



[00:00] [Rob Campbell] [RC] Wen Quan Cheong is back on the podcast with fresh insights from Taiwan and South Korea on some of the concerns and opportunities in the AI build-out: how safety considerations might reshape the outlook for compute demand, which supply chain players have the strongest structural tailwinds, and, outside of AI, why a port operator in the Philippines might be one of his most exciting ideas right now.

[00:26] [Disclaimer] This podcast is for informational purposes only. Information relating to investment approaches or individual investments should not be construed as advice or endorsement. Any views expressed in this podcast are based upon the information available at the time and are subject to change.

[00:43] [RC] Wen, welcome back.

[00:44] [Wen Quan Cheong] [WC] Thanks for having me again. Always a pleasure to have this exciting conversation.

[00:52] [RC] Maybe hold that thought, because you might not like my first question. I want to ask you about some of the bigger headlines that we've seen more recently and how they connect to a good chunk of the emerging markets universe. A lot of them have to do with concerns about AI. We had the Hugging Face incident where a bunch of AI agents went rogue. Most recently, we had a researcher at Anthropic quit quite publicly amid some ethical concerns that are quite hyperbolic in terms of human civilization.

We're gearing up for midterms in the U.S. and clearly AI has been a big topic there. I'm wondering, in your research, how do you view these concerns?

[01:30] [WC] Maybe I'll start off by sharing a bit about what's happening now and some of the implications, and what that could mean for chip demand and semiconductor stocks, which, as you rightly said, is a big part of the emerging markets universe. I think there are several reasons behind the push for greater oversight and a slower pace of AI development. Safety is the clearest explanation, but there are talks about competitive pressures and commercial considerations that also matter, and a ton of other conspiracies.

The first one is on alignment and safety, basically to catch up with AI's capabilities. The big topic here is alignment, and alignment is making sure AI pursues the goals people intend, stays within acceptable boundaries, and responds appropriately to human direction. That becomes especially important when AI can use tools or take actions by itself. Keeping people in control also requires monitoring, limits on access, and reliable ways to intervene.

For example, as you mentioned, [Dario Amodei's recent essay](#) highlighted a July incident where AI systems were being tested by OpenAI and they coordinated an unauthorized attack on Hugging Face, which is a platform used by AI developers. The independent investigators found roughly 700 AI agents participated together to achieve the



goal. There were very interesting conversations that happened if you follow up with the reading, and Anthropic themselves also reported similar situations that happened during test evaluations. Dario's concern here is that similar behaviour from more powerful systems could cause much greater damage.

Interestingly, I also by chance came across an article providing somewhat of an inside view that I thought was interesting because it was written by one of the lab researchers. In that article, it claimed that the lab researchers were not suddenly frightened because of some public model that was released. They were frightened because of their internal results, suggesting that several largely untapped scaling methods could combine to produce much more capable models very quickly, especially in dangerous areas such as cyber operations, before reliable control systems exist. So I think they see frontier pacing as a genuine attempt to buy time for alignment and security rather than an attempt by leading labs to lock out competitors.

This entire thing about pacing, and what it means for compute demand, is interesting as well. On the alignment front, the chief scientist at OpenAI, Jakub Pachocki, recently mentioned a couple of days back that ensuring AI safety and alignment requires significantly more spending on compute as models grow more and more capable.

[04:37] [RC] I would imagine, at least when looking at many of the companies that you cover in the emerging markets space, a lot of them are providing the ingredients that go into that compute and enable it. One version of considering this slowdown would be, hey, maybe there's less demand for that compute. If I heard you right, though, you're saying that might be possible, but on the other hand, if there is greater investment in alignment, that's going to require even more compute.

[05:03] [WC] I would break compute demand into three different buckets now. The first one is model training, then you have this alignment bucket, and then there's an inference bucket. Of course, you develop models, and that's the training aspect that might slow with pacing. I'm not sure how much will be offset by the increased compute demand generated from more alignment and safety.

