

TOP MARK CAPITAL MANAGEMENT LLC

Report For the Quarter Ended

June 30, 2024

416 13th Street

2nd Floor


San Diego, CA 92101

E: contact@topmarkcapital.com


T: +1 (646) 470-3016

INDEX

MANAGEMENT’S LETTER TO PARTNERS.....	4
TOP MARK CAPITAL PARTNERS LP.....	6
PARTNERSHIP LETTER & PERFORMANCE UPDATE.....	7
PORTFOLIO UPDATE & COMMENTARY.....	13
TOP MARK HEALTH PARTNERS LP.....	16
PARTNERSHIP LETTER & PERFORMANCE UPDATE.....	17
PORTFOLIO UPDATE & COMMENTARY.....	20
APPENDIX.....	27
HYPERSCALE CAPEX AND NVIDIA’S 5 RISKS.....	28
A BRIEF FOLLOW-UP ON YIELD CURVE INVERSIONS.....	44
IMPORTANT DISCLAIMER.....	46
REFERENCES.....	47



ASSETS UNDER MANAGEMENT



Dear Partners and Prospective Partners,

We hope this letter finds you and your families well and having thoroughly enjoyed the summer. At Top Mark Capital, we've been particularly immersed in researching healthcare and biotech businesses. You'll find excerpts from our findings in the Top Mark Health Partners (TMHP) section of this letter. In the Top Mark Capital Partners (TMCP) letter, we provide our portfolio performance and updates, and details on our private credit allocation - including the latest loan repayment.

In the appendix, you will find two memos, both "Thematic Updates" subsection, sharing our latest thoughts on key themes. In the first memo we discuss "Software is Eating the World", a theme we've been following for quite some time. The memo focuses on hyperscaler capex as it relates to the top 5 risks to Nvidia's Dominance in the data center. The second memo revisits and updates our yield curve inversion analysis from prior letters.

We're excited about our portfolios and are eternally grateful for your continued support.

Performance Highlights

Both TMCP and TMHP delivered strong results in the second quarter of 2024. TMCP delivered a net return of 11.7% and TMHP posted a net return of 11.1%¹²³. As always, we encourage you to focus long-term, cumulative results rather than quarterly fluctuations. You'll find these cumulative results in each partnership's respective section.

Investment Philosophy and Partner Outlook

Concentrated equity investing has proven successful for long-term outperformance. Top Mark Capital's approach is unique, reflecting your manager's operational and investment

¹ Net Performance figures are presented at the fund level, after management fees and performance fees

² Gross Performance figures are presented at the fund level, before management fees and performance fees

³ Please review the Important Disclaimers on Page 46

experiences, enhanced by our ongoing efforts to expand our competencies. We welcome deeper conversations with prospective partners about either partnership and appreciate referrals to like-minded individuals.

As our Partner, you'll receive these quarterly letters offering insights into our strategies and performance. Formidium Fund Management Services, our administrator, will continue to provide quarterly account statements and annual audited financials. Please inform us if you miss any communication.

Given our long-term focus, we believe less frequent updates may better serve you by minimizing short-term market distractions. For those interested in more frequent updates, we invite you to listen to the Telltales podcast. Subscribe at telldetails.us for episode notifications.

Our edge lies in our capital allocation skills and your patience as investors. This forward-looking perspective distinguishes us from short-term-focused investors and defines our Partnership's essence.

We're fortunate to be surrounded by like-minded, long-term investors. Outperformance comes from looking further ahead than the short-term crowd. Remember, we're an investment partnership, not a hedge fund, and we seek a different relationship with our Partners.

If you share our temperament and perspective, we invite you to consider subscribing to one of our partnerships. The next contribution window for both TMCP and TMHP is October 1, 2024. Please contact us to begin the process.

As always, we deeply appreciate your confidence and support.

Sincerely,

Michael Nicoletti

Jason Wallace

TOP MARK CAPITAL PARTNERS LP

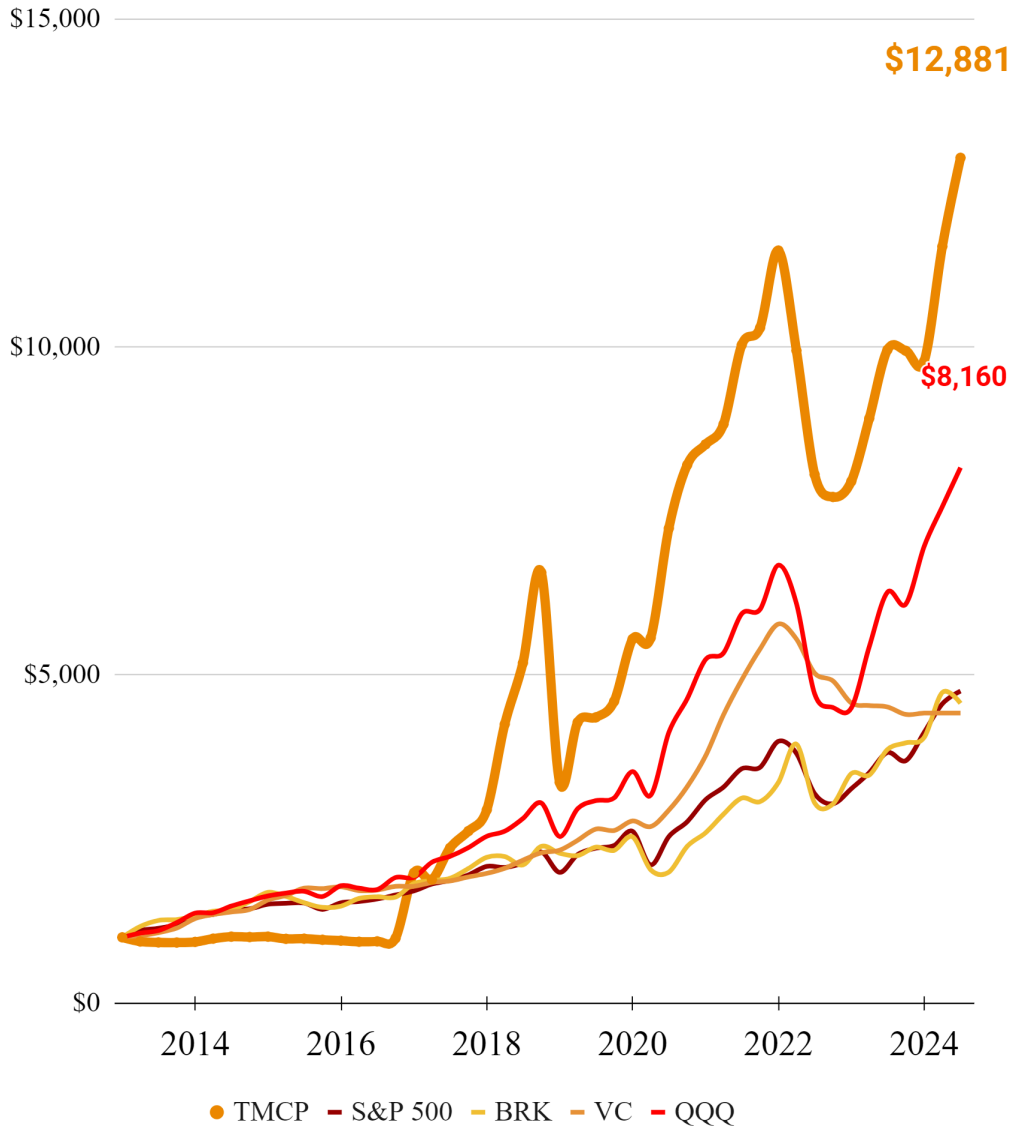
Partnership Letter

For the Quarter ended June 30, 2024



PARTNERSHIP LETTER & PERFORMANCE UPDATE

CUMULATIVE RETURNS ON \$1,000⁴



⁴ TMCP figures are presented at the fund level, net of fees (i.e., management fees, costs, and performance fees). SPXTR: S&P 500 Total Return. BRK: Berkshire Hathaway Stock. QQQ: Nasdaq 100 Index ETF. VC Index: Cambridge Associates US Venture Capital Index. Please review the Important Disclaimers on Page 46.

