolio will be down meaningfully and another 1/4 to 1/3 will be up meaningfully. Similar to the old marketing truth that "you know you're going to waste about half of your advertising budget, you just don't know which half," you don't know which 1/4 to 1/3 of the portfolio will outperform and which will underperform because it is typically based on no real news at all.

We continue to think about the mental warfare a more hardware-oriented world can wage on fundamental investors and how many are truly wired to embrace it if that is what the work supports. It can be a great behavioral challenge to allocate high percentages of your funds to businesses that you know will be highly cyclical, but it's a very important time to remain open-minded in our view. For the first time in history, we have companies that have figured out how to *manufacture intelligence*. This is worth repeating: we now have companies that manufacture intelligence. For \$20 per month, individuals can access AI systems that perform above the level of most people on a wide range of projects, and in some tasks can match or exceed highly experienced professionals. For those with big budgets or

working within frontier labs, the level of intelligence available today is higher still. And all indications are that we're still early on the manufactured intelligence capability curve. Model improvements are still being released at a fast pace, and as shared by OpenAI's Noam Brown below, we likely underestimate the productivity value of today's models because allowing them to think for much longer would lead to better outcomes but is currently too inconvenient and costly, particularly as new, better models are surfacing quickly.



What this means is we will need way more compute over the next five and 10 years and beyond. Take a recent interview with the founders of Etched, which has developed a rack-scale solution for inference (the output of the AI models) featuring more tokens per second (throughput), lower latency (speed), lower cost per token (affordability), and better tokens per watt (efficiency). Their argument is that inference demand is compounding faster than the industry's ability to serve it. Today's deployed compute could satisfy only a fraction of the demand that would materialize if agents were allowed to run as long and as often as their users would like, and every meaningful reduction in cost per token to date has been met with a more than offsetting increase in consumption.

Invest Like the Best Podcast Episode with Etched Founders Gavin Uberti and Rob Wachen

Name	What They Said
Gavin Uberti	Power is the gating constraint. What matters is “power availability and time to power,” and the more power needed, the more constrained supply becomes.
Rob Wachen	Future models may not just be bigger in parameter count. They may run many parallel copies, activate many more experts, and use huge experts spread across multiple server racks.
Rob Wachen	Future AI systems may be able to hold essentially all written knowledge in short-term working context, implying much larger context windows, active memory, and inference systems.
Gavin Uberti	Inference could become a huge part of the economy, potentially even a majority of global GDP over time.
Gavin Uberti	Countries may eventually dedicate a large share of total energy to data centers doing inference. Data center efficiency would then determine how many AI agents a country or company can run.
Rob Wachen	Trillion-dollar data centers are likely a matter of time. The analogy is semiconductor fabs: scale keeps increasing because unit economics improve as facilities get larger.
Gavin Uberti	Intelligence is becoming cheap to produce relative to its value, creating a potential multi-decade shortage of tokens.
Gavin Uberti	The largest companies may be the ones that produce most of the world’s token supply and own much of the token supply chain.

Maybe this level of optimism for the growth ahead will prove to be misplaced euphoria, and there are clear data points to monitor along the way, but the undeniable trend thus far has been better models in rapid succession, agent harnesses now allowing for the rise of automated tasks and work completion, all translating to an explosion of manufactured intelligence, and consumers and enterprises alike voting with their wallets that they are seeing enough productivity gains, or at least productivity potential, to want to consume more and more tokens. In our view, manufactured intelligence is already allowing for enough productivity across a large enough and growing set of applications and tasks to conclude that demand for manufactured intelligence should be some level of enormous.

What this means is we are putting a lot of compute in the ground, and we will ultimately need a lot more infrastructure. While there is a wide range of industry forecasts in terms of just how fast this will ramp, it seems clear to us that the ramp will be steep. The world should add something like 20 GWs of AI-related compute this year and over 30 GWs next year, even with the ongoing difficulties in getting sites permitted, powered, and operational. With even modest assumptions about agent adoption, the installed base can grow from tens of GWs today, to a few hundred GWs by the early 2030s, to potentially 1,000+ GWs by 2040 (which would make the chart below look conservative).