An interesting stat is that scuttlebutt with one of the frontier labs showed that two, three years ago, 80% to 85% of compute actually went to training and the remaining went to inference, and today that number is 50-50. On the model training side, Dario also mentioned that pacing does not mean stopping model training or technical progress. In fact, as I mentioned earlier, the AI safety and alignment portion might increase compute and may make up for the pacing of model development.

Judging from that, even if the next model arrives later, compute demand will then depend primarily on inference demand, which is observed to be skyrocketing, especially in a world of agentic AI. There are tons of token stats out there that prove that. Another concern that people might have, or the broader concern, is that slower improvement in models could delay useful applications and weaken future demand. But for now, the existing models are really generating insane demand on the inference side. That's one to keep watch on, one to monitor as well.

For semiconductor stocks, my concern would initially be lower growth expectations and lower valuations, but share prices can fall even while revenues keep growing, if investors had expected way faster growth and that didn't turn out as expected or baked into expectations. So I would watch the actual infrastructure spending, the chip orders, whether there are delays, utilization, customer usage, and this kind of stuff. But for now, I don't see



that as a major concern.

In terms of portfolio construction, as TSMC always says, it's important to get the long-term structural direction right. It's hard to predict the near term, but the longer-term structural direction is that compute demand is getting higher and higher, and with the longer-term direction of models towards more recursive self-improvement capabilities, new models should arrive materially quicker. So over the longer term, more compute demand on the model training side. Over the long term, there's demand generated for AI safety and alignment, and then you have the inference bucket, which is increasing at a very rapid rate today.

[07:58] [RC] Even before some of these concerns got more attention, the question of demand, how much, and how sustainable for how long, that's been a major topic. I know we've talked about it. Can you share some of the things that you do to help ascertain what that looks like, which then informs your assumptions around value and trajectories? What does the day-to-day look like for you guys on the ground working on our emerging markets team?

[08:23] [WC] AI demand, in my opinion, is a very blue ocean market. It's really hard to know. As I mentioned just now, getting the general direction right is way more important, in my opinion.

There are a couple of stats that support the argument that we are largely still very under-penetrated today. There were stats published by Ramp or Andreessen Horowitz on the spending per employee per month of the top 1%, top 10%, and the median. All three stats showed that there's still a lot more room to grow.

I think AI infrastructure demand compounds across two things. The first is the number of users, which is penetration. I think we are quite penetrated today. But the second aspect of it is the tokens consumed per user and per agent. That second piece is the one that has been increasing very rapidly and is the driver of compute demand today. With agentic AI and the ability to run many agents at the same time, that consumes a lot more incremental inference compute demand as well. For now, that's the general direction. I have no ability to predict what happens in the near term, but from the evidence that we see, both the near- and long-term trend themes still seem to be playing out.

[09:51] [RC] You've just come back from a research trip in Taiwan and South Korea, where a lot of these companies are based, as well as their supply chains. I'd love if you could share a little bit more about that trip. What prompted it in the first place?

[10:06] [WC] For us, a research trip helps to do two things. In today's world, it helps us to build a better understanding of changes earlier, and it also allows us the opportunity to gather different perspectives to challenge our views. I think that matters a lot more in today's world. Easy sources of alpha, in my opinion, are much harder to find nowadays. Technology makes information available to almost everyone instantaneously. A useful insight can appear very quickly on X and be widely discussed within hours. So alpha from a lot of short-term trading has an even shorter shelf life today.

Against that backdrop, I see several ways we can build our edge, especially as longer-term investors. Gain early understanding. This is especially true for areas that evolve very quickly, like AI, and perhaps even the biotech theme that we have been invested in as well. Josh, I think, will be speaking about that in several weeks.



The aim of this is to get out there, speak to more people, and identify these frontier changes before they become widely understood. Understanding these changes can provide a longer-term direction from which we can position and profit. And not only if we understand earlier, but also gain conviction to be opportunistic when market dislocation occurs. We can't predict the short term, but we hope to study enough to gain early understanding and be directionally correct in the long term, and then use that to our advantage to identify opportunities and invest in them.

