

Insights from the US

Source: Adobe Stock

I attended the Goldman Sachs 'Communacopia' technology conference in San Francisco in September with panel sessions from, among others, OpenAI, Nvidia, SpaceX and Google Cloud.



Emma Fisher: Obligatory Jensen Huang fangirl photo

The key takeaways coming out of the 40 sessions I attended over three days were:

1. Token-cost economics

If 2025 was the year of "token-maxxing", with corporates aiming to drive LLM usage across companies and get employees using the tools for anything and everything, then 2026 has been about token-cost optimisation. There was much discussion of using the right model for the right task; not every task requires frontier intelligence, and there was a broad consensus that the rise of (significantly) cheaper open-source, open-weight models is a good thing for the AI ecosystem and broader adoption.

The Pinterest CEO was blunt: "Any CEO who is not heavily leveraging open-source models is almost certainly wasting shareholder money...we are able to get cost at less than 8% of the comparable frontier model by using open source."

2. Cybersecurity: the agents are here

Cybersecurity is clearly the next "killer app" for the frontier labs after coding: every company in the world will need the greatest defender tool set, and this will drive significant agentic workflow. While computing demand to date has been driven primarily by consumer queries, which are relatively low token intensity, agentic workflows such as cybersecurity run 24/7. Jensen Huang called this out: "Cybersecurity will be the next big market [with] agents running all the time." This links to a broader theme: agents are here. Meta released its personal agent 'Muse' while I was in the US, offering a consumer-facing agentic application that could further compound token demand, as agentic workflows run 24/7.

3. The models will likely keep improving as scaling laws work across multiple variables

"Moore's Law is coming to the end. Moving down the process node cannot drive double the performance. We are trying to substitute Moore's Law with something else. Moore's Law on one semi chip is limited by the radical law of physics on optics, so you can only cram so many transistors into one. Now you can put two die together into one chip, gets you double the performance. Current generation has four die, the next will have eight die...this will drive exponentially better LLMs...driving more and more value out of superintelligence. We do not see a limit at all yet." CFO of Broadcom.

Jevons paradox: what does it mean for our portfolio?

The question that is fascinating global markets right now is whether or not we're in an AI bubble. The case for a bubble is not hard to make. The scale of the AI build-out is historic- the largest US tech companies are spending c2% of US GDP on capex this year, a comparable level to the building of America's railroads in the 1850s and approaching the peak of the 1990s internet boom. More importantly, funding has changed. Until recently, the build out was paid for out of the hyperscalers' enormous free cash flow. It is now increasingly funded by debt, and the credit market has noticed: default protection on heavy borrowers such as Oracle and CoreWeave has widened sharply.

The case against is that the demand is real, and it far exceeds compute supply today. Wherever we look, compute is short: customers are committing, in some cases prepaying, to secure access to energised GPUs; capacity is being sold before it is built, and the revenue that can be earned from a megawatt of AI capacity is rising: SemiAnalysis reports that one-year rental contracts has risen sharply over the past year: SemiAnalysis reports that one-year rental contracts for Nvidia's H100 rose almost 65% between October 2025 and August 2026. This is a classic price signal that the market needs more supply. On the usage side, the cost of each token is falling fast- Epoch AI estimates the price to achieve GPT-4's performance on a set of PhD-level science questions fell by 40x per year. Meanwhile, total token demand, driven by more tasks and more users, is increasing- a classic "Jevons paradox"- a term used to describe what often happens with new technology: it gets cheaper and more efficient, so people end up using far more of it. The phenomenon is named for William Stanley Jevons, who predicted in 1865 that more efficient steam engines would increase rather than decrease coal consumption in British factories.

Overall, from a portfolio perspective, we watch the debate rage and are more interested in the question of what the diffusion of this extraordinary technology will mean for company moats across the world. Everything is changing, and at an accelerated pace. The two big themes of the conference; that is, token cost optimisation, and the rise of open-source, open-weight models, are related themes that in our view broaden the health of the AI ecosystem and will drive adoption. Knowledge is becoming a commodity, and change is coming for every company.