Plaintext				
REGION	2026	2030	2035	2040
USA	20	65	115	168
China	9	24	45	72
Europe	7	18	32	47
APAC	6	19	38	64
RoW	3	11	24	48
GLOBAL	45	137	254	399

(All values in Gigawatts. Sources: BloombergNEF, S&P Global, IEA, McKinsey, Gartner)

What might this look like from an average daily active user level? Very roughly, daily active users today might be consuming 10,000-35,000 tokens per day, or 300,000 to one million tokens per month. Heavy knowledge workers, coders, analysts, students, investors, etc. will be meaningfully above that in many cases. That consumption is mostly chat-style questions and search-adjacent queries, with some drafting, summarization, and analysis work layered on top, and with only the earliest adopters running agents in the background.

Active Daily User Today	
Daily AI Use	Token Equivalent / Day
Quick questions / explanations	1k-3k
Writing, rewriting, tone help	2k-6k
Email / text / message drafting	1k-3k
Research / web-style synthesis	3k-10k
Reading / summarizing articles, docs, PDFs	3k-10k
Work analysis / spreadsheet / decision help	2k-8k
Coding / troubleshooting / formulas	2k-10k
Follow-up conversation / iteration	2k-8k
Hidden system/tool/context overhead	1k-5k

Several years from now, as more of the population becomes more familiar with more powerful tools, including agents that can run on a loop and handle far more tasks across areas like shopping, travel, finance, health, and education, monthly token usage might be closer to 50 million to 200 million tokens per month, or 50-200x today’s levels. The mix should also invert over time, with always-on agentic workflows (research running overnight, purchases negotiated, finances and health tracked, work products drafted and revised while you sleep) consuming multiples of what direct chat will.

A plausible 2040 heavy user might look like this:

Monthly AI Use	Token Equivalent
Direct chat, writing, research	5M-20M
Email / calendar / admin agent	5M-20M
Personal research / news / investing agent	10M-50M
Work agent monitoring docs, meetings, Slack, CRM	10M-50M
Coding / analysis / spreadsheet agent	20M-150M
Shopping / travel / finance / health / education agents	5M-50M
Hidden reasoning / retries / verification	2-10x multiplier

What this means is we could need hundreds of GWs of compute to serve billions of people who want to use AI in these ways as the tools improve and as more enterprises and people alike better understand how these tools can be used. Some including SemiAnalysis have suggested that we could be over one terawatt of compute by 2040.

What 500 GW Implies

If **250 GW** of a 500 GW data-center world were actually AI workloads, and real-world efficiency were **10-20x worse** than NVIDIA's clean benchmark, then 3B active users could average roughly:

AI Capacity	Real-World Efficiency Haircut	Tokens / Month / Active User
250 GW	10x worse than benchmark	~60M
250 GW	20x worse than benchmark	~30M
500 GW	10x worse than benchmark	~120M
500 GW	20x worse than benchmark	~60M

But the road from here to there is unknowable, and the behavioral dynamics involved mean it will be difficult (but possible) to avoid running into periods of overbuild along the way. This means investors like us who are positive on the long-term trajectory but are aware of the history and cyclical nature of these businesses can feel almost certain that the share prices of the companies leading the charge will be very volatile. There are likely to be quarters like 1Q 2026 when many of these companies show exponential growth that makes us realize our prior expectations were too conservative, only to feel in other periods that things have plateaued sooner than expected or run into a surprise obstacle, leading us to question whether this infrastructure buildout will prove less impressive than previously believed.