The second way we can build our edge is to be more curious, to be more open-minded, ask better questions, explore unfamiliar ideas, and take opposing views from others seriously. That includes the opportunity to scuttlebutt on competitors and get different perspectives from other investors as well. A research trip, in my opinion, fulfills those two things that are more important today as longer-term investors.

[12:04] [RC] In terms of the mechanics of actually planning the thing, I assume if you're going to hop on a plane and go on a research trip, you're going to pack your days with meetings and know exactly what you want to do to really maximize the time. Do you explicitly leave room for more of that curiosity?

[12:26] [WC] In planning research trips, I tend to be curious and open to anything. I would fill up my schedule with portfolio companies that we invest in and their competitors, but also random companies that I see. Some of them don't even look optically attractive on screeners or just by looking at their financials. But I just have them out there to learn something new, and aside from the structured conferences, I also add randomly screened companies just to go out there, explore, and hopefully find something.

You don't always know which conversations, at the end of the day, will turn out to be very valuable. I just rather lead this by curiosity. Interestingly, some of that played out quite well. Let me use AI as an example. Before the AI theme actually became the topic of the day, we spent quite a lot of time in Taiwan getting to know companies, understanding their businesses, and building relationships with management teams. So when this theme of AI was incepted and started accelerating, the initial groundwork built years ago helped us to connect the dots and act faster. We were already familiar with some of the businesses and the people running them, and could quickly get up to speed and either invest ahead of the curve or as things were accelerating. That's something I've come to appreciate about research trips and being curious: the value of the work isn't always obvious when you're doing it.

Another example I can share: one of our larger positions in the portfolio is Acter Co. I met this company in the early part of 2024. I was attending a conference in Taiwan and happened to be arranging additional meetings on the side, and for some reason, I added Acter onto the list. Interestingly, Jeff Liang, who was a special assistant to the chairman and handles the IR function as well, was very generous with his time. I had very little clue, given it's such a niche business, but he patiently spent more than an hour walking me through the business. It's one of a kind in the world. You don't really see this kind of business, so he had to really induct me into the business. It took a while and several meetings after for me to understand it.

Later on, months after the meeting, I think a year after the meeting, as the AI build-out continued, Trump's tariffs and incentives obviously encouraged more manufacturing, especially Taiwanese manufacturing investment, in the U.S. Thinking along those lines of development and how we can be opportunistic about this transition or reshoring, companies like Acter that we had met previously came up. And having done the work, it helped us to



identify the opportunities much more quickly.

[15:20] [RC] Just to identify for listeners who may be less familiar with Acter, can you explain what they do?

[15:27] [WC] Acter builds clean rooms and infrastructure that allows high-tech factories to operate. 70% of their business relates to semiconductor facilities, and that includes things like clean room construction, air conditioning, power, MEP (mechanical, electrical, plumbing), and the water systems. Also things like connecting production equipment and then managing the engineering and construction process for these clean room facilities.

It already has a very established position in the semiconductor ecosystem, including work for TSMC. It's one of the three qualified clean room contractors for TSMC, which says a lot about the quality and the barriers to entry in this industry, because if you get the clean room wrong, billions of dollars of chips get wasted.

In addition to serving TSMC, they also serve customers along that entire supply chain, from PCBs [printed circuit boards] to even the advanced packaging lines that are used to make the AI chips that we see today. Across much of that supply chain, clean room capacity has become a bottleneck, and that has supported very strong demand and pricing for Acter. That's leading to margin improvement and a huge and growing backlog. We can see very strong structural growth over at least the next three to five years.

[16:58] [RC] You mentioned margins, Wen, and I came across a chart recently that I found fascinating to think about. Typically, when you think about something like Apple and the iPhone, most of the margins go to Apple, and as you go further and farther out the supply chain, the margins get smaller and smaller. At least that's my understanding. This chart I came across, though, showed that it was the reverse for the AI space, where actually the further you went down the supply chain, often the bigger the margins got. Do you have any insight into why that is and how sustainable it is?

