J. STERN & CO.

Investing in Quality for the Long-Term

Investment Insight

EXPERIENCING THE AI PARADIGM SHIFT FIRSTHAND: NOTES FROM A TRIP TO THE WEST COAST

We recently spent a week on the West Coast of the US meeting with a number of technology companies in Silicon Valley and Seattle, all deeply involved in artificial intelligence and the broader AI ecosystem.

When a new theme like AI reaches the forefront of public markets, it creates a rush to call the top. Negative headlines, particularly around Nvidia, travel quickly. We base our investment decisions on facts and our fundamental research into the companies driving innovation and disruption has led us to consistently maintain our conviction in AI. In September 2023 (We are not selling Nvidia) and again in February 2025 (DeepSeek: New AI apps and the inexorable demand for computing capacity), we provided arguments and evidence for our positive long-term view amid the waves of negative sentiment.

Our recent company meetings in the US and the breadth and depth of our discussions have reinforced our conviction that AI is still in its early days and will continue to be one of the most compelling structural investment opportunities of the coming decade.

Source: Giles Tulloch

Giles Tulloch, partner and analyst, at Nvidia's HQ in Santa Clara, CA

Nvidia's competitive moat

Our recent meetings, including conversations with Nvidia and several of its customers and competitors, have deepened our confidence in both the company and the opportunity in AI. Nvidia continues to dominate the AI headlines. We first bought it in February 2022 and it is the largest holding in our World Stars Global Equity strategy.

We have written extensively about Nvidia's competitive positioning. While the technology sector is often seen as highly competitive with low switching costs, Nvidia has built a robust and durable ecosystem that is misunderstood by many outside the industry. There is a popular tendency to compare today's leaders with past giants like Kodak, Blockbuster or BlackBerry, companies that were disrupted and displaced as they failed to adapt to innovation and change. Such analogies ignore that companies like Alphabet and Meta have continued to innovate and improve what they do and have been able to grow revenues and profits in their core businesses despite competition being 'just one click away'. The threat to leaders is not incremental competition but entirely new paradigms. That is why companies like Alphabet and Meta are pouring billions of dollars into AI to ensure that they lead the next wave of innovation and are not overtaken by it.

Our meetings with Nvidia have reinforced several key components of its competitive moat and why competition is not 'just one chip away.'

First, Nvidia's pace of innovation. Customers and competitors consistently praised Nvidia's rapid product cadence and clear long-term roadmap, which enables them to plan their infrastructure spend with visibility and confidence. Nvidia's backwards compatibility and hardware fungibility extend the lifecycle of its products, giving users flexibility as their workloads evolve from training large models to deploying them for inference.

Second, Nvidia's software platform CUDA. Its deep integration across the developer ecosystem creates significant inertia. Competitors openly acknowledged the network effect this has created because switching to an alternative stack is prohibitively costly. Only the largest companies like the hyperscalers Alphabet, Amazon and Microsoft can afford to dedicate engineering teams to building or using other architectures.

Third, Nvidia's full-stack approach from silicon to software to networking. With inference workloads becoming more compute-intensive, driven by complex reasoning engines and real-time responsiveness, the need for high-performance networking and tight integration is rising. Nvidia's rack-level solutions were repeatedly cited as a competitive differentiator.

As part of our investment process, we continuously look for facts to falsify our hypotheses. For Nvidia, the biggest threat is a new product that would reduce the demand for its high-performance computing products and solution. The most credible threat comes from 'custom silicon' developed by the hyperscalers. Alphabet, for example, created its own AI accelerator application Tensor Processing Units (TPUs) in 2015, while Amazon acquired Annapurna Labs in the same year to help build its Trainium and Inferentia chips. More recently, third-party chipmakers like Marvell and Broadcom have partnered with the hyperscalers to co-develop specialised chips.

Nvidia's GPUs are general-purpose accelerators, with AI being one of their most transformative applications. They are different from custom chips (ASIC or application-specific integrated circuits) that are designed for discrete and highly optimised workloads. For example, Amazon might deploy a specific chip to accelerate webpage load times on its cloud computing platform AWS, an isolated task that benefits from bespoke hardware. These ASICs can and often do operate alongside Nvidia GPUs. As the GPU is highly fungible and versatile, it remains essential for more dynamic and evolving workloads, especially those related to AI inference where flexibility is key.

In our view, custom silicon will be increasingly valuable to hyperscalers, particularly for large, repetitive workloads that justify the cost of specialisation. We expect hyperscaler spending to be balanced between internally developed ASICs and external 'merchant silicon' like Nvidia's GPUs, which remain indispensable for general-purpose compute and inference applications.

Physical AI and synthetic data

At this early stage of AI, it is important not to let a short-term market focus on news flow, company-specific dynamics or competitive threats lead to overlooking the more fundamental, much broader and more important explosive growth in demand for high-performance compute. Across industries, new AI use cases are emerging rapidly.

'Physical AI,' like robotics, humanoid and autonomous systems, are a significant emerging source of AI-driven demand. Amazon, for example, recently surpassed one million robots deployed across its operations, marking a major milestone more than a decade after it acquired Kiva Systems in 2012.

When it comes to transformative technologies, Amara's Law applies: we overestimate the impact in the short term but underestimate it in the long run. After years of development, we are now approaching a tipping point where AI is beginning to influence physical systems, specifically around perception, motion and autonomous decision-making. We referred to Amara's law in our Insight in April 2019 about driverless cars. (Driverless cars: Where are they heading?)

