OVERVIEW

The Geometrica Fund returned **-1.0%** in July, net of all costs bringing our calendar year to date return to **+10.7**%.

We discuss the contributors and detractors below, but it is worth touching on the sharp moves in USDJPY and short end US interest rates, which triggered a short and sharp dislocation in US and Japanese equities in late July and early August, which we took advantage of to add to the portfolio.

In mid-July both the S&P500 and Japan's Nikkei index hit record highs, but by August 5th, the former would be down 8.5% and the latter down an eye watering 25.5% from its peak.

The market has become increasingly concerned that the Fed is overplaying its hand and keeping interest rates too high in the face of moderating activity. So, when weak US economic data seemed to crystalise fears of a possible US recession, US short end bonds spiked and yields fell as speculators pulled out of equities and piled into bonds. This then reportedly triggered a sharp unwind of the Yen carry trade[1].

It is hard to be precise on the size of the Yen carry trade but as interest rates have risen in a post ZIRP world, it's a sure bet that so has the size of the USDJPY carry trade (some estimate up to US$1 trillion). The flow into it is reflexive, so the mental model for any unwind is a stadium crush...lots of people all wanting to exit at precisely the same time.

USDJPY shot from 160 to 145. But rapid Yen appreciation wipes out the future earnings of Japanese exporters, so Japanese equities then dutifully went into a tailspin.

The panic selling that ensued threw up a few opportunities we took advantage of and we've added around 10% of capital to our longs since July month end.

One new stake is in a Japanese maker of specialised bandwidth products used in data centres. The market has been slavishly focused on the exponential growth of GPUs and generative AI data centres. But as these AI data centres scale, the call on bandwidth scales faster and if it doesn't, then idle time for all those expensive GPUs increases. Our company is rapidly expanding capacity, meaning its sales and earnings are also about to shoot up. We think they'll need further expansions after this wave to keep up with demand.

| US 2yr T bond yield | USDJPY | Nikkei Index |
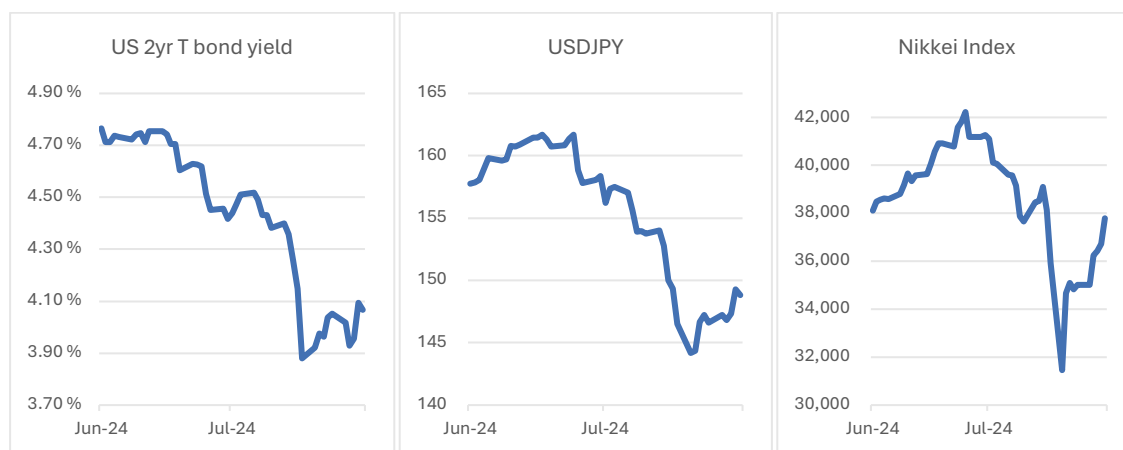
Source: Bloomberg, Geometrica

---

[1] The Yen carry trade involves selling low yielding Yen to buy a higher yielding currency such as the AUD or more recently the USD. The typical expression recently has been: sell Yen short end rates and buy US$ short end rates...picking up the interest rate differential.

PORTFOLIO

Carnival Cruises (CCL.US, mkt cap US$18.5bn) fell -11.0% in July on no news. The timing of the fall dovetailed with CCL being caught up in the general growth scare / Yen carry trade unwind.

Typically, late in an economic cycle, consumer cyclical stocks are on the nose with investors wary of the impacts of a slowing economy. We have already seen typical evidence of this in statements made by a few airlines and hotel stocks, which have cited softening forward bookings.

In Carnival's case, structural changes in the business and industry mean forward bookings, the highest on record and still building, bridge across near term uncertainty, especially given valuation has fully derated to prior pre-covid lows.

The lowest level of industry supply growth in a decade, coupled with a demand shift from higher priced vacation forms to cruise and a structural change in forward bookings means Carnival has the best forward demand visibility it has ever had. Better still, high frequency demand indicators remain rock solid.

Prospective free cashflow yield is north of 20%, backed by customer deposits. So whilst we have used the sell off to downsize cyclical exposure and upsize duration across the portfolio, Carnival has been retained. We think on any reasonable timeframe there is material upside.

DOF Group (DOFG.NO, mkt cap €1.6bn) announced the acquisition of Maersk Supply Services, as a means to add young and highly specialised additional Offshore Support Vessel capacity at a very attractive price in July. Market reception was enthusiastic, with the stock closing the month up +16.7%.

The acquisition was partly funded by DOF issuing US$500m in 7 year amortising bank debt on terms that would not have been possible 12 months earlier. Capital markets are once again open to DOF.