This will likely mean several days and weeks when these stocks are up meaningfully and several times when they fall 10%+ in a single day. Understandably, many seasoned investors, at a human level, aren't wired to deal with this type of volatility or have mandates that make it very difficult to do so and will largely sit this out. Even more investors will have some tolerance for the volatility but will have fairly low ceilings for where portfolio weights exposed

to these dynamics can go. We are willing to endure volatility and have complete flexibility here, but there must be strong offsets to the factors causing it. For example, we believe cyclical companies deserve a penalty in a portfolio construction model that determines inclusion and sizing, because all else equal, cyclicity makes it harder for operators and investors alike to understand the true earnings power of the business. We believe sizing decisions should also have a timeliness consideration, as we believe sizing should involve both long-term conviction and considerations for the path along the way. To get to large position sizes, this means we would need to believe cyclical businesses like semis and industrials are that much stronger than the non or less cyclical alternatives in fundamental areas like competitive advantage, growth opportunity, governance, and valuation. We have gotten there on some of these, and there is a newer crop of companies (at least in terms of finding newfound strength) that we believe are beginning to demonstrate similar-type characteristics. Similar to the software boom that started circa 2014 and lasted ~10 years, when it felt hard not to allocate substantial chunks of invested capital to software companies like Salesforce, ServiceNow, Adobe, Workday, Microsoft, and Autodesk that were developing better business models and better products that were boosting productivity at the same time, our portfolio construction model has continued to tilt more towards the leaders we see that are shaping this productivity boom. At this stage, we'd be surprised if our work suggested that should drastically reverse, but we will continue to monitor our pre-mortems, such as model breakthroughs that would be massively more compute efficient, regulatory issues that would block progress, or a realization that model scaling laws and meaningful improvements are mostly in the rearview mirror. In the absence of these types of events, just how much we participate (and there can be a wide range when looking at portfolio weight) will depend on factors such as variant perception around valuation and our ongoing work that determines which of these companies are further expanding their product portfolios, further separating themselves from the pack, and becoming even more difficult to compete with.

All of that said, we continue to find opportunities elsewhere, although they have basically been a bag of detractors through June. We do not currently see a world in which the models get to a level where humans become unimportant, or where the frontier labs, namely Anthropic and OpenAI, eat every industry and eliminate the ability of others to establish leadership positions within them. This means that Category #1 opportunities that we referenced earlier remain a clear ideas source and are still producing new holdings. As we've noted in prior memos, in a few cases, we are seeing prices for businesses we admire that we flat out did not expect to see due to AI-related fears that we believe are misguided despite how positive we are on AI more broadly. With the caveats again that trying to call a bottom in these types of situations is a fool's errand, and that there are always pre-mortems to monitor, we believe there is already unusual opportunity in some of these companies.

These are the current battleground stocks. We recounted last quarter just how violent the moves here have been, including a two-week stretch from late January through the first week of February when several digitally oriented, asset-light businesses like Salesforce, ServiceNow, Workday, CoStar, Moody's, S&P, FICO, Thomson Reuters, Wolters Kluwer, and AppLovin fell 15-30%, and 115 stocks in the S&P 500 declined 7% or more in a single day over a mere eight trading sessions, all while the S&P 500 itself fell only 1.5%. That dynamic has kept rotating beneath a surface that has looked calm. Last time, we noted a couple of potentially interesting areas within Category #1 and share some updated thoughts here:

Ecommerce

In our last letter, we walked through the market's fears here in detail: that third party apps like ChatGPT, Claude, and Gemini could soon provide incredible shopping agents that allow consumers to shop and transact within those apps instead of directly with the platforms, putting lucrative components of these business models like high

