

AI – WHERE WILL MOST BUSINESS VALUE ACCRUE?

CARINE SMITH IHENACHO
Chief Governance Officer,
Norges Bank Investment Management

November 14, 2023





Good morning!

How many of you remember the Y2K problem? You must be born in the last century to do that, so it does not go for all of us, I guess, but I remember it. The Y2K problem was the fear of computer errors as we turned from 31st December 1999 to 1st January 2000. Could the whole infrastructure for computer reliant industries be brought down on New Years eve? Just think of it, a digital doomsday!

The answer was: no. Hardly anything happened, despite of warnings and doomsday predictions.

It was not the first time anyone said or meant something that in hindsight would turn out to be incorrect.

In 1977, Ken Olsen, the founder of Digital Equipment Corporation, that later was going to be acquired by Compaq, said: “There is no reason anyone would want a computer in their home”.

Or the former Chief Technology Officer in Microsoft, Nathan Myhrvold, who in 1997 said: “Apple is already dead”.

Of all “sciences”, predictions are the most difficult.

History has a lot to teach us. That is why it is so fascinating. Three months into 2000, we reached the peak of the dot-com boom. What was then the status of three of the big long-term winners from the internet – Meta, Google and Amazon?

Meta - and social media in general - did not even exist.

Google did exist - but was just an 18-month-old start-up.

Amazon was listed - but was a simple retail business. The cloud business and extensive logistics that drive much of Amazon's current value, was non-existent.

In addition, some of the stocks expected to be the winners in the dot-com boom, turned out not to be such good investments. Cisco, the "picks-and-shovels" of the internet, has not been able to recapture the value it had at the peak of the boom. This is in sharp contrast to Amazon, for example, that has far exceeded that level.

Why am I telling you this? This was supposed to be about AI, right?



Well, if we are to learn anything from the history, the AI investment landscape will change dramatically over the next twenty years. Many winners might look very different from their current form. Some might not yet exist at all. And some companies that now look like winners, could end up much less successful than many thought.



As investors, we should acknowledge this uncertainty and not overestimate our predictive abilities.

When making this presentation, I wanted to find out what ChatGPT could do for me. So, I asked: "what is the future of AI"? The answer was pretty good: "The investment outlook for AI in 2024 is incredibly promising". Hard to disagree.



But I am not going to give you the ChatGPT-answer. Instead, I will give you our thoughts and deliberations on the investment outlook for AI.

I work at the Norwegian Sovereign Wealth Fund, a global, long-term investor, with around USD\$ 1.5 trillion under management. We own on average 1.5% of the world's listed companies. We have more than USD\$ 100 billion invested in AI related business, although that's not a clearly defined term.

When we talk about AI related investments, this covers our investments in companies that today are the main drivers of the AI development, through their involvement in semiconductors,

software industries, data services and data centres. Out of this, our largest investments are in semiconductor and software companies.

And how do we see the investment outlook in these sectors?

“ The short answer is that the investment outlook involves significant uncertainties. We are at the beginning of the AI journey. The only thing we know is that events will be unpredictable. ”

But how do we best manoeuvre in this uncertain landscape?

We have identified three principal sources of uncertainty that we believe will shape the AI investment outlook. As investors, we should follow the developments of these uncertainties and plan for various outcomes.

“

The first area of uncertainty is the ownership both of foundational models and software infrastructure.

”

Will they be proprietary – so that only the companies that develop them can adapt and modify them – or will they be open source?

For foundational models, the question of proprietary versus open source, can be illustrated by OpenAI’s ChatGPT versus Meta’s LLaMa.

ChatGPT is free for anyone to use, but the actual model is proprietary to OpenAI and cannot be modified to fit other software systems. Meta’s LLaMa, on the other hand, is free for anyone to use, modify and adapt.

And similar for software infrastructure, there is Nvidia’s CUDA, that can only be used in connection with a Nvidia GPU, versus OpenAI’s Triton that can fit other systems.

We believe this issue of ownership of foundational models and software infrastructure is the area that will have the most impact on companies’ long-term value creation.

“

The second area of uncertainty is the competitive landscape of silicon infrastructure.

”

Today, Nvidia is pretty much the only game in town. But will it continue to be a concentration in one single provider? Or will it be a competition among many? Will the long-term winner be Nvidia’s GPUs, or will there be competing GPUs from other providers, such as AMD or Intel? Or will it be Nvidia’s GPUs versus custom AI processors, like Alphabet’s TPU or Amazon’s Trainium and Inferentia?



The third area of uncertainty is the ownership of intellectual property and data. AI generates content with higher quality and speed of creation compared to average human performance.




But to do so, the models rely on huge training data sets by mining publicly available content.

This use and output of data creates many issues around intellectual property and ownership that are not yet resolved. Will using copyrighted data to train generative AI models be covered by the “fair use” doctrine? Can you copyright the output of a generative AI model? Who would own it? What remedies can be introduced to allow generative AI to flourish, while giving credit or compensation to the creators? These questions are currently being discussed and disputed.

With all these uncertainties, it is useful to assess these six scenarios.