The acquisition sees a general reduction in leverage relative to earnings (DOF also raised equity) and the potential to accelerate distributions to shareholders via the addition of an asset grouping not subject to a creditor cash sweep.

As the company's cost of capital normalises (distribution reinstatement will be a key milestone here likely in 2025) we would expect that to flow through positively to valuation. The stock is trading on a forward free cashflow yield of over 20%.

We used the strong share price surge of recent months to trim position sizing, noting that the current sizing is not too different to when we established the position at a far lower share price.

Nvidia (NVDA.US, mkt cap US$3.0 trillion) fell -5.3% in July. Nvidia is possibly the most polarising stock in the market. We added back to our position as the recent sell off ebbed.

We look for stocks offering a steep discounts to their underlying value due to one of two inefficiencies as depicted below. The bar bell on the right, where something is truly exponential, is rarely seen.
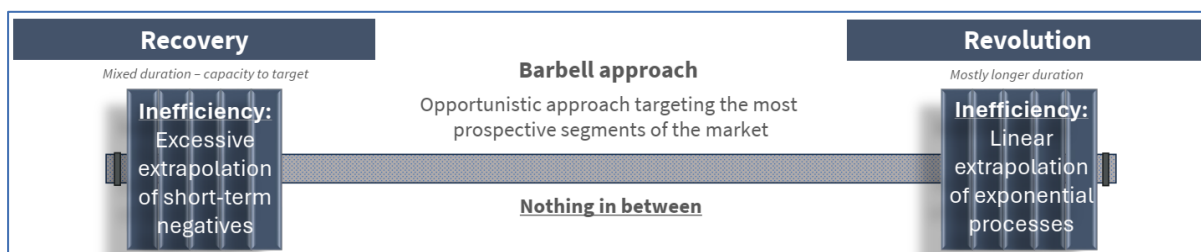
The quote below describes an earlier demand inflection in semiconductors that flowed from exponential improvement in processing capacity, which in turn drove down unit cost and saw explosive adoption. In this example, it was the CPU:

*"I said to him, "Oh, **exponential phenomena are pretty rare, pretty dramatic**. Are you serious about this? Because this means, in effect, we can think of computing as free."*

And it was really his insight that because of semi-conductor improvements, things would just keep getting better[2].

Bill Gates, 1993 interview.

The recurring question thrown at every major hyperscaler in the recent reporting season could be summarised as "*is all this spending on GPUs and AI really going to pay off?*".

Most companies answered in the affirmative by increasing their capex spend on AI and articulating a multi-year process, not too different to cloud and early cloud focused data centre investments made where returns lagged the initial investment.

Applovin (APP.US, mkt cap US$30.1bn) fell -7.4% in July. Another victim of the market selloff, APP reported particularly strong earnings in early August and has subsequently recovered all of its July share price decline.

APP has been rapidly taking market share in its main vertical, which is the historically fast-growing area of mobile game advertising.

Similar to Meta, but on a smaller scale, APP is offering buyers of advertising capacity an enhanced ability to target prospective customers and precisely measure ROAS or return on advertising spend. This is a seismic shift from the traditional model where an advertiser allocated a fixed budget to a wide audience and spent with limited capacity to measure return on investment. It is clearly driving spend down the funnel to more targeted forms of advertising, a boon for APP.

APP is now attempting to move outside its core vertical. There is nothing in the share price for success; so in essence the set up is asymmetric. We are watching this very closely.

Micron (MU.US, mkt cap US$ 105bn) fell -16.5% in July. Micron is one of three makers of dynamic random access memory (DRAM) which is a faster, impermanent version of hard drive memory used in every device with any computing capability from calculators and iPhones to laptops and high performance servers.

When you look at the semiconductor sector today, the current upswing in activity is extremely narrow and largely confined to generative AI.

DRAM manufacturers are a key beneficiary of this generative AI demand via a specific type of memory called high bandwidth memory (HBM). As the name suggests, HBM has better throughput which is crucial for AI algorithms.

The crux of our thesis here is that if demand for GPUs continue, then so too does demand for memory. As this demand is derived from bigger models, the **effect on memory bandwidth requirements and therefore gigabytes of on-chip memory should be multiplicative**.

We discuss this in greater detail in the appendix at the end of this letter.

Rightmove Plc (RMV.LN, mkt cap £4.2bn) rose +7.3% in July. We invested in the stock back in late 2023 when it swooned on the news that CoStar (CSGP.US, mkt cap US$29.5bn) had acquired the #3 player OnTheMarket in the UK).

The negative market reaction to CoStar's entry provided us the opportunity of buying in at a steep discount, at a point in time likely to prove to be at or close to the cyclical low in UK housing.

---

[2] Bill Gates, 1993. Emphasis added.

In July, the Bank of England cut its official interest rate by 25 basis points. Housing and housing activity tends to be highly interest rate sensitive, albeit with a lag. Since that time, Rightmove has reported an increase of 19% in buyer inquiry volumes.

Resmed (RMD.AU, mkt cap A$48.6bn) rose +11.6% in July. The last two results from Resmed have positively surprised, as gross margins have recovered in tandem with normalising inventory levels.

In this most recent result, the key positive was an acceleration in mask resupply rates. Masks carry much higher margins than machines and Resmed has been adept and lifting mask order rates, driven by what is reimbursable under insurance plans and also driving greater usage compliance for insurers.