TMCP CUMULATIVE RETURNS TO DATE⁵

To: 6-30-2024, From:	TMCP % Gross Return	TMCP % Net Return	S&P 500 Total Return	Berkshire Hathaway	VC Index	QQQ
One Year	31%	29%	25%	18%	25%	30%
Three Years	34%	28%	33%	46%	33%	38%
Five Years	252%	196%	102%	92%	102%	165%
Ten Years	1820%	1175%	235%	222%	235%	453%
Since Inception	1866%	1191%	374%	373%	374%	718%
<i>Annualized Since Inception (October 1, 2012)</i>	<i>28.8%</i>	<i>24.3%</i>	<i>14%</i>	<i>14%</i>	<i>14%</i>	<i>20%</i>

Performance Update

Over the last five years Top Mark Capital Partners (TMCP) delivered a robust 196% return⁶, outpacing the S&P 500's 102%. Since our inception in September 2012, our capital has grown at an impressive annualized rate of 24%, significantly surpassing the S&P 500's 14%. To date, TMCP's total return stands at 1191%, dwarfing the S&P 500's 374%. The table above illustrates this point. For those interested in annual or quarterly returns data, please refer to the tables on the following pages. However, focusing solely on short-term periods can lead to misconceptions about the success of any investment strategy.

As I write, the market is reacting strongly to various news items, ranging from the reversal of a popular carry trade to Berkshire Hathaway's sale of approximately half of its enormous

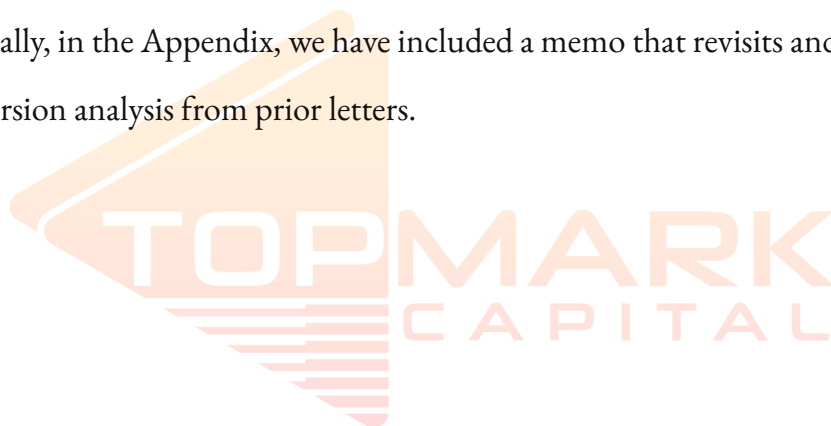
⁵ TMCP % Gross Return is at the fund level before management and performance fees; TMCP % Net Return is at the fund level, net of all fees. SPXTR: S&P 500 Total Return. BRK: Berkshire Hathaway Stock. QQQ: Nasdaq 100 Index ETF. VC Index: Cambridge Associates US Venture Capital Index. We include the S&P 500 Total Return Index, a Nasdaq 100 index fund, Berkshire Hathaway stock, and the Cambridge Associates US Venture Capital index as gauges against other potential investments. These are not explicit benchmarks but provide a cumulative perspective. Please review the Important Disclaimers on Page 46.

⁶ Performance is at the fund level, net of all fees (i.e., management fees, costs, and performance fees) and rounded to the nearest whole percentage. Please review the Important Disclaimers on Page 46.

Apple position. These gyrations and aberrations create opportunities for long-term investors like us.

During the quarter we made minor additions to existing portfolio positions and had one of our private credit facilities repaid in full. This leaves us with a very strong cash position in the event of future ‘opportunities’. Details on these changes as well as our cash position are highlighted for partners in the *Portfolio Update and Commentary* section.

Regular readers of our letters are familiar with our thematic approach to research, which broadly aligns with our desire to foster and grow our circle of competence. In the Appendix to this document you will find two memos for your perusal. The first explores hyperscaler capital expenditures, analyzing how they have changed and why we can expect them to continue to grow. Additionally, in the Appendix, we have included a memo that revisits and updates our yield curve inversion analysis from prior letters.



TMCP ANNUAL RETURNS DATA⁷

Period	TMCP % Gross Return	TMCP % Net Return	S&P 500 Total Return	Berkshire Hathaway	VC Index	QQQ
2024 1H	32%	32%	15%	13%		17%
2023	24%	23%	26%	16%	-3%	55%
2022	-30%	-31%	-18%	4%	-21%	-33%
2021	41%	35%	29%	30%	53%	27%
2020	64%	54%	18%	2%	36%	48%
2019	79%	65%	31%	11%	19%	39%
2018	15%	14%	-4%	3%	18%	0%
2017	58%	48%	22%	22%	11%	33%
2016	133%	109%	12%	23%	0%	7%
2015	-5%	-6%	1%	-12%	13%	10%
2014	9%	9%	14%	27%	22%	19%
2013	-6%	-7%	32%	33%	29%	37%
2012Q4	0%	0%	0%	4%	7%	0%

⁷ TMCP % Gross Return is at the fund level before management and performance fees; TMCP % Net Return is at the fund level, net of all fees. SPXTR: S&P 500 Total Return. BRK: Berkshire Hathaway Stock. QQQ: Nasdaq 100 Index ETF. VC Index: Cambridge Associates US Venture Capital Index. Please review the Important Disclaimers on Page 46.

TMCP QUARTERLY RETURNS DATA ⁸

Quarter Ended	TMCP % Gross Return	TMCP % Net Return	S&P 500 Total Return	Berkshire Hathaway	VC Index	QQQ
June 30, 2024	12%	12%	4%	-3%	n/a	8%
March 31, 2024	18%	18%	11%	17%	n/a	9%
December 31, 2023	-1%	-2%	12%	2%	0%	14%
September 30, 2023	0%	0%	-3%	3%	-3%	-3%
June 30, 2023	12%	12%	9%	11%	0%	15%
March 31, 2023	12%	12%	7%	-1%	-1%	21%
December 31, 2022	3%	3%	8%	15%	-7%	0%
September 30, 2022	-4%	-4%	-5%	-1%	-2%	-4%
June 30, 2022	-19%	-19%	-16%	-23%	-9%	-23%
March 31, 2022	-13%	-13%	-5%	17%	-4%	-9%
December 31, 2021	13%	11%	11%	10%	7%	11%
September 30, 2021	3%	3%	1%	-2%	10%	1%
June 30, 2021	16%	14%	9%	9%	12%	11%
March 31, 2021	5%	4%	6%	11%	17%	2%
December 31, 2020	5%	4%	12%	9%	14%	13%
September 30, 2020	15%	13%	9%	20%	12%	12%
June 30, 2020	35%	30%	21%	-2%	9%	30%
March 31, 2020	0%	0%	-20%	-20%	-3%	-10%
December 31, 2019	24%	21%	9%	9%	6%	13%
September 30, 2019	7%	6%	2%	-2%	-1%	1%
June 30, 2019	2%	2%	4%	6%	7%	4%
March 31, 2019	33%	27%	14%	-2%	6%	17%

⁸ TMCP % Gross Return is at the fund level before management and performance fees; TMCP % Net Return is at the fund level, net of all fees. SPXTR: S&P 500 Total Return. BRK: Berkshire Hathaway Stock. QQQ: Nasdaq 100 Index ETF. VC Index: Cambridge Associates US Venture Capital Index. Please review the Important Disclaimers on Page 46.