Source: Giles Tulloch

During our trip to the US in June, we had an opportunity to visit Pony.ai, a leading Chinese autonomous driving company. We took a test ride in one of its vehicles still with a human safety driver, highlighting both the complexity of full autonomy and the massive compute resources required. This contrasted with our rides in Waymo vehicles in Los Angeles and Phoenix over the past year, which operated fully autonomously and were already integrated into Uber's platform. Waymo is owned by Alphabet and its first-mover advantage is beginning to show tangible network effects. It is taking market share in the San Francisco taxi market from cars with human drivers and we fully expect it will continue to do so as it expands across new markets.

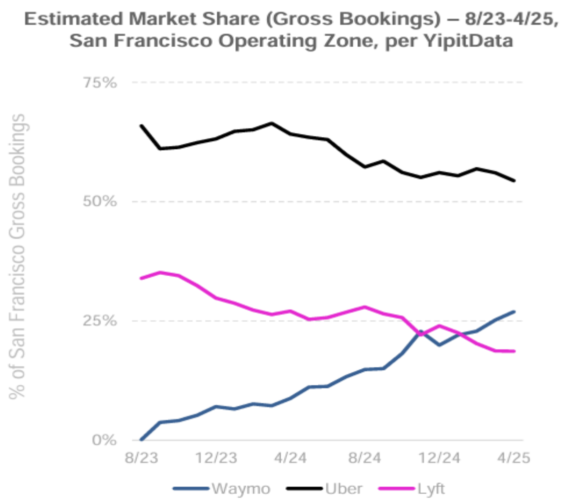

Source: Bond Cap

One particularly interesting takeaway from our meeting with Pony.ai was its use of synthetic data to train its autonomous driving models. Rather than relying solely on real-world driving data, which is time-consuming and expensive to collect at scale, Pony.ai generates its own simulated data. This allows the company to model a wider range of driving scenarios, including rare or 'tail' events, far more efficiently. Rather than waiting for thousands of hours for a real-life event like a sudden obstacle or a herd of cows in the road, the simulation can create the problem and the model can learn from it.

This approach highlights an important consequence of developing autonomous systems and AI more broadly. The need to create representative and diverse datasets adds a significant layer of complexity to the development process. The scale, quality and structure of data can now define the upper bound of AI model performance. Increasingly, synthetic data is not just helpful, but essential, particularly when real-world data is scarce, fragmented, or sensitive. It also means that we have to differentiate between the different values of real-world data. Specific and proprietary data will be of great value for example for analysis, efficiency and productivity in research, healthcare, finance or public services but generic data like how cars interact with roads, other cars and people will be of limited value so implementation, scale and network effects will be of greater importance.

Constraints and opportunities

This shift toward synthetic data has parallels with Meta's experience following the Apple IDFA changes in 2022 and is evidence of how adversity can accelerate innovation. The IDFA changes disrupted Meta's access to user-level data and limited the ability of its advertisers to target potential customers based on their mobile phone journey. The changes were a blow to Meta's business and the share price reacted negatively. They forced the company to rebuild its data infrastructure, essentially creating a new signal from alternative sources and synthetic data. As we highlighted at the time, the constraint has become a strategic advantage. Today, Meta is one of the leaders in AI, allowing it to achieve tangible improvements in both user engagement and monetisation. AI-driven ad tools have significantly improved targeting and conversion, delivering higher ROIs for advertisers and growth in advertising revenues.

Meta's realization of the importance of data is reflected in its recent announcement of a US$14bn strategic investment in ScaleAI and the hiring of its founder Alexandr Wang to lead a new 'super intelligence' group at Meta. ScaleAI specialises in building high-quality training datasets for AI models. Data infrastructure is now as important as model architecture in shaping AI outcomes and as foundational pillars in the development and deployment of AI systems.

During our trip, we also met with Salesforce in San Francisco. The company said that enterprise adoption of AI remains slow, primarily constrained by internal data complexity. Unlike internet-scale consumer LLMs trained on publicly available and well-labelled data (as with ScaleAI), proprietary enterprise data is often unstructured, poorly labelled and siloed within organisations. This makes training effective AI models far more challenging in corporate settings.

Salesforce emphasised that enterprise AI must meet a much higher bar for accuracy than consumer chatbots. In consumer settings, errors are tolerable and may even be expected. Yet, in an enterprise context, mistakes can have serious financial, legal, or reputational consequences. This higher burden of proof demands both superior model precision and better data integrity. The opportunity for enterprise data is enormous and the data challenge is why Salesforce has announced its intention to acquire Informatica for US$8bn. It may foreshadow broader M&A trends across the AI value chain. As AI adoption expands, we expect growing demand for data infrastructure, governance and management tools, especially those capable of preparing enterprise data for AI systems at scale.

In summary, our trip reinforced our conviction that Nvidia remains the clear leader in AI. It has a deep and defensible competitive position built on technology, ecosystem and execution. While we acknowledge the growing role of custom silicon developed by hyperscalers, we believe the demand for advanced compute is set to expand so significantly that both internal and merchant solutions will find ample room to grow. AI use cases are expanding and include the emergence of physical AI**,** where AI moves beyond the digital realm into robotics and autonomous systems. Across all these domains, data remains the gating factor. The ability to access, structure and augment high-quality datasets will determine the pace of enterprise and industrial AI adoption. As companies overcome these data challenges, we expect the next wave of AI deployment to accelerate meaningfully.

Giles Tulloch
July 2025