The stock remains cheap relative not just to its history but also its forward earnings prospects. GLP-1s do represent a long-term threat but in the short to medium term are likely to drive greater awareness and thus greater adoption of obstructive sleep apnea treatments. CPAP remains standard of care and we think we'll see that reaffirmed in the earnings of Resmed in the years ahead.

Ambu (AMBUB.DC, mkt cap €5.0bn) rose +4.4% in July. Ambu recently rolled out several new endoscopes targeting niche uses. Underlying growth for Ambu's products is accelerating, evidenced by two sequential revenue and earnings upgrades thus far in 2024. Two new products are likely to further accelerate revenue and earnings growth into fiscal 2025 given they address predominantly new procedures.

Ambu is the largest player in single use endoscopes, a market it pioneered. Growth rates for single use endoscopes, which offer significantly lower risk of infection, are around 3x reusables. We continue to see material upside to Ambu, given still immature margins, and a very long structural growth runway.

Novo Nordisk (NOVOB.DC, mkt cap €534bn) fell -9.3% in July, after falling victim to the market wobbles.

On August 7[th] Novo released earnings which saw the stock down 7%. The negative reaction stemmed principally from two perceived issues.

First, operating profit guidance for 2024 was lower than expected. This stems from Novo including in the guide certain one-off costs that won't repeat in 2025. One off costs are generally something we look through.

Second, price realisations for Wegovy were lower than expected. Rebates are prevalent for larger payers in the industry, with focus on so called "gross to net" pricing adjustments. And rebates are typically larger for government linked payers.

In the case of the US, a market almost as large in annual value terms as the rest of the world combined, Medicare Part D precludes coverage of weight loss drugs.

However CMS (the umbrella organisation for Medicare and Medicaid) declared recently that where a GLP-1 drug is also indicated for say, kidney disease or heart disease, it will be covered by Medicare for those indications. This is a major boost for Novo's US addressable market, and a clear positive.

The negative is the price they'll pay is lower than rack rate. But having CMS, which covers over 30% of US lives, in the GLP-1 tent is a massive positive, given a year or so ago, it was thought CMS would never reimburse weight loss drug use.

Paladin (PDN.AU, mkt cap $3bn) We exited most of our positions in Paladin and other uranium linked equities during the month as spot prices fell ahead of production guidance from Kazatomprom (KAP) which surprised to the upside. Paladin and NexGen fell -8.7% and -4.4% respectively and have continued their declines in August.

From entry to exit three uranium equities, Paladin, NextGen and Sprott, have contributed +1.66% of the fund's returns. We really do work hard to avoid the profitless "round trip".

Kazatomprom remains the key swing factor in uranium supply against a backdrop of steadily growing demand. Their second quarter trading update saw increased production guidance for the remainder of the year. It was known that KAP had procured "enough" sulphuric acid for 2024, but this may imply an easing acid supply picture since last quarter.

To be clear, KAP's updated production forecast remains materially below their original plans for the year, and we have not seen any significant changes to Kazakhstan's acid imports, domestic production or consumption. The increased production guidance did not translate to higher sales as KAP looks to take pressure off its inventory which on our math, was a touch below eighteen months. And in any case, production on an attributable basis, despite the upgrade, remains well below sales volume.

It remains to be the case that what KAP can produce in 2025 is the lynch pin in the uranium thesis, either long or short. Along with clarity on ease of access to waivers on the United States partial ban on Russian enriched uranium, nuclear utilities appear likely to wait for KAP's guidance in late August before re-entering the term market, activity in which has been particularly muted in recent months.

We maintain a watching brief.

APPENDIX: DRAM

Semiconductors are cyclical and DRAM yet more so. Over the years the downcycles have seen many DRAM makers weeded out, to the point now where the sector is highly consolidated.