Quarter Ended	TMCP % Gross Return	TMCP % Net Return	S&P 500 Total Return	Berkshire Hathaway	VC Index	QQQ
December 31, 2018	-39%	-35%	-7%	8%	7%	-10%
September 30, 2018	61%	54%	11%	7%	12%	16%
June 30, 2018	89%	76%	3%	-5%	10%	11%
March 31, 2018	73%	62%	6%	9%	7%	11%
December 31, 2017	28%	25%	11%	17%	6%	14%
September 30, 2017	46%	39%	8%	10%	5%	10%
June 30, 2017	23%	19%	9%	4%	5%	17%
March 31, 2017	111%	91%	10%	16%	3%	12%
December 31, 2016	135%	111%	8%	13%	3%	11%
September 30, 2016	7%	6%	6%	1%	4%	9%
June 30, 2016	-1%	-1%	4%	10%	-3%	-3%
March 31, 2016	-3%	-3%	8%	9%	-2%	8%
December 31, 2015	-3%	-3%	0%	-3%	1%	5%
September 30, 2015	-1%	-1%	-6%	-10%	6%	-3%
June 30, 2015	-3%	-3%	1%	-9%	11%	4%
March 31, 2015	-2%	-3%	6%	5%	15%	8%
December 31, 2014	0%	0%	6%	19%	13%	11%
September 30, 2014	3%	2%	6%	10%	6%	13%
June 30, 2014	9%	9%	7%	7%	8%	8%
March 31, 2014	7%	6%	13%	10%	18%	12%
December 31, 2013	1%	1%	16%	6%	20%	24%
September 30, 2013	-1%	-2%	8%	9%	11%	15%
June 30, 2013	-7%	-8%	14%	26%	7%	10%
March 31, 2013	-6%	-6%	11%	17%	3%	6%
December 31, 2012	0%	0%	0%	4%	1%	0%

PORTFOLIO UPDATE & COMMENTARY

<u>Positions</u>	<u>% of Portfolio</u>		<u>Change</u>	<u>Primary Theme</u>
	<u>3/31/2024</u>	<u>6/30/2024</u>		
NVIDIA Corp	20%	23%	3%	Software is Eating the World
Lantheus Holdings Inc	17%	19%	2%	Health
Celsius Holdings, Inc.	10%	9%	-1%	Health
Microsoft Corp	9%	8%	-1%	Software is Eating the World
CLI Studios**	5%	6%	1%	TIIA
Amazon.com Inc	5%	4%	0%	Software is Eating the World
EchoStar Corp*	5%	4%	-1%	Open RAN, 5G, & the Future of Wireless Networks
Altria Group Inc	4%	3%	0%	TIIA
Tesla Inc	2%	3%	1%	Metals Powered Economy
Taiwan Semiconductor Mfg. Co. Ltd.	2%	2%	0%	Semiconductors are the Teeth
Snowflake Inc	1%	1%	0%	Software is Eating the World
Zan Compute**	1%	1%	0%	TIIA
Lifelink**	3%	0%	-3%	TIIA
Stocks, Convertible Bonds, & Private Credit	82%	83%	0.6%	
<u>Cash</u>	<u>18%</u>	<u>17%</u>	-1%	
Total Partnership	100%	100%		

* Includes both equity and convertible debt

** Private Credit

The table on the previous page outlines TMCP's positions, along with the investment theme most applicable. The following section details portfolio changes.

Portfolio Updates

Nvidia: No changes. Stock split 1 for 10 during the quarter. See the Appendix for insights on hyperscale capital expenditures and Nvidia's top 5 risks.

Lantheus: No changes.

Celsius: Increased exposure via common shares and options.

Microsoft: No changes.

CLI Studios: Negotiated extension and increase in loan facility size, including additional warrants. Pleased with CLI's performance and ongoing relationship.

Amazon: No changes.

Echostar: No changes.

Altria: No changes.

Tesla: Increased exposure by selling puts and buying shares. Some puts recognized gains during the quarter; the remainder expire in September.

Taiwan Semiconductor: No changes.

Snowflake: No changes.

Zan Compute: No changes.

Lifelink: Loan repaid on April 16 after facing challenges at March 31 maturity. IRR: 16.35%. Warrants and put rights remain intact.

Private Credit

Our Partners are familiar with our allocation to private credit through Agility Capital, a Santa Barbara based venture debt fund. While not initially planned when we launched TMCP, these deals have been beneficial so as to expand your manager's circle of competence under experienced mentorship.

In April, our second venture loan was repaid by the borrower, Lifelink, yielding a 16.35% gross IRR. Like our first repaid loan (Gento) we continue to hold warrants that could prove valuable in the event of a liquidity event. This mid teens return compares favorably to the performance of the indices we present in these letters, as well as bond indices, considering the high equity valuations and low interest rates in effect at the time of origination.

These investments offer a secondary benefit: active participation in the underwriting process, thereby expanding our circle of competence in credit, technology, and startups.

Looking ahead, we may need to cease private credit activity within TMCP if an institutional investor requires it. Until that time comes, we will continue evaluating and participating in these deals opportunistically, further expanding our circle of competence.

TOP MARK HEALTH PARTNERS LP

Partnership Letter

For the Quarter ended June 30, 2024



PARTNERSHIP LETTER & PERFORMANCE UPDATE

Scientific discovery throughout history has often come in bursts. Every so often, a new idea comes along that completely changes our understanding of the world and leads to an explosion of discovery. Examples would include Copernicus' model of the solar system (centered around the Sun rather than Earth), Maxwell's theory to unify electricity and magnetism into a single force, Darwin's theory of evolution, or Einstein's general relativity (superseding Newtonian physics).

While perhaps not as Earth shattering as the above, a new idea came along in the 1970s that would rapidly move medicine forward. Prior to that time, new discoveries were made by chance, foraging for compounds existing in nature and testing their effects on the body. Then in 1976, Genentech was founded based on the novel idea of using recombinant DNA, allowing scientists to isolate and control gene expression which provided for the production of proteins as a therapy. The age of biotechnology began.

At the risk of overstating current discoveries, it feels as if we're at a new inflection point in medicine - a new age where a treatment for every affliction is possible. Recent discoveries have demonstrated techniques to not just splice DNA, but to precisely edit the genetic code itself. Coupling this with technologies such as organoids and AI, the ability to rapidly design wholly new proteins, print their genetic instructions, and test the interactions in organ models is within reach. In the not too distant future, we will have a computational model for every protein in the human body and the ability to predict both intended and unintended interactions of new proteins or molecules.

As we enter this new era of medicine, we hope these developments increase the success rate of clinical trials, reduce the overall cost for drug development programs, and enable us to target cures for previously incurable diseases. Top Mark Health Partners aims to be a participant along this reinvigoration of discovery.