margin ads in jeopardy, plus the company-specific concerns swirling at the same time (e.g., autonomous vehicle competition for DoorDash, competitive intensity for MercadoLibre). We also shared why this could be a bigger source of opportunity compared to SaaS. The addressable consumer retail markets are measured in the trillions of dollars versus hundreds of billions for traditional enterprise software, ecommerce penetration remains well below 40% of total retail spend across most key geographies, and certain ecomm platforms could be well positioned to “box out” third party shopping agents that refuse to partner in ways that create win-wins. For example, if a marketplace has >55% share in its core U.S. market and deep operational relationships with hundreds of thousands of required partners across hard-to-do areas like merchandising, payments, refunds, and customer service, a third party agent would likely need this marketplace’s cooperation to deliver a winning and consistent consumer experience in the category, and certain leading marketplaces would likely cooperate only on terms that protect their economics. Marketplaces remain the area within Category #1 where we believe good ideas can be found, both in terms of defensibility and the ability to keep growing. The valuations we did not expect to see earlier in 2026 haven’t really bounced back yet in certain cases, and the businesses we admire in this area have continued to perform well in our view while we see the stocks continuing to trade primarily on the market’s AI narratives and its short-termism that has a distaste for higher levels of reinvestment for future growth.

Other Consumer Software

Last quarter, we wrote about how the Moat War had extended into consumer software. We saw potential ideas in areas like music streaming where certain companies like Spotify might be back to ~2024 prices despite having bigger businesses with higher levels of free cash flow today. You might find yourself staring at a company that has hundreds of millions of monthly active users, hundreds of millions of premium subscribers, the scale to negotiate more favorable economics with key partners that might increasingly rely on them for distribution, track records of moving first on new features and formats, and demonstrated pricing power. You will likely need a view on the AI disruption threat. For example, you might conclude that AI-generated music fears are misguided for the same reasons AI-generated video hasn’t undone Meta: this is the type of content that might still largely require a human connection, any proliferation of content might still need distribution to find an audience, and certain companies might have that type of leading distribution platform.

These types of ideas might extend into areas like long form video. The narrative headwinds here are also plentiful. The view for companies you might be interested in could be that revenue growth is still healthy but more mature, which the market is having a difficult time paying attention to in a world where an entire corner of the tech sector is growing gangbusters. The view might be that a company of interest is now reaching for inorganic growth and isn’t yet ready to deliver an advertising business that is large enough or fast-growing enough to move the needle. There might be management turnover with a visionary founder stepping aside. And perhaps a lack of big hit content in recent periods. You will hear that people’s attention spans are only getting shorter, so there will be fears that long form video content is no longer in demand. And AI-related content proliferation fears might hover over the industry: if anyone can soon produce watchable long-form content with AI tools, what happens to the barriers to success?

To the extent you reach different conclusions, there could be opportunity. For example, that producing winning long-form original content across genres, languages, and cultures year-in and year-out might require things that are extraordinarily difficult to replicate; that it might require a global audience large enough that the world’s best content creators feel their works of art can be distributed in ways that allow them to be enjoyed by as many people as possible while also allowing the creators to be appropriately compensated; and that it might require all sorts of domain expertise in production and otherwise to understand how to deliver a consistent flow of winning long form

content to such a large audience, where emotional connection with an audience is everything and where AI tools might assist the humans involved far more than replace them. On advertising, you might believe that there is nearly zero precedent for a highly engaged platform to forever fail at eventually creating a large advertising business, and that the numbers to date suggest the company is squarely on that trend. You might agree that people's attention spans are shorter now, but you might see that time spent on certain platforms continues to rise nonetheless. When the story is compelling enough, and therefore the entertainment value is high enough, viewers might be definitionally feeling connected to the storyline at minute seven, minute 24 and minute 45, with potentially no data supporting that the minutes passing themselves are outcome determinative. All of this might sit within a business that continues to expand free cash flow margins and at a valuation that is roughly half of where it has been in recent years.