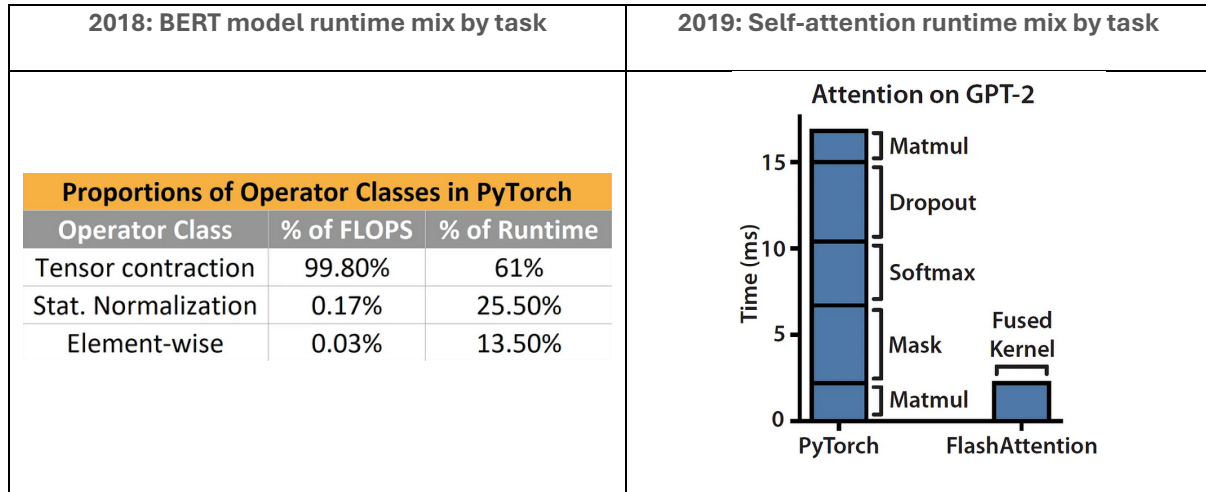**Number of DRAM producers**
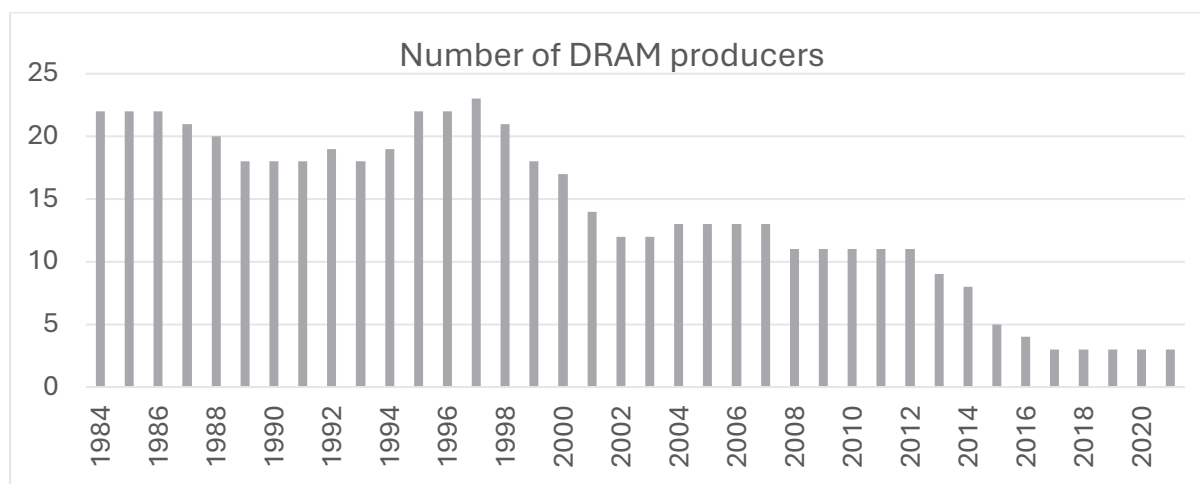
Source: Geometrica

When you look at the semiconductor sector today, the current upswing in activity is extremely narrow and largely confined to generative AI.

DRAM manufacturers are a key beneficiary of this generative AI demand via a specific type of memory called high bandwidth memory (HBM). As the name suggests, HBM has better throughput which is crucial for AI algorithms.

Demand for GPUs has thus far been exponential, driven by the transformer model and internet sized dataset scaling paradigm.

People often talk about how much compute will be required for AI, but as a percentage of training runtime, memory-bound tasks have been increasing rapidly in the mix.

You can see below that even back in 2018, *memory-bound processes made up almost 40% of LLM runtime despite only using 1% of the compute*. It has become more difficult to calculate in general today but, zeroing in on self-attention (the beating heart of the transformer model) we can see that *memory bound tasks make up ~85% of runtime*.

2018: BERT model runtime mix by task	2019: Self-attention runtime mix by task

Proportions of Operator Classes in PyTorch

Operator Class	% of FLOPS	% of Runtime
Tensor contraction	99.80%	61%
Stat. Normalization	0.17%	25.50%
Element-wise	0.03%	13.50%

Source: https://arxiv.org/pdf/2007.00072, https://arxiv.org/pdf/2205.14135

NB: tensor contraction or matmul (same thing) are compute-intensive tasks that do not require large amount of memory whereas stat. norm., element-wise (general term), dropout, softmax and mask are largely memory bound. Breakthroughs like flash attention help this but they just outsource what would be done by DRAM to SRAM, which is orders of magnitude more expensive per GB.

Consider a bus that arrives every five seconds but that only fits one person and a bus that arrives every thirty seconds but that fits six, each minute these busses transport the same amount of people. When calculating how many people you could transport, increasing the speed at which a bus arrives is meaningless without the context of how many people it can fit. **A fast GPU can only process as many calculations as it can access the data for**.