TMHP CUMULATIVE RETURNS TO DATE⁹

To: 6-30-2024, From:	TMHP % Gross Return	TMHP % Net Return	S&P 500 Healthcare
6/30/2023	20%	18%	12%
Since Inception	26%	23%	17%
<i>Annualized Since Inception</i>	<i>20%</i>	<i>18%</i>	<i>13%</i>

Performance Update


Partnership interests in Top Mark Health Partners (TMHP) increased in price by 11.4% gross and 11.1% net in Q2. Over the last 12 months TMHP returned 20% gross and 18% net, compared to the 12% gain of the S&P 500 Healthcare Sector index.

During the quarter, the sole action taken by your managers was to continue building upon an existing position. Fortunately for us, the market did not begin to notice the value that we saw in this company until after this action was completed. That changed in the later portion of the quarter as I will expand on in the partners section below.

⁹ TMHP % Gross Return is at the fund level before management and performance fees; TMHP % Net Return is at the fund level, net of all fees. S5HLTH: S&P 500 Health Care. We include the S&P 500 Health Care Index as gauges against other potential investments. These are not explicit benchmarks but provide a cumulative perspective. Please review the Important Disclaimers on Page 46.

TMHP QUARTERLY RETURNS DATA¹⁰

Quarter Ended	TMHP % Gross Return	TMHP % Net Return	S&P 500 Healthcare
6/30/2024	11.4%	11.1%	-1%
3/31/2024	5.8%	5.2%	9%
12/31/2023	2.2%	1.7%	6%
9/30/2023	-0.1%	-0.5%	-3%
6/30/2023	4.6%	4.2%	4%



¹⁰ TMHP % Gross Return is at the fund level before management and performance fees; TMHP % Net Return is at the fund level, net of all fees. S5HLTH: S&P 500 Health Care. Please review the Important Disclaimers on Page 46.

PORTFOLIO UPDATE & COMMENTARY

Below you will find commentary on each of the Health Partnership's holdings.

Additional context can be found in previous partnership letters.

<u>Positions</u>	<u>% of Portfolio</u>		<u>Change</u>
	<u>3/31/2024</u>	<u>6/30/2024</u>	
Lantheus Holdings Inc	24%	28%	3.8%
Vertex Pharmaceuticals Incorporated	23%	23%	0.2%
Celsius Holdings, Inc.	12%	8%	-4.5%
Axsome Therapeutics Inc	7%	7%	-1%
Harrow Inc	6%	13%	7%
Exelixis Inc	2%	2%	0%
Stocks	74%	79%	5.3%
Cash	26%	21%	-5%
Total Partnership	100%	100%	

Vertex Pharmaceuticals

In the first quarter of 2024, Vertex publicly announced positive phase 3 study results for their experimental non-opioid pain treatment, VX-548. Subsequently, the FDA agreed to accept Vertex's new drug application as a rolling submission. A rolling submission allows the company to pass fully analyzed data to the FDA as it becomes available, rather than waiting for a polished FDA submission to be drafted. This puts data into the FDA's hands sooner so they may begin their process at an accelerated pace. That rolling submission was completed on July 30th with the FDA responding that they will make a decision by January 30th, 2025.

Also in the quarter, Vertex has made their largest acquisition to date with the purchase of Alpine Immune Sciences for \$4.9b. The main asset acquired is povetacicept

which treats IgA nephropathy, an autoimmune disease affecting the kidneys. This adds to Vertex's pipeline which includes inaxaplin, a potential treatment for APOL1-mediated kidney disease. Both diseases manifest with inflammation disrupting the kidney's ability to function. Additionally, both diseases are undertreated with no approved therapies for either that address the underlying mechanisms of the diseases.

The patient population size and value to the lives of these individuals is very similar to the population impacted by cystic fibrosis. If Vertex can bring these two drugs to market (one in phase 3 trials, the other entering phase 3) this could potentially launch a kidney franchise of a similar impact to the business as the CF franchise.

The last Vertex update of note pertains to a potential cure for Type 1 Diabetes. Vertex released very early stage results from a small phase 1/2 study on their stem cell derived, insulin-producing islet cells. The initial phase of study saw twelve patients treated, with eleven of them fully insulin independent, functionally cured of T1D. I'll reiterate that this is very early stage, but worth highlighting due to the impact it can make on the two million Americans that have this disease.

Lantheus Holdings

Our thesis for Lantheus' flagship product, Pylarify, continues to hold. We believe that advanced imaging use for prostate cancer diagnosis will continue to grow, and that regardless of which next-generation prostate cancer therapies are approved, this trend will continue.

An aspect we enjoy about this company is their efficiency with research spending. So when they reported a quadrupling of capital allocated to R&D this quarter, we were a bit taken aback. However, the business has grown their cash flow to a point where even at this elevated level, they are still spending right at our desired amount - no more than 50% of free cash flow.

However, this does shift our thinking about what we own in this business. Rather than owning a growing stream of cash flow, we are now attributing more value to their pipeline of products. We've come to believe that Lantheus is making the correct decision, and certainly a rational one.

Lantheus has established themselves as a leader in radiopharmaceuticals. And as a leader vying to become the leader, they should leverage their position into a robust product portfolio in order to enable scale efficiencies in both manufacturing and distribution and to further corner expertise in this burgeoning field. In addition to Pylarify, their pipeline includes several other oncological radioligand therapies and diagnostics as well as new acquisitions to build a radiodiagnostic portfolio for Alzheimer's disease. By virtue of many of the larger pharmaceutical companies not entering this space, it leaves the opportunity open for Lantheus to capture the market.

Celsius Holdings

The energy drink market has been under pressure as of late, shrinking for the last two months as the consumer pulls back. This is highlighted by the largest convenience store operator, key to energy drink distribution, seeing overall sales down 4% in North America.

In the quarter, Celsius market share dropped from 11.5% to 11%. This is obviously very disappointing for a growing energy brand. We started to see anomalies in the weekly reported industry sales data whereas Celsius and Monster saw decreases in share, with Red Bull rapidly capturing that lost share. Later we found that a video had been circulating social media promoting a pseudoscientific opinion on cyanide content in both Celsius and Monster. Poisoning your customer does not seem like an optimal business model, and fortunately both companies have debunked the conspiracy theory. At this time, the loss of share has stabilized, and we are keenly anticipating growth once

again. We believe this will be a short-term setback and will continue to keep a close eye on the situation.

A bright spot continues to be the Amazon and club channel sales. In the quarter those were up 41% and 30% respectively with Celsius remaining the #1 seller in Amazon's energy category. We view this as the potential of where Celsius' overall market share could go. Additionally the past quarter saw the start of their international expansion into Canada, the United Kingdom, Ireland, Australia, and New Zealand. Canada was the first to launch with the brand rapidly achieving 5.5% market share in just one quarter.

Harrow Health

Harrow continues to make excellent progress in the commercialization of their drug portfolio. Sales over the first half of the year have grown in the low-40% range with management's expectation that growth will accelerate into the back half of the year. To meet the company's 5-year strategic plan, growth needs to continue at this pace for several more years. If that does occur, the upside in this investment is exceptional.

Why might that lofty 5-year plan succeed? Harrow is in the early stages of commercialization where prescriber adoption of new drugs typically follows an S-curve. Studies have shown that 64% of physicians will not even consider prescribing a new drug in the first 15 months on the market¹¹. Iheezo has now been prescribed for four quarters and Vevye for just two. Additionally, only half of commercial insurance provides coverage and not all state medicaid plans provide access. As Harrow continues its commercialization efforts and physicians find these to be differentiated products, they will begin to tip into the more rapid adoption rates of the S-curve.