The models are getting better and cheaper

The first clear area of change will be in the operating costs of enterprises. While many companies are still struggling (or reluctant to) publicly outline clear ROI on AI spend, in our one-on-one reporting season catch ups with companies we noted a marked shift between February and August. While most companies in February pointed to the ability to accelerate product roadmaps with existing resources (i.e., use AI to "do more" with the same resources), by August more than half were conceding they could likely do more with fewer resources over the coming years. As the models get better and token costs fall, the cost to achieve an outcome falls dramatically, driving ROI on the spend.

Who banks the savings?

We believe it is reasonable to assume that AI will lead to greater productivity for most companies. Some companies will choose to bank the savings; others will reinvest proceeds into revenue-

generating activities. However, this does not logically lead to higher margins for all. The extent to which shareholders enjoy banked savings and hence higher margins will depend on the pricing power of the business, as it always has with any technological advancement.

Taking banking as an example: in the 1970s, banking meant visiting a branch during limited hours, queuing for a teller who wrote deposits into a passbook, and waiting days for cheques to be sorted by hand. Since then, ATMs, core banking systems, the internet and mobile banking have cut the cost of a transaction to a fraction of what it was. Yet banks' returns haven't structurally risen, because every bank adopted the same technology at roughly the same time, so savings were passed onto consumers via lower fees, as well as paying for the investment in technology. More companies now "clip the ticket" between a customer and the bank. A clip is paid to cloud service providers, payment rails (Visa and Mastercard), Apple Pay, core banking vendors like Oracle and Fiserv, SAP, per-seat SaaS fees to companies like Microsoft, cybersecurity vendors and credit data providers. The same dynamic is likely to be true in an AI world.

Companies with unique pricing power are likely to hold onto more of the gains of AI. Businesses like James Hardie, with its near 100% share of the fibre cement siding category, The Lottery Corporation, which has exclusive monopoly licences to run state lotteries in Australia, and Aristocrat, which can save on game design and development spend while charging casino floors the same clip of the ticket, have an attractive combination of unique pricing power and (likely) undisrupted business models that mean they should be able to hang onto any productivity gains from AI.

Rapid change in moats

If knowledge is a commodity, there is no longer any moat in how hard it is to produce that knowledge. Coding was the first advantage to get wiped out: the cost of producing code has collapsed and legacy companies that find themselves with messy data and legacy 'tech debt' can have a disadvantage vs AI-native companies in the new world. The question shifts to what becomes scarce in an AI world. From an intangible perspective, companies that own proprietary data and own the customer relationship should be able to use AI to produce better revenue-generating products. Car Group (CAR) is an example of a business in the portfolio that is almost as much a critical B2B business in car dealer supply chains as it is a meeting place for buyers and sellers.

The biggest disruption risk for CAR isn't a competitor building a classified site, it's the consumer interface moving to AI agents. If agent search proliferates, the Carsales 'meeting place' will lose some of its pricing power. We believe the best hedge for this risk is for Car Group to move deeper into dealers' systems, thereby ensuring it stays essential, and monetises its proprietary data, even if parts of search shift to agents (something we see no evidence of to date). We believe management recognise this: at the August result, the company announced the launch of Nexgate, an AI-powered dealer platform that connects inventory, leads and workflow.

The portals have the advantage of the fact that they are two-sided marketplaces, and they are free for buyers to use. This means that you cannot 'dislodge' a portal by offering a cheaper product.

On the seller side, you can't dislodge a portal by offering cheaper


listings, as the sellers go where the buyers are. By contrast, we are concerned that Xero's pricing power may erode over the coming years. Xero users are paying customers who see the bill every month. The bear case for Xero is that increasingly capable AI agents could automate many of the bookkeeping, reconciliation and administrative tasks that have traditionally underpinned accounting software. A big chunk of Xero's current competitive advantage is the high switching costs of the user experience. Busy small business owners simply cannot be bothered learning a whole new accounting software interface. The more workflow that becomes agentic, the lower the switching costs of the product. While we continue to track how this space may evolve, we have exited Xero in the portfolio as we consider some of these long-term risks.