And indeed, we can see that bandwidth on chip has been scaling but has increased significantly less (~10x) than compute (~200x). And this has been true for a while; over the past two decades, CPU and then GPU compute has increased by 90,000x but memory bandwidth has only increased by 30x.
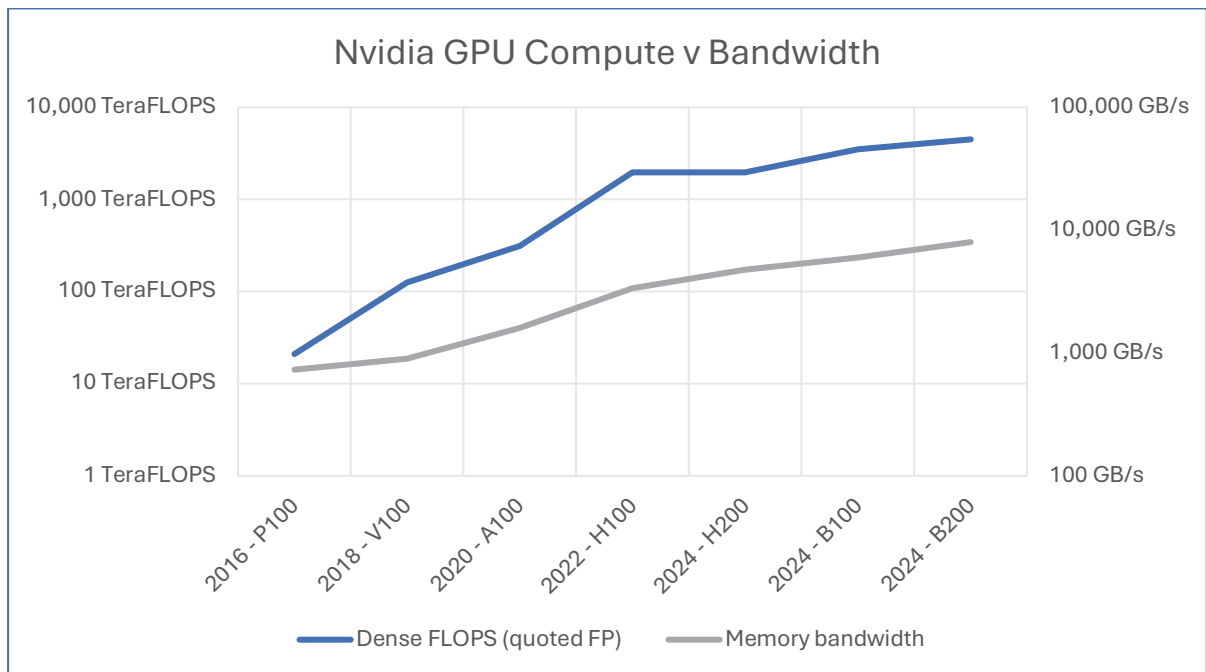

Source: Geometrica, Nvidia NB: chart is in log scale

To illustrate this in a way that highlights GPU performance more accurately, we calculate FLOPS x bits / bandwidth. On this metric, you measure the ratio of how much data your chip can process per second against how much data it can transfer per second. So, a higher ratio implies worse compute utilisation. Ever since the H100, Nvidia has been unable (or unwilling... in order to sell more chips!) to keep this metric below one.
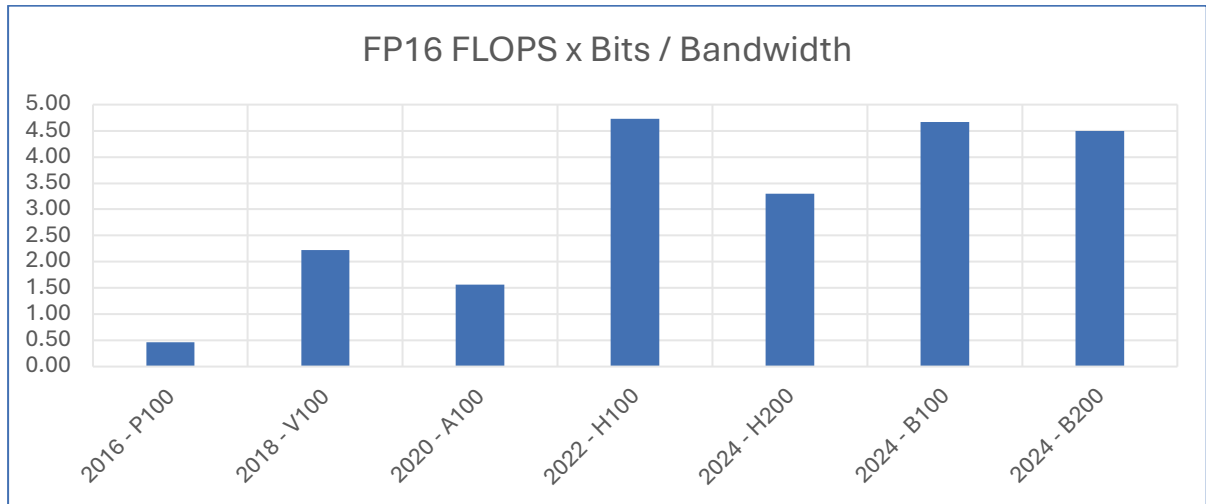
Source: Geometrica, Nvidia