¹¹ Anderson et al., "Patterns and Predictors of Physician Adoption of New Cardiovascular Drugs."

As mentioned in the previous section, we were fortunate enough to build our position in Harrow without time to spare. The market began to realize the value held in this company with shares nearly doubling from our basis by the end of the quarter. Strong prescription data was likely the driver with patients exhibiting a much stronger refill rate with Vevye than those of competitors. Another explanation, of course, is that investors have been listening to Mike and I on the Telltales podcast.

Axsome Therapeutics

Axsome's lead product, Auvelity, continues to grow new and continuing prescriptions at a steady pace. Analysis of insurance claims data now backs up physician's anecdotal evidence of patients continuing on with Auvelity at a higher rate than when prescribed an SSRI.

Axsome has recently re-submitted their new drug application for AXS-07, a novel drug to treat acute migraine. This market is similar to depression treatment with 70% of patients dissatisfied with their current therapy. And like finding the SSRI that works best for you, most patients cycle through migraine treatments hoping one provides relief.

The new drug application for AXS-14 is currently being compiled by the company and is expected to be submitted in Q3. AXS-14, if approved, will be used to treat fibromyalgia. The drug met its primary study endpoints with the company stating that it provided positive and statistically significant results in participants. Fibromyalgia remains a challenge to treat with most approved therapies being pain relievers, antidepressants that also have an effect on the condition, or behavioral changes, and having no novel treatments approved in almost 20 years. If AXS-14 does provide a

meaningful improvement over other treatment options as the team believes, the market potential is large with roughly five million Americans affected by the disease.

The pivotal study of AXS-05 (Auvelity) in the treatment of Alzheimer's agitation continues to have strong enrollment. The approval of a competing product, Rexulti, was anticipated to slow enrollment in this study but it does remain on track with data to be read out later this year.

There are a number of other programs also expecting to complete this year or next that only add to the value of Axsome's pipeline. Taking the industry average of 45% success rates in phase 3 studies, we would expect Axsome to bring three new treatments to market. This leads us to believe the potential of this pipeline, on top of Auvelity and Sunosi, presents a good value.

Exelixis

Exelixis is an oncology-focused biotech company that is presently a research position in TMHP. The company's goal is to improve the standard of care for cancer patients. Its sole FDA approved product, Cabometyx, is a tyrosine kinases inhibitor (TKI). TKIs block the signaling pathway for cell growth and division. This has the effect of slowing or stopping tumor growth and potentially putting a cancer in remission.

Cabometyx is a key treatment for a number of cancers, however, it has a patent dispute that must be settled before a full position can even be considered.

Cabozantinib, the compound behind Cabometyx, has a patent portfolio that extends to 2031, but an initial chemical compound patent that expires in 2026. A compounding pharmacy is attempting to invalidate the later patents in order to produce the generic version in 2026. If upheld, Exelixis has already set the precedent for generic production

in 2031 through an agreement with Teva Pharmaceuticals, the world's largest generic producer.

Exelixis needs the patents to be upheld in order to maintain cash flows that are financing their ongoing research, including a newer variation of Cabometyx called Zanzalintinib. In testing, it is proving to be just as effective, but with better tolerability and a shorter half-life. The half-life of Zanzalintinib is one-fifth of Cabometyx, allowing the patient to cycle off and on the drug faster when other treatments or diagnostic procedures are required. Exelixis believes this new drug has the potential to be the best in class TKI if approved.

So while we like the potential for Exelixis, the Delaware court needs to issue their ruling before we consider a full position. We expect that the ruling will come this quarter.



APPENDIX




HYPERSCALE CAPEX AND NVIDIA'S 5 RISKS

Theme: Software is Eating the World

Frequent readers of our letters will recall that we have been beating the drum on the prospects for artificial intelligence (AI) to reshape the software industry since 2015.

Today, generative AI is at the cutting edge of this trend, and Nvidia's jaw dropping cash flow growth has stunned even us, its most ardent supporters. Nvidia represents at least ¼ of data center market¹², captures ~50% of total hyperscaler capital expenditures¹³, and

Figure 1: Data Center Market Share



Source: Digits to Dollars

¹² D2D Advisory, "No Going Back."

¹³ Note that Meta is the only one of these companies that is not a public cloud provider.

achieves margins in excess of 75%. Needless to say, many investors that missed the trend are scratching their heads.

“Trees don’t grow to the sky.”

- German Proverb

Trees don’t grow to the sky. Or do they? Nvidia reports earnings next week and is expected to hit a run rate of \$100 billion in data center revenue. **This article will assess the state of data center spending, through the lens of the hyperscalers (Amazon, Google, Meta, Microsoft, and Oracle), and suggest our top five risks to Nvidia’s dominance in this market.**

“Software is eating the world” is the phrase we adopted to represent the long term trend of software encompassing all aspects of our professional and personal lives¹⁴. The shift from CPU based computing to GPU based computing fits squarely within this trend, and artificial intelligence is its largest accelerant. Nvidia CEO and co-founder Jensen Huang, has suggested that his vision for this shift in workloads dates back to the company’s founding in 1993. But, back then, and for the next 25 years or so, Intel’s CPUs would dominate the industry with market share figures similar to those of Nvidia today.

Traditional CPUs execute tasks sequentially using a single processing core. This is very effective for sequential programs with interdependencies. But for large tasks that can be divided into smaller sub-tasks, parallel computing increases overall processing speed and efficiency by spreading the work across multiple cores. Graphical rendering,

¹⁴ Andreessen, “Why Software Is Eating the World.”

which requires individual calculations for each pixel on a screen, was an early use case - hence the name Graphical Processing Unit or GPU. The computing shift from CPU to GPU wouldn't gain much steam until 2012 when AlexNet, a convolutional neural network written with CUDA¹⁵, running on a Nvidia GPU, trounced the ImageNet Large Scale Visual Recognition Challenge (ILSVRC)¹⁶. ILSVRC pitted software programs against each other, to see which was most capable of correctly classifying objects and scenes in images. Many of the leading minds in the field of artificial intelligence can trace their experience, in one way or another, back to this AlexNet project.

Through the 2010's, Nvidia, and notably CUDA, chipped away at the CPU market, shifting workloads that were entirely impossible or exorbitantly expensive to run on CPUs, like weather forecasting, fluid dynamics simulations, and other scientific problems. AI and machine learning (AI/ML) were also major factors in this steady climb. Meta's recommender system, for example, predicts user behavior and ranks content in real-time, across the company's "Family of Apps". It is likely the most profitable AI workload in production today. But it wasn't until OpenAI released its large language model (LLM), ChatGPT, in November 2022, that the world woke up to the power of "AI" models. This triggered data center spending shifted quickly to Nvidia as the core of the data center shifted to GPUs rather than CPUs (see Figures 1 and 2).