[17:30] [WC] I think it comes from two main things. The first one is consolidation across that value chain. Because requirements for AI accelerators are getting higher and a lot more complex, the competitive landscape has become a lot smaller; the ones that can target those areas are fewer. So it's a more consolidated space with better margins.

The second part, I think what's leading to a lot of this margin expansion today is the astronomical demand that we see.

Yes, the bottleneck effect, but more importantly, we have very rich end customers. In the past, a lot of these semiconductor supply chain beneficiaries were mainly to do with consumer electronics, where demand is a bit more uncertain and a lot more price elastic. Whereas what we see today is a lot more price inelastic, richer customers. We see the end customers, the hyperscalers, getting good ROI on their CapEx investment, and that's increasing the propensity for them to pay more and get more supply to the market. With that, we see higher margins and very strong growth along the supply chain.

[18:50] [RC] I know on this trip you spoke to a number of major players. Any specific insights from that with respect to the demand question earlier, or the capital cycle?



[18:59] [WC] I had a conversation with TSMC and the TSMC supply chain, as well as the major memory players, namely SK Hynix and Samsung. My main takeaway is that there are reasons this investment cycle could be more durable. Of course, we also have to respect the possibility of potential overbuilding down the road.

The conversation with TSMC was particularly useful in thinking about the quality of the investment and the entire build-out that we see today. Today's build-out has a stronger financial foundation than the dot-com boom: the biggest cloud companies have substantial, profitable existing businesses that generate the cash to support the investment. But having money to spend is only one part of the equation, because the other part is that spending ultimately needs to produce an economic return.

What I found interesting was how TSMC approached that question. They described looking beyond the forecasts from their direct chip customers and speaking with the customers further down the chain, including cloud providers, AI developers, AI frontier labs, and other infrastructure buyers, and using that to try to understand the underlying demand, building a bottom-up analysis on it: who is funding it, and whether those companies have the financial capability to follow through. Having heard that gives me a lot more confidence in the process behind their investment decisions. Obviously, they also have lots of information that they can see that we can't. And because they are conservative, it still leaves a buffer and room for demand to arrive later than expected, or for the industry to add too much capacity.

But their broader point was about getting the longer-term direction right, and the longer-term direction is the need for more powerful, energy-efficient computing. That has guided their investment through both the smartphone era, then HPC, high-performance computing, and now the AI era. With that, they have the conviction to spend on CapEx, which is a huge sum for them. They always have to be very conservative in their planning as well. There might be short-term surprises and short-term inventory corrections that we might not expect.

When speaking to the memory players, it was interesting to hear their perspective, because both companies align on the same kind of argument for why they think it's different this time. Historically, a memory producer might build capacity based on relatively short-term expectations, and most of the time that relates to consumer electronics, like PC and smartphone demand, and those expectations can change very quickly. But with AI infrastructure, this allows customers to plan several years ahead. For example, a hyperscaler committing to data centers, AI accelerators, and networking, and the memory supplier investing in the fabs and clean rooms that need to support that build-up. So there is more stability and more long-term planning.

Another aspect of that would be longer-term agreements, and in some cases deposits or prepayments, that help to connect the investment plans both of the memory players, in terms of capacity, and their customers' investment plans as well. That gives suppliers better visibility and can potentially also reduce the risk of expanding against speculative orders or double ordering. My interpretation is that these commitments could also help lengthen the cycle and provide a stronger floor for demand, not only for AI, but for memory demand as well.

[22:54] [RC] We've seen some pretty big stock price movements across that value chain. Obviously it's been a lot on the way up, which has been great, but more recently a little bit on the way down. How have you and the team dealt with the signal coming from the market? Has it changed, or have we rethought any of our holdings in light of that?



[23:13] [WC] I think it's telling to note that the entire AI theme was crowded at one point. Going to conferences, I see that with the number of participants. At the most recent conference I attended, at any time in the room there were 70 people. But I think what has given me a bit more comfort today is there is an even more healthy tension between the bulls and the bears.