SaaS

A reminder of the market's five fears in enterprise software from our last memo: code generation democratizing product development, AI-first entrants capturing the agentic growth layers, seat-based pricing doom, evaporating switching costs, and management teams that do not understand the importance of becoming leaner and returning capital. We continue to believe a smaller group of incumbents should remain well positioned in an increasingly agentic world and continue to apply the criteria we articulated: deeply embedded solutions, business logic domain expertise, room to grow within legacy offerings, a track record of winning product development, a right to win agentic workflows within their areas of expertise, demonstrated progress bringing agents to market that are straightforward to implement and customize, and maniacal management teams.

We believe this is the toughest area within Category #1, and you don't get extra points as investors for degree of difficulty. Barriers to entry have come way down. It is unclear what human capital requirements will look like at end customers. It is unclear where gross margins might need to go to serve customers agentic solutions in ways that encourage widespread adoption. And with the rise of AI native customers and more competition, it is unclear how much growth is left for the industry even where disruption risk might be low. What this means is we are confident valuation will be on our side in SaaS, probably for the foreseeable future, but valuation cannot be the only thing or even the most important thing. We need to see a business that should prove resilient, and in this case, a path to revenue growth acceleration that we have conviction in.

A potential idea here could be one of the most controversial of all currently, Salesforce. Agentforce is still very much a work in progress, but its ARR reached \$1.2 billion as of Q1 FY27, up 205% year-over-year, with 3.8 billion Agentic Work Units processed. Bookings from premium, AI-embedded SKUs were up ~60% year-over-year. The April launch of Headless 360 made the platform accessible as an API, MCP tool, or CLI command. Early adoption disclosures (4.5 million MCP calls into the platform since launch, ~1 trillion API calls across Core products in Q1) suggest the market's "the interface is the moat, and the interface is dying" framing misses where Salesforce's value resides. Business logic expertise, the permissions, compliance guardrails, and embedded workflows could prove critical. Management guided FY27 revenue to \$45.9-46.2 billion (10-11% growth) with acceleration expected in the second half and raised its FY30 revenue target to \$63 billion in February, supported by RPO of ~\$68 billion growing 11% and current RPO of ~\$33 billion growing 14%. At \$2 per conversation for customer-facing agents (or ~\$0.10 per action under Flex Credits), a customer whose seat count shrinks can still become a larger customer as agents take on more of the work, which is why a focus on seat attrition alone could misread how the economic relationship with existing customers might evolve. Slack could also be more strategically important for the agentic world than many appreciate. With >45 million DAUs, >75% of the Fortune 100, and employees spending hours per day signed into the platform, Slack is where work happens for many companies. Slackbot could be a routing layer for an enterprise's agent network,

helping to connect employees to apps and an expanding catalog of agents. Each enterprise Slack seat should be a potential Agentforce consumption entry point. Perhaps the Slack MCP surpassing one million active users within 6 weeks of launch was an early signal that Salesforce could be moving itself more “in the token path.”

Portfolio Turnover in the Modern Era

It has been tough sledding in public markets for many investors who focus on “quality” companies (based on how they define that term). Public market structure has evolved, and the types of companies one views as “quality” could be evolving as well. Quality-focused investors who haven’t been open to how they might need to evolve their respective processes while staying true to their investment belief systems could be disadvantaged.

One of the pitfalls to be mindful of that we continue to flag is maintaining a strategy of buying and keeping the same type of position size over, let’s say, the next year or two. In the past, when market dynamics were different, investors could reasonably only consider long-term conviction when making sizing decisions. Today, we believe a very different market structure and much higher levels of volatility should mean that the path also factors into sizing models. The rise of passive, quant, and multi-strategy “pods” has given non-thinking and short-term focused decision-making a disproportionate share of the pie, and the market uncertainty that AI brings to more businesses could be just the fuel this market needed to create a potentially lasting volatility bomb. AI spawns new narratives around AI winners and losers almost weekly, and some of these debates (AI-related or otherwise) become protracted in ways that can cause multiples to change dramatically in short order. Particularly to the downside, we do not believe it is prudent to sit with large position sizes if and when we have identified a company issue or debate that we believe will ultimately be resolved favorably but that we also believe could be a losing battle for some time. This doesn’t mean we exit or even reduce position sizes to very small, but it could involve more trading around the core position. If your portfolio sizing model contains factors for timeliness and company-specific controversies as we believe it should, then these sizing decisions should happen naturally within the process.