Through the lens of history, this is not a new trend. In the chart below, only the chips in the bottom green quadrant have balanced their computing resources with their memory bandwidth. For non-HPC applications, this is probably fine but when we are talking about multi-billion-dollar CAPEX buildouts where every second equals margin, this is a problem.
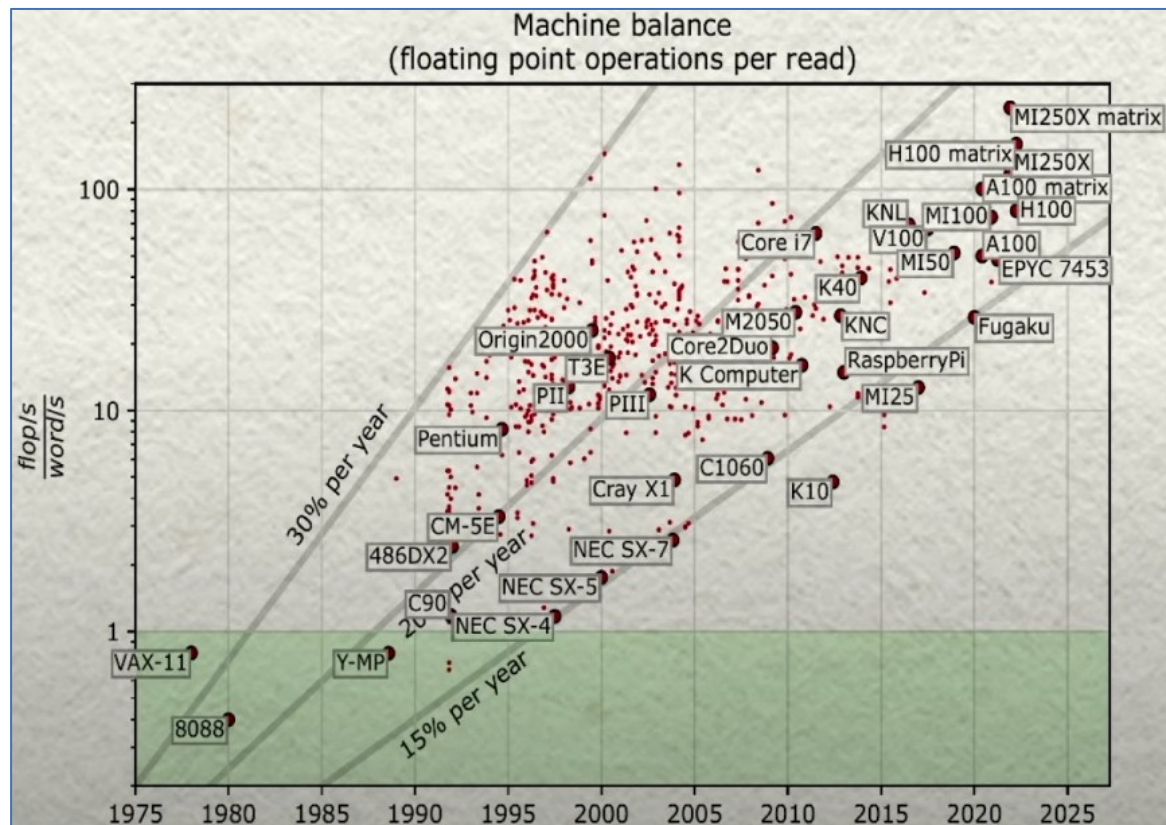
Source: https://www.youtube.com/watch?v=cSO0Tc2w5Dg see timestamp 48:29

If Nvidia wanted to balance their engine with enough fuel injectors to achieve maximum utilisation for, say, the B100, the total cost of HBM would be equal to ~80% of the indicated RRP of the whole chip.

The HBM tail is wagging the GPU dog...in practice, so long as Nvidia's quasi-monopoly status is maintained in AI accelerator installations, it doesn't matter so much, evidenced by their rapidly expanding gross margins i.e. they seem to have been able to pass this cost on.

If demand for GPUs continue, then so too does demand for HBM. As this demand is derived from bigger models, the **effect on bandwidth requirements and therefore gigabytes of on-chip memory should be multiplicative**.

SK Hynix has to date been the principal beneficiary of this AI related demand as it was the first DRAM producer to achieve satisfactory HBM performance to be qualified for Nvidia's chips. It is thus the main supplier of HBM or high bandwidth memory used in AI chips today. Micron is also a qualified supplier but until next year, has a much smaller scope of product. Samsung has recently all but confirmed that their qualification is nigh. Market shares are key, and Micron appears confident in its ability to reach a significantly larger share over the coming twelve months.

Outside of HBM, this is leading to capacity shifts in standard DRAM as the silicon intensity of HBM is ~3x that of standard DRAM. Because HBM is in demand and making a lot of money for Hynix in particular, producers are repurposing their standard DRAM production capacity to HBM production.