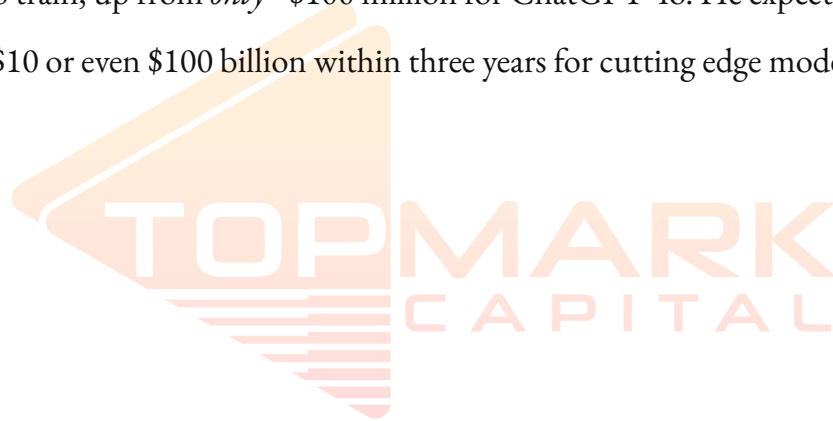
Inference - the actual running of these LLMs - is the fastest growing parallelized workload today. These include ChatGPT, Claude, Gemini and other 'assistants', as well as APIs that OpenAI, Anthropic, Google, and others provide for developers to incorporate into their products. Walmart, for example, utilizes generative AI to

¹⁵ CUDA is Nvidia's proprietary software language used to program the company's GPUs

¹⁶ AlexNet was not the first AI model to run on GPUs, but it was the most notable given the *relative* fame of the ImageNet competition. All other competitors were running CPU based models.

dynamically edit product descriptions. Some companies, like Goldman Sachs, are using a combination of open source models (like Meta’s Llama) and Nvidia hardware to develop custom solutions using proprietary data¹⁷. Regardless of the implementation method, these LLMs act like a computer version of the human brain. They break down data into tiny pieces, analyze them in hidden layers, and make predictions based on what they’ve learned.


Training these models is the highest value and most compute-intensive task. Importantly, training heavily influences future hyperscale capital spending. Anthropic’s CEO, Dario Amodei, says large language models in development today can cost up to \$1 billion to train, up from *only* ~\$100 million for ChatGPT-4o. He expects the cost to balloon to \$10 or even \$100 billion within three years for cutting edge models¹⁸.



¹⁷ “Goldman Sachs CIO on How the Bank Is Actually Using AI.”

¹⁸ Morales, “AI Models That Cost \$1 Billion to Train Are Underway, \$100 Billion Models Coming — Largest Current Models Take ‘only’ \$100 Million to Train.”

Figure 2: Hyperscale Capex & Chipmaker Share



The Hyperscaler Investment Decision

For hyperscalers to justify the costs associated with investing in the chips and infrastructure required to support each successive generation of models, they will need to continue to see an attractive return on investment (ROI) and a short payback period. Colette Kress, CFO of Nvidia made this abundantly clear on the company's most recent earnings call:

“Training and inferencing AI on NVIDIA CUDA is driving meaningful acceleration in cloud rental revenue growth, delivering an immediate and strong return on cloud provider's investment.”

So what is the ROI and payback period for a Nvidia H100 cluster for a hyperscaler? Jensen Hwang has suggested that hyperscalers can expect a 400% ROI over the course of five years. A more thorough analysis suggests that Microsoft, for example, expects an internal rate of return (IRR) between 38% and 110% over five years, and a payback period between 10 months and 2.1 years. The variation depends on if the usage is ‘pay as you go’ or on a five year fixed contract¹⁹.


Demand ultimately drives hyperscaler purchasing decisions. GPU availability at public clouds can give us indications as to how much demand there is for GPUs. Figure 3, on the following page, shows A100 demand moderately easing as H100 capacity started to come online. What is quite remarkable is that the A100 first started to ship in May 2020! This means, four years into its lifecycle it was still practically impossible to get access to one. Every hyperscaler with A100s likely earned phenomenal ROIs.

The public cloud hyperscaler investment decision will be heavily weighted toward payback period. Having line of sight on the payback period creates a “heads I win, tails I

¹⁹ “Return on Investment (ROI) on AI for Cloud Hyperscalers.”

don't lose" scenario. If average payback periods creep much beyond two years, I would expect that the CFOs of the hyperscalers will tighten the reins on capital spending. This also means that so long as there is demand for GPUs, all public cloud hyperscalers will want as many Nvidia GPUs as they can get their hands on.

Figure 3: Mean GPU Availability Percentage by Hour/Day/Month



Nvidia's 5 Biggest Risks:

Today, Nvidia dominates not only the market for GPUs, but also nearly everything inside the data center, thanks to its acquisition of Mellanox. The company's enviable position today is largely dependent on Jensen Hwang's vision for the shift from CPU to GPU based computing that lead him to invest in CUDA, Nvidia's programming language that developers use to code parallelized applications that run on Nvidia GPUs²⁰. CUDA gave Nvidia a 10+ year lead on competitors and created network effects in that every parallelized computing programmer has used CUDA. Defending this position into the future will depend on Nvidia maintaining its competitive advantage versus a litany of would be competitors, while simultaneously navigating the natural ebbs and flows that result from the underlying applications running on the company's GPUs. Below we list and detail our top five risks to Nvidia's dominance. Since the hyperscalers are ultimately the primary purchasers of these data centers, our analysis will be through the lens of a hyperscaler. Risks 1 and 2 cover the supply side decisions hyperscalers face, and risks 3-5 cover the demand side, which the hyperscalers will have to estimate.


1. Hyperscalers Overestimate Demand

The case of investing in a Nvidia GPU cluster has been a very easy decision because the payback period is so short (<2 years, assuming the aforementioned assessment of a Nvidia H100 cluster at Microsoft is remotely accurate). When these servers are new, more of the resources are dedicated to training, therefore having line of sight on demand for training enables hyperscalers to confidently make this investment decision. This confidence breaks down if the hyperscaler doesn't have line of sight on the demand

²⁰ "NVIDIA's CORNERSTONE for UNMATCHED DOMINANCE in AI (CUDA)."

needed to cover its cost. Lower utilization and less confidence in the first two years of the forecasted demand will make this investment decision less obvious.

Training is the primary driver of adoption when new GPUs hit the market due to the level of compute intensity required to train frontier models. As those models are



Source: Nvidia Company Overview February 21, 2024

deployed into production, the inference process can run on the same GPUs²¹.

Estimating this inference demand is much harder, given that we don't even know the capabilities of the next generation of models²². As the supply of leading edge Nvidia GPUs normalizes, hyperscalers may be tempted to overinvest in this equipment and accept lower rates of return. Finance departments at every hyperscale public cloud look

²¹ There are of course exceptions to this rule, but in general this is what we are seeing

²² One could argue that we don't truly understand the capabilities of the current frontier models.

at their investment decisions similarly. They estimate the cost of the project, then estimate the cash flows expected over the life of the investment. Typically, they have an internal hurdle rate that a project should meet. Since estimates in years 3-5 greatly influence the ending IRR, there is potential that hyperscalers could get too aggressive with their GPU buildouts. If these investments fail to meet prior target hurdle rates, hyperscaler finance departments may become more conservative in adopting the next generation of GPUs.

2. Hyperscalers are Overestimating the Useful Life of GPU Servers

Historically, servers are depreciated in a straight line basis over a three year period. However, starting in 2020, hyperscalers began stretching their useful life to as much as six years. This made sense. Intel had fallen off the leading edge and successive generations of CPUs were only marginally better than the prior generations. Moore's law was dead. However, as documented in this memo, hyperscale capex has shifted to Nvidia GPUs, which are increasing in performance at almost twice the rate of CPUs. Meaning, Nvidia's innovation cycle may render its older GPUs worthless faster. The implications of this would likely mean higher prices for cloud based computing, and cause

3. The Performance of Successive Generations of Frontier Models Do Not Scale

Successive generations of LLMs are generally improved today by increasing the amount of data (tokens) used to train the model and the parameters (complexity) of the model. As a result, new generations models are dependent on new generations of Nvidia hardware (A100 -> H100 -> B200). But, if performance increases of successive generations of models begin to taper, then we can expect that the appetite for investors to finance the costs associated with training new models will wane. At least until a new

approach arises that significantly moves the ball forward. This is a risk to Nvidia's data center sales, as AI could hit a proverbial "brick wall".