- If proprietary software infrastructure and foundational models dominate, the owners of these models are likely to attract an outsized share of the value. Big Tech will give way to Big AI. Big AI may comprise the current big tech companies, or new may appear.
- If open-source software and models dominate, the value will be more evenly distributed across the economy among users.
- If a single silicon platform dominates, this will earn very high returns on investment. Nvidia’s enormous valuation increase may then just continue.
- If multiple silicon platforms compete evenly, value will be accrued more with the model owners or the end users. AMD and Intel may capture some of Nvidia’s value creation, or new companies may emerge.
- If we get stringent copyright laws, this will limit the training data sets. This will again benefit content owners, who can license and create walled gardens.
- But with a more relaxed copyright regime, companies with platforms relying on volume for user engagement, such as Meta and Google, may be the winners.



And this is probably the best investment advice I can give – at least much more concrete than I got from ChatGPT: Follow closely the developments of the following areas and adapt your portfolio to what you believe are the most likely outcomes:

- The adoption and success of open-source software and models.
- The extent to which cloud players will invest in silicon to compete with Nvidia.
- The extent to which specialist processors for AI applications can compete with more general-purpose GPUs.
- The development of the various generative AI regulatory schemes

under consideration across the world to determine how the copyright regime will develop.

These are the developments we watch. As of today, this is how we approach investments in AI:

- We take a portfolio approach.
- We invest in companies that can win in multiple scenarios.
- We invest in multiple companies that can win in discrete scenarios.
- And we have reduced exposure to companies with highly vulnerable business models.

I have now looked at some of the investment opportunities in companies that will be deeply involved in the development of AI. In addition, AI can be a powerful tool to enhance business models, increase efficiencies and reduce cost across a variety of industries. As an investor in more than 9,000 companies, we urge these companies to utilise the great potential of AI.

We will probably first and foremost see this in companies that depend on many manual routine processes based on knowledge or data intensive work. Chatbox can be used for customer services. AI can draft legal and technical documents. What a relief for paralegals and junior lawyers who no longer will have to review thousands of pages of legal precedence. Software engineers may be replaced by AI coding. HR processes, such as payroll, may be totally automated. IBM reported that they have reduced their HR admin cost by 90% due to AI. Even the creative industries, such as advertising agencies and stand-up comedian, may benefit from a bit of AI prompt.



The same applies to us. As a fund, we have an ambition of 10% efficiency gain the next year by using AI. We do not think it can replace investment professionals, but it can make us trade more efficiently, automate many processes, and facilitate our risk monitoring and reporting.



But some companies' gain may be other companies' loss. Companies with revenues predominately generated by services that can be substituted by AI, may find their business model challenged. This could include, among many, payroll processing companies, call centres, IT consultants, ride-hailing companies and companies providing project management services. Watch out for that.



But AI will have much more impact on people and society than companies' revenues and costs. AI also comes with severe and uncharted risk. AI now develops at a pace where it is challenging to predict and manage its impact.



Beyond regulatory, operational, and reputational risks to companies, development and use of AI systems can impact society at large and even human rights. Issues such as privacy, security, personal freedom, and non-discrimination are at stake. AI will also increase the risk of large-scale misinformation, deception, or manipulation.

As a long-term, diversified financial investor, we believe that our investments will benefit from the development of comprehensive and cohesive regulation of AI. This will contribute to safe innovation, market predictability and create a basis for long-term value creation.

But while AI legislation is developing, companies should view responsible development and use of AI as a core element of responsible business conduct. Companies' own responsible use of AI should be a necessary complement to the emerging regulatory landscape.

We have developed and published clear views on what we believe are key elements of responsible AI. This includes three areas:

- The first is board accountability. Boards should ensure that the company balance competitive use of AI against potential risks – including risks to people and wider society. This will require board expertise. Boards also play an important role in overseeing that AI policies and guidance are implemented across the business.
- The second is transparency and explainability. Companies should be able to explain how the AI systems they develop, or use, have been designed, trained, and tested. Stakeholders should be able to assess the potential impacts of AI and understand its accuracy, efficiency, and reliability. Companies should also provide information to relevant and trusted third parties, such as auditors or regulators, to

allow them to verify the AI system and assess its risks. Companies should inform people of how AI is used in outcomes that affect them and provide appropriate access to remedy. We're very far from that today.

- The third is robust risk management. Risk management processes should identify, assess, and mitigate risks to business, people and society. Risk measures should include evaluations of system limitations and the potential consequences of system failure. AI systems that can pose severe risks to societal and human rights, should be subject to additional controls.

I mentioned earlier the developments to watch on the investment side. Let me also part you with some developments to watch when it comes to responsible use of AI.

- The regulatory landscape is evolving quickly, so keep abreast of developments of international standards and regulation.
- Also, best practices are also evolving where we do not yet have regulations. Be informed of the latest developments.
- And as AI develops, be updated on the many risks involved, in

particular risks related to privacy, security, misinformation, deception and manipulation.

One of the most famous quotes ever made in history was the former president in IBM, who in 1943 reportedly stated: "I think there is a world market for maybe five computers".

When he said this, computers were pretty much as large as a house, and nothing close to what we got today. Watson had no idea that he was at the very dawn of the computer industry. And I can't blame him.

“ It is hard to prepare for something you cannot even imagine. But we must try. That is how we move forward. ”

Thank you!