The return of the physical economy

Sometimes, when considering change, it's also worth inverting the question and considering what is unlikely to change. The physical world of mines, commodities, physical supply chains, energy, land and power remains unchanged by artificial intelligence, and in fact the AI expansion is highly dependent on pulling on stuff in the physical world.


It takes around W15 years on average to bring a mine online, which means the commodities available today to meet demand are a function of the projections made about demand 15 years ago. It would have been difficult 15 years ago to accurately predict the electrification of the energy grid, the reshoring renaissance of industrial capacity in the US (incentivised by tariffs), and the mother of all AI-capex builds. Additionally, these demand drivers have arrived at a time when physical assets have languished under years of low commodity prices, and a long period where shareholders have demanded more disciplined capital allocation from management teams.

The mining explosives market in Australia is a perfect case in point. Ammonium nitrate (AN) is a controlled chemical made from gas-backed ammonia, and plants take years to gain permits and to build. The 2010s saw a wave of supply added by Wesfarmers, Dyno Nobel and Orica. Supply has been flat since 2017, yet the addition of new greenfield capacity led to declining industry returns from 2013 onwards and it took a decade for demand to grow into current supply.



Source: Company data, Airlie Research

We believe the market has now moved from balance into tightness, driving a pickup in EBIT/t that we expect to continue as no new supply is planned to come online, while elevated global pricing (due to higher global gas prices) plus shipping costs mean the economics of importing AN to Australia are unattractive unless pricing lifts.



Source: Company data, Airlie Research

This restricted supply dynamic has played out across many commodities in a similar fashion-portfolio beneficiaries of this dynamic include BHP, Rio, Orica, Dyno Nobel, Santos, and Ampol.

Emma Fisher

Deputy Head of Australian Equities

Important Information: This material has been delivered to you by Magellan Asset Management Limited ABN 31 120 593 946 AFS Licence No. 304 301 trading as Airlie Funds Management ('Airlie') and has been prepared for general information purposes only and must not be construed as investment advice or as an investment recommendation. This material does not take into account your investment objectives, financial situation or particular needs. This material does not constitute an offer or inducement to engage in an investment activity nor does it form part of any offer documentation, offer or invitation to purchase, sell or subscribe for interests in any type of investment product or service. You should obtain and consider the relevant Product Disclosure Statement ('PDS') and Target Market Determination ('TMD') and consider obtaining professional investment advice tailored to your specific circumstances before making a decision about whether to acquire, or continue to hold, the relevant financial product. A copy of the relevant PDS and TMD relating to an Airlie financial product or service may be obtained by calling +61 2 9235 4760 or by visiting airlifundsmangement.com.au.

Past performance is not necessarily indicative of future results and no person guarantees the future performance of any financial product or service, the amount or timing of any return from it, that asset allocations will be met, that it will be able to implement its investment strategy or that its investment objectives will be achieved. This material may contain 'forward-looking statements'. Actual events or results or the actual performance of an Airlie financial product or service may differ materially from those reflected or contemplated in such forward-looking statements.

This material may include data, research and other information from third party sources. Airlie makes no guarantee that such information is accurate, complete or timely and does not provide any warranties regarding results obtained from its use. This information is subject to change at any time and no person has any responsibility to update any of the information provided in this material. Statements contained in this material that are not historical facts are based on current expectations, estimates, projections, opinions and beliefs of Airlie or the third party responsible for making those statements (as relevant). Such statements involve known and unknown risks, uncertainties and other factors, and undue reliance should not be placed thereon. No representation or warranty is made with respect to the accuracy or completeness of any of the information contained in this material. Airlie will not be responsible or liable for any losses arising from your use or reliance upon any part of the information contained in this material.

Any third-party trademarks contained herein are the property of their respective owners and Airlie claims no ownership in, nor any affiliation with, such trademarks. Any third-party trademarks that appear in this material are used for information purposes and only to identify the company names or brands of their respective owners. No affiliation, sponsorship or endorsement should be inferred from the use of these trademarks. This material and the information contained within it may not be reproduced, or disclosed, in whole or in part, without the prior written consent of Airlie. (080825-#A10) UID500