One should also remember that Moore's law was not a law at all. It was a voluntary challenge that Intel managed to meet year after year, until 2016 when it struggled to implement 10 nanometer technology²³. Performance scaling of generative AI models will similarly be 'optional' and driven by the few companies that are developing foundational models. Improving capabilities are a necessity for Nvidia to maintain its share of data center, let alone expected growth.

What would cause the performance of successive generations of models to taper? The primary concern today is data availability. Successive generations of Large Language Models (LLMs) have shown that increasing the size of the training data set also increases the quality of the model output. Today's models are like giant vacuum cleaners, sucking up every bit of data they can. Web 2.0 companies are quickly realizing that they hold significant value for these large language models, hence the \$60 million per year license fee Google is paying Reddit. Expect to see more deals like this, but don't underestimate the long term implications of this shift. In one sense, the internet as we know it today - the free and open internet - is dying. Interestingly, the sheer volume of data on the free and open internet is what made training these models feasible in the first place. Innovation in software, like synthetic data generation, transfer learning from data-rich domains, and data efficiency improvements are seen as potential solutions to overcome these data limitations. Fortunately, the total amount of data in the world doubles every two years, so we can expect that there will be lots of new information to train new models on. We hope that future AI models will have access to this data.

²³ Woods, "The Death of Moore's Law."

4. Hardware is No Longer the Limiting Factor for AI Scaling

Assuming that data availability and hardware capability (see 3 above) are not barriers to AI scaling, what other factors may hinder AI scaling? There are many, but we will focus on two here: training efficiency and power availability.

Training efficiency is the amount of compute and training time needed to achieve superior model performance versus the current state of the art. Increasingly powerful compute resources can mask inefficient software, as it did for Intel so long as Moore's law maintained its pace. Only when Moore's law started to break down (and total cost of ownership became increasingly important) did ARM based RISC CPU architectures become competitive. Are there opportunities for improving the software related to training that would render as good or better results than the current transformer architectures? It is fairly well documented that there are significant opportunities for improvement in the efficiency of AI software, but it is hard to ignore the history of Intel and Moore's law. We would argue that we are most likely to see efficiency innovations in AI software come from China. China is limited in its ability to adopt the latest semiconductor technology due to export restrictions. These export restrictions create a strong incentive for innovation within the Chinese domestic chip industry, but they also create an incentive to eke out improvements in other areas of the AI stack - notably software. These circumstances put China in a position where scarcity (of compute) may breed innovation²⁴.

A second factor that may hinder AI scaling is power availability. New data centers are facing significant delays of up to 5 years in getting power interconnections due to transmission infrastructure constraints, massive interconnection queue backlogs, and surging data center power demand. While some data centers are exploring alternative

²⁴ Notably, this creates the exact opposite effect intended by those who implemented the export restrictions

solutions like on-site generation²⁵ and co-locating with power plants²⁶, major transmission upgrades and interconnection process reforms will be critical to support the industry's long-term growth. The rapid increase in data center energy consumption increases risks associated with already problematic grid reliability. These issues are already becoming a major bottleneck for new data centers, and could ultimately hinder Nvidia's growth.

5. Competition

Nvidia currently dominates the data center market (see Figures 1 and 2). There are a litany of competitors that are attempting to gain an edge. AMD, Nvidia's long time competitor in the gaming GPU market, has released the MI300X GPU, Intel has launched the Gaudi 3 GPU, and Microsoft, Google, Amazon, and Meta are developing in house chips. Nonetheless, customers still heavily favor the H100 and only resort to alternatives if they cannot get access to H100s. These other products are currently not competitive with Nvidia. Nvidia has established process leadership and scale, enabling them to own practically all of the market. They also have strong network effects due to the fact that the entirety of AI software to date has been developed on CUDA. For Nvidia to maintain its dominance, it needs to maintain a rapid pace of innovation, releasing successive generations of chips that are not only better than previous generations, but also well ahead of their competition in the same way that Intel was able to dominate the x86 CPU market for decades (see figure 4).

A final, and frankly more likely, though longer term, competitive threat would come from an entirely new paradigm shift. In the same way that the CPU gave way to the

²⁵ Morales, "Elon Musk Powers New 'World's Fastest AI Data Center" with Gargantuan Portable Power Generators to Sidestep Electricity Supply Constraints."

²⁶ Fitch, "Why the AI Industry's Thirst for New Data Centers Can't Be Satisfied."

GPU due to the transition of workloads, quantum computing is widely seen as the next potential disruptive technology in the data center market. Quantum computers leverage principles of quantum mechanics like superposition and entanglement to perform certain computations exponentially faster than classical computers. Many computationally demanding problems don't parallelize efficiently due to sequential dependencies and communication bottlenecks between parallel processors. Quantum algorithms can perform certain complex computations with far fewer operations, reducing latency and energy consumption compared to massively parallel classical solutions. Problems with high combinatorial complexity, like optimization, molecular simulation, and machine learning, could see transformative speedups with quantum computing. However, quantum computing still faces many challenges in scaling qubits, improving error correction, and integrating with classical infrastructure. Nonetheless, the technology continues to advance, and it represents both an opportunity and a competitive threat that Nvidia will need to navigate strategically. With that said, Jensen's vision being forever forward looking, has pushed Nvidia to already begin work on CUDA-Q, their library for quantum workloads²⁷.

²⁷ "CUDA-Q."

Takeaways

Before we close, I'd like to draw your attention back to our initial premise: the data center market has shifted dramatically due to competition and changing workloads. Currently, there is only one player of importance in the data center market: Nvidia. As long as Nvidia stays substantially ahead of would be competitors, in the same way Intel stayed far ahead of AMD for decades, their position as the leader in AI hardware is secure. On the topic of changing workloads, we expect more AI/ML and training workloads in the future than we have today. In the same way that the typewriter became obsolete as a result of the PC and the iPod became obsolete because of the iPhone, we view generative AI as distinct from earlier versions of "AI" (machine learning, computer vision, etc) in that it has general purpose applications. As a result, generative AI will render many things obsolete (assuming performance continues to scale). Today's copilots and other AI based features (which are becoming table stakes in most software) will lead to completely different processes and business models across industries. This is a technological paradigm shift. The hyperscalers see this and are intent on investing heavily²⁸. Generative AI has the ability to completely change the technology landscape, and ultimately render the cornerstones of cash flow generation in the technology

²⁸ Hyperscaler quotes from most recent earnings calls...

Mark Zuckerberg, Meta - *"Our expectation, obviously again, is that we are going to significantly increase our investments in AI infrastructure next year, and we'll give further guidance as appropriate."*

Sundar Pichai, Google - *"...the risk of under-investing is dramatically greater than the risk of over-investing for us here, even in scenarios where if it turns out that we are over-investing, [inaudible] these are infrastructure which are widely useful for us"*

Satya Nadella, Microsoft - *"To meet the growing demand signal for our AI and Cloud products, we will scale our infrastructure investments with FY '25 capital expenditures expected to be higher than FY '24. As a reminder, these expenditures are dependent on demand signals and adoption of our services that will be managed through the year."*

Elon Musk, Tesla - *"I'm quite concerned about actually being able to get [inaudible] GPUs and when we want them."*

industry (Google search, Microsoft Office, and Facebook/Instagram) obsolete. For each of these companies²⁹, their relentless pursuit of AI is purely selfish - they believe that the cost of being left behind far outweighs the cost of investing.