Outlook

Investing is humbling. The world continually evolves. Innovation continually disrupts the status quo. Global economic events that can be painful but ultimately finite can move the share prices of companies that are bigger than the macro more than you’d expect in the shorter-term. Company share prices can, and often do, move in both directions for reasons that have little to do with underlying business performance. Our mandate isn’t to know what will happen next. It is to consistently back companies that we believe, based on everything we understand today, have high probabilities of delivering healthy returns and low probabilities of being disrupted over multi-year periods.

So, in this market environment, we continue to execute the process. As it relates to AI, we continue to be open to the viewpoint that some of the best businesses for the next 10 years are more cyclical and capital-intensive than we’ve seen historically, while remaining alert to the risks involved. This environment could also mean that we will stay involved in certain areas that are out of favor and currently classified as “AI losers” if our research continues to conclude that current narratives leading to the share price weakness are misguided.

Despite the added volatility and almost constant macro uncertainties, what we do stays consistent at its core. We are dogged hunters for next Fortune 100s, we narrow to the even smaller group that we believe can grow free cash flow per share in ways that can allow us to exceed our cost of capital, and we aim to only back the even smaller group that falls within one of our core mispricing setups. Very importantly, we do it all within a structure that allows us to express true conviction and challenge the direction that most public market participants have chosen to go. With AI

progressing at rates that we still view as incredible, and with the share prices of so many companies outside of AI infrastructure moving out of step with underlying business performance this year, we see opportunities ahead.

Sincerely,

/s/ Brandon Ladoff

Brandon P. Ladoff
Founder, CEO & Portfolio Manager
DENMARK Capital

Important Disclosures & Legal Information

For Informational Purposes Only. This document and the information contained herein are provided for informational and educational purposes only. This content does not constitute an offer to sell, or a solicitation of an offer to buy, any security, financial product, or instrument. Any such offer or solicitation would only be made through separate, formal confidential offering documents.

No Investment Advice. The views expressed in this memo represent the opinions of Brandon Ladoff and DENMARK Capital as of the date of publication. These views are subject to change at any time without notice. Nothing in this document should be construed as investment, legal, or tax advice. You should consult with your own professional advisors before making any investment decisions.

No Solicitation. This memo is not, and should not be construed as, an offer or solicitation to invest in any fund or account managed by the author or any affiliate. No fund managed by the author is being offered or marketed by means of this memo.

Disclosure of Interests. The author and/or entities affiliated with DENMARK Capital may hold positions in the securities mentioned and discussed or in other securities discussed from time to time and may purchase or sell such securities at any time without notice. The mention of specific securities is solely to illustrate an investment philosophy and research process and should not be viewed as a recommendation to buy or sell. Securities mentioned are discussed solely to illustrate research process and are not necessarily held by any fund or account managed by DENMARK Capital. The securities discussed do not represent all of the securities held, purchased, or sold, and it should not be assumed that any security discussed was or will be profitable.

Third-Party Information. Certain data and charts in this memo are sourced from third parties. While obtained from sources believed to be reliable, this information has not been independently verified and no representation is made as to its accuracy or completeness.

No Performance Information. This memo does not present the performance of any fund or account managed by DENMARK Capital.

Forward-Looking Statements. This document contains forward-looking statements based on current expectations and projections about future events. These statements are not guarantees of future performance and involve risks and uncertainties that are difficult to predict. Actual results may differ materially from those expressed or implied.

Risks of Investing. All investing involves risk, including the possible loss of principal. Past performance is not indicative of future results. Public market volatility can lead to significant price fluctuations that may not align with the fundamental analysis presented here.