Internally, during the drawdown, the long-term theme remains intact. I think that's the first point I want to make, and that guides our investment in the companies. The long-term theme's intact. Yes, we realize it is a crowded trade, so a lot of our stocks within that one month were up very strongly as well. We took the opportunity to be disciplined and reduce some of our exposures.

To your question, yes, we did have somewhat of a temperature check. It's obvious when you go on research trips: the entire AI theme was getting very crowded. There were several meetings that I attended where I was quite surprised to see a lot of investors who weren't particularly focused on Taiwan. They were very region-specific investors, for example, in Thailand, Vietnam, and India. I saw them all at a Taiwan conference. In many meetings, there were 50 to 70 people. So obviously Taiwan and the entire AI theme was getting a bit crowded.

We did take the opportunity to trim some of our names on the way up as well. We also did a sanity check on our models, and the ones that were a lot more expensive and harder to justify, we trimmed them. Maybe not at the peak, but we made a lot of money from these names, and we trimmed them just from that temperature check exercise.

More importantly, I believe the long-term structural theme of AI continues to exist. We had trimmed earlier on and took down some of our exposures, and we trimmed some of the names that didn't go down too much as well. We were able to redeploy money back into some of the names that we found to be quite attractive at the trough of that entire drawdown. So yes, we did trim and invest, but at the end of the day, it all goes back to fundamentals, it all goes back to risk-reward. We size appropriately and take the opportunity if we believe that the reward outweighs the risk. We're still actively on the hunt, and there are several names still on our radar that we hope will be cheaper, and several names that we have been adding over the past couple of weeks.

[25:48] [RC] Whether on this most recent trip or otherwise, what's got you most excited that's not in the AI bucket?

[25:55] [WC] Most of the companies on this particular trip were AI-related, because AI is obviously quite a big chunk of the portfolio, and obviously I was in Taiwan.

[26:03] [RC] There's not a whole lot that's not AI in Taiwan.

[26:06] [WC] Yes, but on the side, we have been doing a lot of work on the non-AI stuff as well.

Two names that we've recently initiated in. The first one is Asymchem Laboratories. I don't want to spoil Josh's story. We'll wait for his podcast, and you can hear all about that theme as well. I think it's really interesting.

Another company that we invested in earlier this year would be a company in the Philippines called International



Container Terminal Services, ICTSI. They're port operators, and they specifically focus on origin and destination ports, O&D ports for short, which have historically been a bit more of a price setter and display a bit more defensive characteristics compared to a port like a transshipment port, which tends to be more of a price taker.

Why we like it is, it's a good, stable business, but more importantly, it's a big bet on management. Because AI has made things a lot more efficient, I find this year I spend a lot more time on the road doing a lot more interesting things, in my opinion. So Shan and I flew to the Philippines just to attend the full-day Investor Day, the first time the company has held an Investor Day, and we were really blown away by the quality of the management team. Valuations were very undemanding.

They have a structural runway, not only from the existing portfolio that they have and the reshoring that we see benefiting the ports in countries like in Latin America. But more importantly, there are also hidden growth options, somewhat like a free option. They were bidding for a concession in Brazil. They're one of the two players left in the race, and management thinks they have a good chance of winning. It's going to be material to their financials if they do win such a concession, and they do have a very strong track record of operating an asset really well. So they have strict requirements. If this works out, it would be very wealth creating, and it would render the company even more attractive. So that's a more recent addition on the non-AI side of things.

[28:12] [RC] Wen, thank you so much for sharing both your AI and non-AI perspectives. As always, it's been great having you.

[28:13] [WC] Thanks.

[RC] Hi, everyone, Rob here again. To subscribe to The Art of Boring podcast, go to [Mawer.com](https://mawer.com). That's M A W E R dot com forward slash podcast, or wherever you download your podcasts. If you enjoyed this episode, please leave a review on iTunes, which will help more people discover the Be Boring, Make Money philosophy. Thanks for listening.

Companies Mentioned:

Hugging Face
Anthropic
OpenAI
Ramp
Andreessen Horowitz -
X
Acter
TSMC
Apple
SK Hynix
Samsung Electronics
Asymchem
International Container Terminal Services (ICTSI)