Standard DRAM pricing is tightening, as supply capacity is tightening and this is occurring in weak demand conditions. If we do get a typical cyclical recovery in non-AI volume demand, there will likely be materially more upside in Micron.
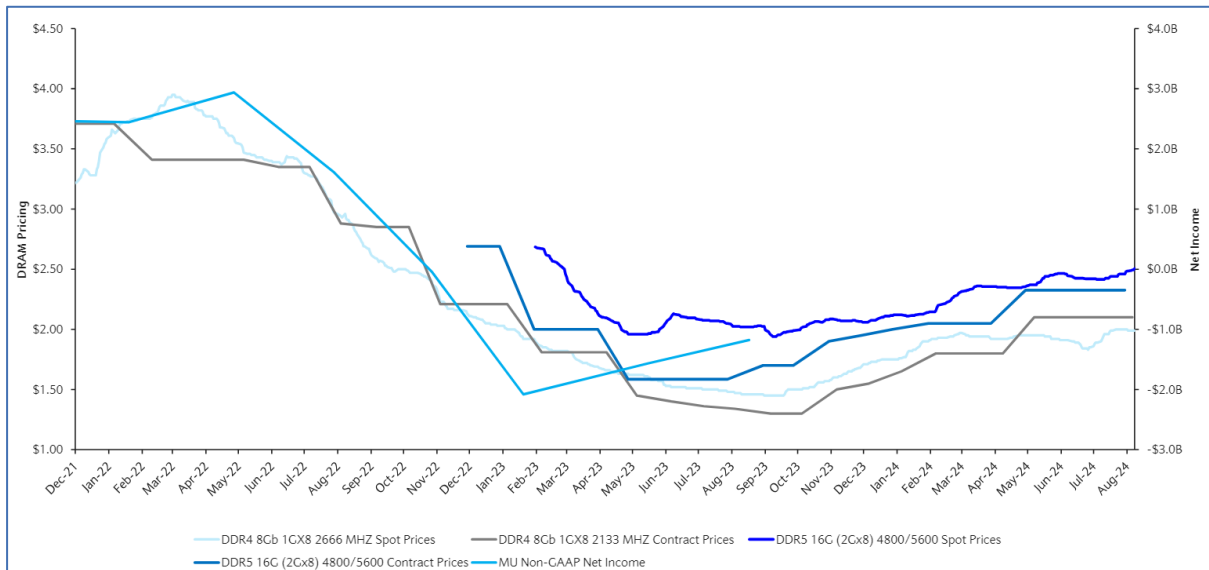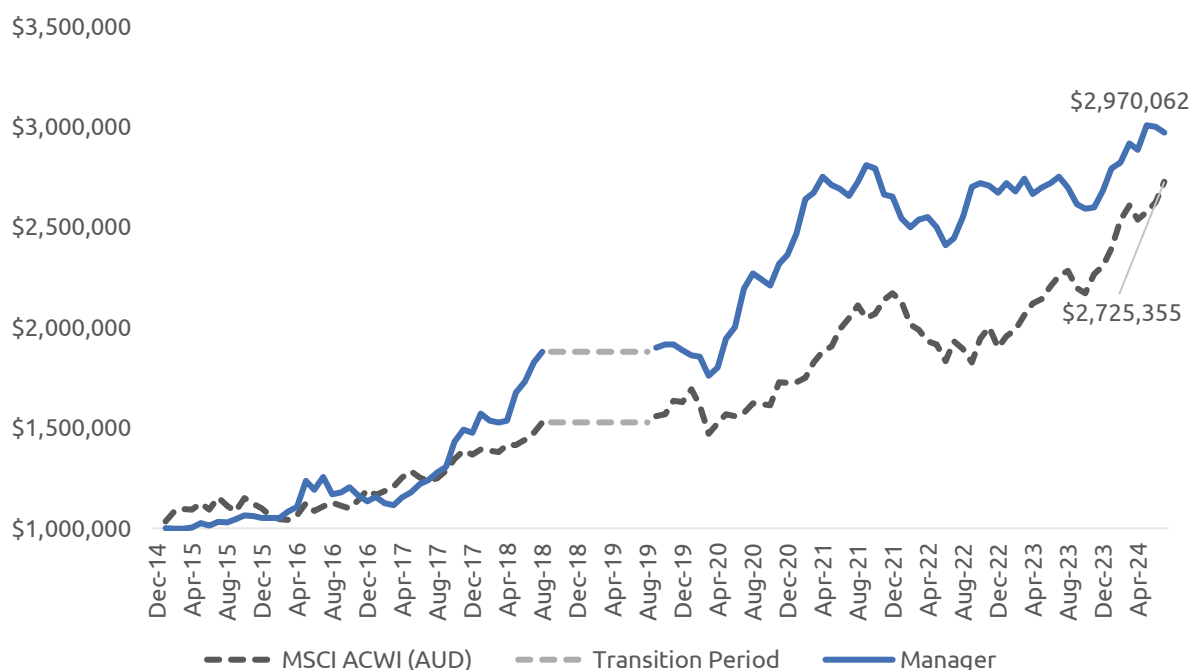

Source: DRAMeXchange, Barclays, Micron

PERFORMANCE (%, NET)[3]

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Fund
2015	+0.1	-0.3	+0.0	+0.7	+2.1	-1.4	+2.0	-0.2	+1.5	+1.7	-0.2	-0.9	**+5.1**
2016	+0.0	+0.0	+2.9	+2.3	+11.8	-3.6	+5.3	-6.8	+0.6	+2.3	-3.5	-2.5	**+7.7**
2017	+1.9	-2.6	-1.0	+3.5	+2.2	+3.5	+1.7	+3.0	+2.4	+9.5	+4.1	-1.1	**+30.3**
2018	+6.5	-2.1	-0.7	+0.6	+9.1	+3.2	+5.6	+2.9	--	--	--	--	**+27.4**
2019	--	--	--	--	--	--	--	--	+1.1	+0.8	+0.1	-1.6	**+0.5**
2020	-1.3	-0.3	-5.2	+2.4	+7.9	+3.0	+9.5	+3.5	-1.4	-1.4	+4.8	+2.0	**+25.2**
2021	+4.5	+6.9	+1.2	+3.0	-1.5	-0.7	-1.4	+2.6	+3.1	-0.6	-4.6	-0.4	**+12.2**
2022	-4.1	-1.7	+1.5	+0.5	-2.0	-3.5	+1.4	+4.4	+5.8	+0.7	-0.4	-1.3	**+0.8**
2023	+1.7	-1.5	+2.4	-2.7	+1.2	+0.8	+1.2	-2.0	-3.1	-0.8	+0.2	+3.3	**+0.5**
2024	+4.0	+1.0	+3.4	-1.1	+4.3	-0.2	-1.0						**+10.7**

2015 – 2018: CVF (same portfolio managers and strategy)
2019 onwards: Geometrica.