In conclusion, we see the market for hyperscale compute evolving rapidly. Nonetheless, our long term view remains intact: Nvidia³⁰ is the clear leader in the space, and there will be significantly more GPU based workloads in the future than there are today.



²⁹ Note that Apple, who generates more free cash flow than all of these companies is not exempt from the existential implications generative AI could have on its core iPhone business. Apple's current strategy is to provide the platform (iPhone) by which users access LLMs. This is a more rational approach from a cash flow generation perspective, but also leaves the door open for missing the (gen AI) boat entirely.

³⁰ Note that this discussion does not cover the all important topic of valuation

A BRIEF FOLLOW-UP ON YIELD CURVE INVERSIONS



In Q1 of 2023 we published an essay titled *A Brief History of Yield Curve Inversions* where we analyze the timelines involved from an inversion to a recession. While not all yield curve inversions result in a recession, all recessions have been preceded by yield curve inversions. And as far as the magnitude of an inversion goes, the present was as deep as they generally get. That initial analysis, and this follow-up, admittedly use a small, but very important, slice of data.

As outlined in the original essay, the stock market does rather well following an inversion, although it is over a finite timespan and typically ends in a steep drop. Over

the last 50 years, the mean gain in the S&P 500 following an inversion was 25%. As we write this update, the SPX is up 40% from that event, but when the market peaked in July it was a loftier 48% gain. This, however, is not the largest post-inversion run of recent times. That distinction belongs to the 1978 inversion that saw the index gain 52% to its peak.

This time does seem different. I imagine in past Fed tightening cycles that Congress has not worked so hard to generate spending counter to their efforts. What this means is anybody's guess. It is plausible that it has acted to extend the market's performance in this particular cycle.

History would say that we've had an exceptionally strong run post-inversion and odds are the market peaked in July. The timing of past cycles indicates that recessions begin within months of the market finding a top (the markets being forward looking), and tends to align with the yield curve normalizing, as it currently appears to be doing. Normalization is typically caused by the 2-year rate falling as the shorter end of the curve begins to anticipate Fed rate cuts. Often the Fed is reactionary, late to cut, and a signal of pain to come. If we look back sometime next year and determine the recession started in Q3 of this year, it will have begun entirely within historical norms being one standard deviation from expectations.

IMPORTANT DISCLAIMER

Past performance is no guarantee of future results. Investing in equities and fixed income involves risk, including the possible loss of principal. The investment performance presentation contains historical data and information relating to the performance of certain investments. These figures should not be considered as a guarantee or a reliable indicator of future performance. Investment returns and the value of an investment can go up or down, and there is no assurance that any investment strategy will achieve its objectives, generate profits, or avoid losses. Investing in financial markets inherently involves a certain degree of risk and speculation. The value of investments, and the income generated from them, can fluctuate due to various factors, including but not limited to market conditions, economic changes, interest rates, and political events. As such, there is always the potential for loss, and you should only invest funds that you can afford to lose. The S&P 500 Index is included to allow you to compare your returns against an unmanaged capitalization-weighted index of 500 stocks designed to measure performance of the broad domestic economy through changes in the aggregate market value of the 500 stocks representing all major industries. The NASDAQ 100 Index measures the performance of the largest 100 stocks on the NASDAQ.

IMPORTANT DISCLOSURES. The information herein is provided by Top Mark Capital Management LLC ("Top Mark Capital") and: (a) is for general, informational purposes only; (b) is not tailored to the specific investment needs of any specific person or entity; and (c) should not be construed as investment advice. Top Mark Capital makes no representation with respect to the accuracy, completeness or timeliness of the information herein. Top Mark Capital assumes no obligation to update or revise such information. In addition, certain information herein has been provided by and/or is based on third party sources, and, although Top Mark Capital believes this information to be reliable, Top Mark Capital has not independently verified such information and is not responsible for third-party errors. You should not assume that any investment discussed herein will be profitable or that any investment decisions in the future will be profitable. Investing in securities involves risk, including the possible loss of principal.

REFERENCES

- AI Investing. “Return on Investment (ROI) on AI for Cloud Hyperscalers.” Substack newsletter, August 8, 2024. <https://aiinvesting.substack.com/p/return-on-investment-roi-on-ai-for>.
- Anderson, Timothy S., Wei-Hsuan Lo-Ciganic, Walid F. Gellad, Rouxin Zhang, Haiden A. Huskamp, Nitesh K. Choudhry, Chung-Chou H Chang, et al. “Patterns and Predictors of Physician Adoption of New Cardiovascular Drugs.” *Healthcare (Amsterdam, Netherlands)* 6, no. 1 (March 2018): 33–40. <https://doi.org/10.1016/j.hjdsi.2017.09.004>.
- Andreessen, Marc. “Why Software Is Eating the World.” Andreessen Horowitz, August 20, 2011. <https://a16z.com/2011/08/20/why-software-is-eating-the-world/>.
- D2D Advisory. “No Going Back: The New Data Center.” Digits to Dollars, December 1, 2023. <https://digitstodollars.com/2023/12/01/no-going-back-the-new-data-center/>.
- Fitch, Tom Dotan and Asa. “Why the AI Industry’s Thirst for New Data Centers Can’t Be Satisfied.” WSJ, April 24, 2024. <https://www.wsj.com/tech/ai/why-the-ai-industrys-thirst-for-new-data-centers-cant-be-satisfied-93c7eff5>.
- “Goldman Sachs CIO on How the Bank Is Actually Using AI.” Odd Lots. Accessed August 13, 2024. <https://omny.fm/shows/odd-lots/080624-odd-lots-marco-argenti-v1>.
- Morales, Jowi. “AI Models That Cost \$1 Billion to Train Are Underway, \$100 Billion Models Coming — Largest Current Models Take ‘only’ \$100 Million to Train: Anthropic CEO.” Tom’s Hardware, July 7, 2024. <https://www.tomshardware.com/tech-industry/artificial-intelligence/ai-models-that-cost-dollar1-billion-to-train-are-in-development-dollar100-billion-models-coming-soon-largest-current-models-take-only-dollar100-million-to-train-anthropic-ceo>.
- Morales, Jowi. “Elon Musk Powers New ‘World’s Fastest AI Data Center’ with Gargantuan Portable Power Generators to Sidestep Electricity Supply Constraints.” Tom’s Hardware, July 24, 2024. <https://www.tomshardware.com/tech-industry/artificial-intelligence/elon-musks-new-worlds-fastest-ai-data-center-is-powered-by-massive-portable-power-generators-to-sidestep-electricity-supply-constraints>.
- NVIDIA Developer. “CUDA-Q.” Accessed August 23, 2024. <https://developer.nvidia.com/cuda-q>.
- Top Mark Capital | LinkedIn. “NVIDIA’s CORNERSTONE for UNMATCHED DOMINANCE in AI (CUDA),” August 23, 2023. <https://www.linkedin.com/pulse/nvidias-cornerstone-unmatched-dominance-ai-cuda-top-mark-capital/>.
- Woods, Audrey. “The Death of Moore’s Law: What It Means and What Might Fill the Gap Going Forward | CSAIL Alliances.” Accessed August 22, 2024. <https://cap.csail.mit.edu/death-moores-law-what-it-means-and-what-might-fill-gap-going-forward>.