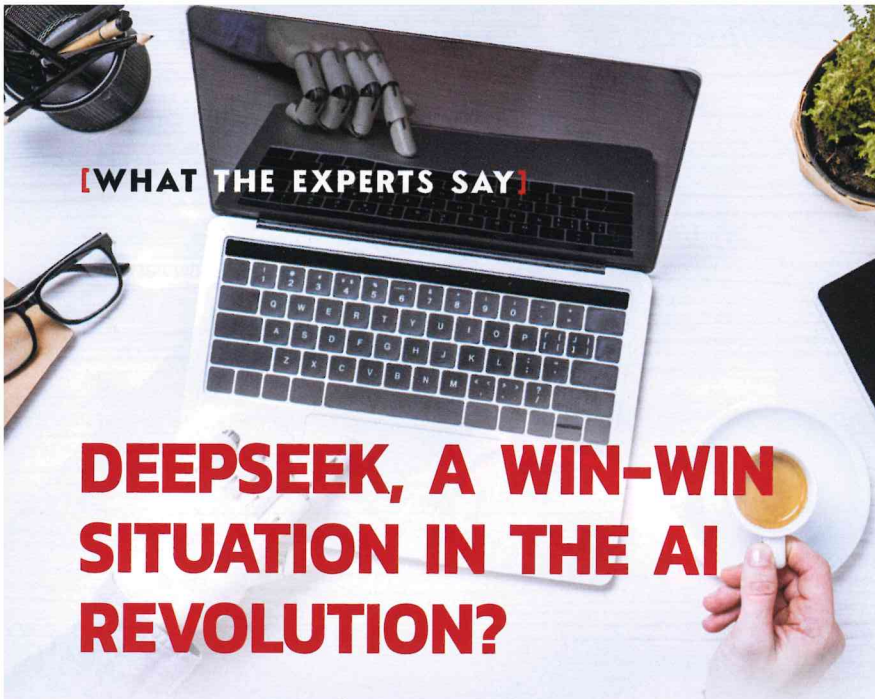

DEEPSEEK, A WIN-WIN SITUATION IN THE AI REVOLUTION?

Benoit Dabertrand
Group CEO and Investment Committee Member
Fisconsult Fund Management Ltd

Seilendra Gokhool
Partner - CEO
Sinews Global Ltd
Director, Fisconsult Fund Management Ltd

A few weeks ago, there came a sudden realization that a Chinese company, DeepSeek, had come out with a few good, efficient AI models at (supposedly) a fraction of the cost of the ones developed by their US competitors, in particular DeepSeek R1.

Since DeepSeek, being Chinese, has less to no access to the latest GPU chips because of US export controls, that came out as quite a shock to investors and American AI companies such as OpenAI. As the saying goes, necessity is the mother of invention and having to make do with less latest generation semiconductors pushed DeepSeek to be more creative in the creation of its AI model.

All of this led to a relatively sharp selloff in everything AI-related, be they Nvidia, Taiwan Semi, Schneider Electric, Vertiv and many others, although names like Meta and Apple were up. The thinking for the latter was that they would have to spend less on AI than what was originally planned, although the latest announcements show that the big companies, Apple included, still plan on spending hundreds of billions of dollars on AI in the coming months and years.

When such a shake-up in the markets occurs, one should always let the dust settle before concluding anything with certainty. Is all this an overreaction or the beginning of the end for the AI trade? We shall see, but we would argue that this was an overreaction. This is because of the new buzzword of the day: Jevons paradox.

The Jevons paradox was first described by the English economist William Stanley Jevons in his 1865 book The Coal Question. Jevons observed that England's consumption of coal soared after James Watt introduced the Watt steam engine, which greatly improved the efficiency of the coal-fired steam engine. Watt's innovations made coal a more cost-effective power source, leading to the increased use of steam engines in a wide range of industries.

This in turn increased total coal consumption, even as the amount of coal required for any particular application fell. Jevons argued that improvements in fuel efficiency tend to increase (rather than decrease) fuel use, writing: "It is a confusion of ideas to suppose that the economical use of fuel is equivalent to diminished consumption. The very contrary is the truth."

This idea, applied to the first industrial revolution, and which turned out to be entirely true, is probably applicable to our current, AI, industrial revolution. As AI tools become less energy-intensive, demand for them will only increase, leading to steady demand for energy, semiconductors and all the other tools needed to power the AI revolution forward. In fact, this is what we are seeing as companies are now moving from Large Language Models to Reasoning Models and AI Agents. Training conventional large language models, the types you see in free versions of most AI chatbots, requires vast amounts of power and computing time. But companies are rapidly figuring out ways to reduce the amount of resources they need to run when a human calls on them. Reasoning models, which are based on large language models, are different in that their actual operation consumes many times more resources, in terms of both microchips and electricity.

And the thing is, owing to their enhanced capabilities, these reasoning systems will likely soon become the default way that people use AI for many tasks. AI reasoning models can use more than 100 times as much computing resources as conventional large language models, Nvidia's vice president of product management for AI, Kari Briski, wrote in a recent blog post. That multiplier comes from reasoning models spending minutes or even hours talking to themselves in a long "chain of thought." The amount of computing resources used by a model is proportional to the number of words generated, so a reasoning model that generates 100 times as many words to answer a question will use that much more electricity and other resources.

In January, it appeared that the cost per token, in both computing power and dollars, would crash in the wake of the release of DeepSeek R1, the Chinese AI model. DeepSeek, with its accompanying paper, showed it was possible to both train and deliver AI in a way that was radically more efficient than the approaches previously disclosed by American AI labs.

On its face, this would seem to indicate that AI's future demand for computing power would be some fraction of its current amount, probably a tenth, or even less. But the increase in demand from reasoning models when they are answering queries could more than make up for that. To look at in the most simplistic way, if new, more efficient AI models based on the insights that went into DeepSeek slash demand for computing power for AI by a tenth, but reasoning models become the standard and increase demand for those models by a factor of 100, that's still a 10-fold increase in future demand for power for AI.

This is why we believe that the news out of China are probably a net positive for the whole AI trade. The only clear conclusion is that China, far from lagging, is just as good as the US on the AI front, whatever export restrictions were put in place. As a matter of fact, when looking at Nvidia financials, one notices that 20% of revenues come from Singapore. That Singapore would need that many semiconductors seems highly improbable, and one can conclude that some companies in Singapore act mainly as fronts for Chinese companies that then re-export those chips to mainland China.