Manager ITD	+197.0
Manager p.a.	+13.5
Geometrica p.a.	+9.8

PERFORMANCE CHART (%, NET)[4]

$2,970,062

$2,725,355

Legend: MSCI ACWI (AUD) · · · · · ·   Transition Period - - - - -   Manager ——

---

[3] Performance is after all fees, Founder Lead Series unit.
[4] 2015-2018: CVF (same portfolio managers and strategy); 2019 onward: Geometrica. Index = MSCI All Country World Index (AUD)
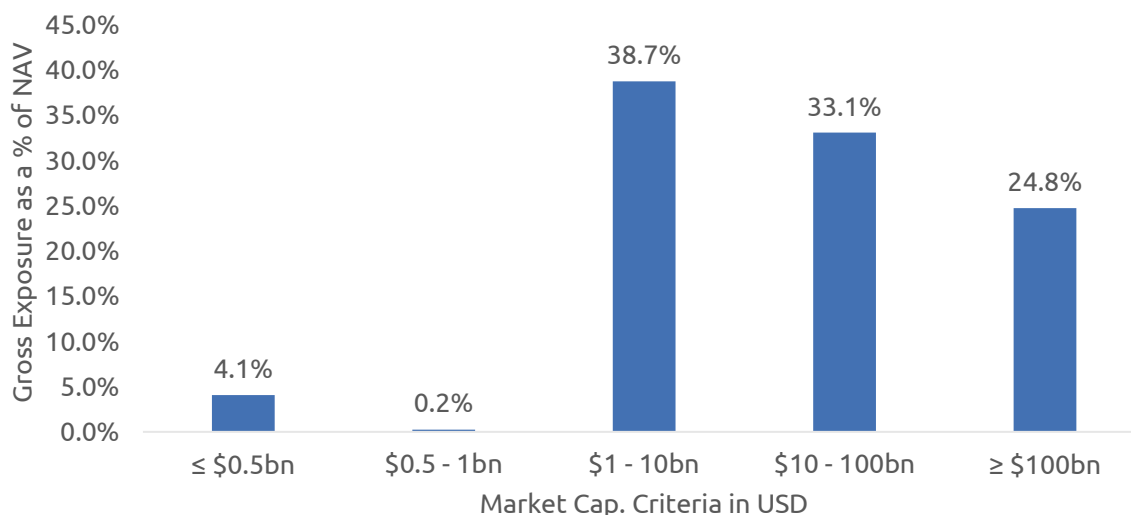
Geometrica Fund

ASSET ALLOCATION

Country	Long	Short	Gross	Net
Australia	3.3%	(3.8)%	7.1%	(0.5)%
Americas	44.4%	(7.0)%	51.4%	37.4%
Asia	4.5%	(3.4)%	7.9%	1.1%
Europe	32.7%	(1.7)%	34.4%	30.9%
Total	84.9%	(16.0)%	100.9%	68.9%

GROSS & NET EXPOSURE
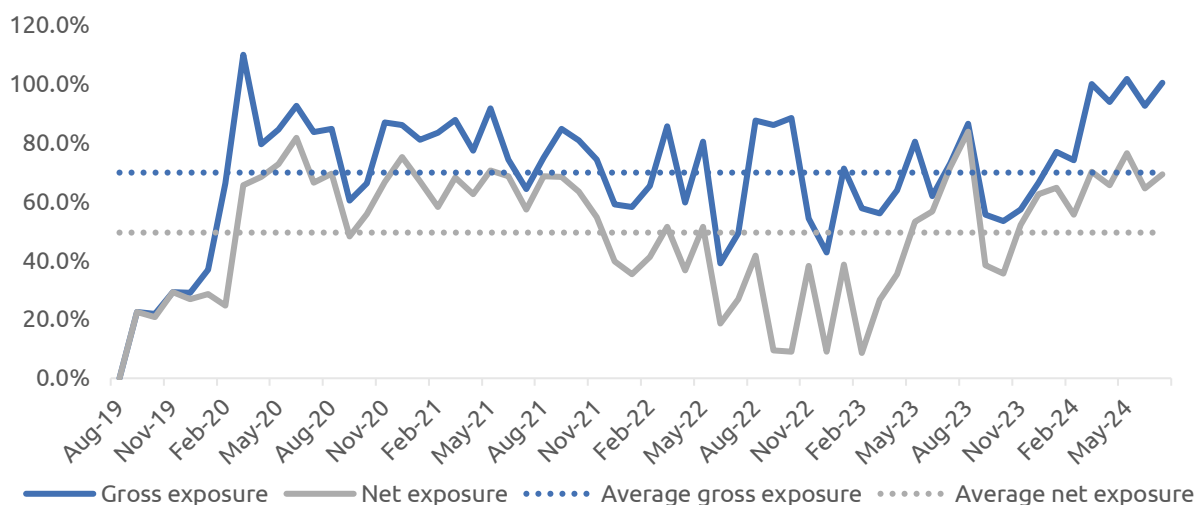
GROSS EXPOSURE BY MARKET CAPITALISATION

FUND OVERVIEW

Fund	Geometrica Fund
Structure	Wholesale unit trust
Mandate	Global long short Mid-cap focus
Gross exposure range	0 - 200%
Net exposure range	up to 100%
Single stock long limit	15% at cost
Single stock short limit	5% at cost
Buy / Sell Spread	Nil / 0.25%
Investor Eligibility	Wholesale only
Platforms	Ausmaq, Hub24, Powerwrap, Netwealth
Fees (Founders Class)	1% management (+GST) 15% performance (+GST)
Benchmark	RBA Cash Rate
High water mark	Yes
Liquidity	Monthly
Administration & custody	Apex
More information	www.GeometricaFund.com