


The Bronte Amalthea Fund is a global long/short fund targeting attractive risk adjusted returns over the long term, managed by a performance orientated firm with a process and portfolio that we feel is genuinely different. Objectives include lowering the risk of permanent loss of capital and providing global diversification while reducing the market/drawdown risks typical of long-only funds. We believe a highly diversified short book substantially reduces risk and enables profits to be made in tough markets

	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	FYTD
FY13											5.4%	1.3%	6.8%
FY14	6.0%	-2.5%	0.4%	3.6%	5.7%	4.3%	-3.7%	0.2%	-2.6%	0.9%	3.4%	-0.8%	15.2%
FY15	-0.9%	-1.6%	2.7%	1.7%	3.4%	4.9%	2.3%	-0.1%	1.7%	-1.7%	4.4%	-1.7%	15.6%
FY16	6.1%	0.9%	-0.2%	3.8%	-1.3%	-1.4%	0.5%	1.8%	-4.1%	-3.4%	5.1%	-3.4%	3.8%
FY17	2.5%	-0.8%	-2.5%	-1.3%	-1.5%	6.1%	-2.0%	1.6%	1.0%	7.0%	7.2%	-3.7%	13.6%
FY18	-0.9%	1.5%	1.1%	5.9%	-1.3%	-1.6%	4.4%	4.1%	1.5%	3.7%	-2.0%	2.9%	20.8%
FY19	0.1%	3.8%	-1.8%	-0.4%	-3.9%	6.5%	-3.6%	3.4%	0.0%	2.2%	0.1%	0.7%	7.1%
FY20	1.5%	-0.4%	1.3%	3.4%	3.1%	-2.1%	4.3%	4.2%	11.0%	-5.1%	-0.1%	-4.8%	16.5%
FY21	-0.1%	-3.9%	1.7%	-0.7%	-5.0%	-5.7%	-7.3%	-3.7%	8.2%	5.5%	3.2%	-2.2%	-10.7%
FY22	9.7%	3.0%	-4.5%	1.1%	1.8%	7.3%	4.4%	-5.6%	-4.6%	5.2%	2.2%	0.1%	20.7%
FY23	-0.7%	-5.9%	6.7%	6.5%	0.8%	1.0%	-1.8%	4.3%	4.9%	2.9%	-2.8%	-1.4%	14.7%
FY24	2.9%	7.3%	-0.4%	2.4%	-1.2%	-2.1%	6.6%	-0.3%	1.7%	-0.0%	0.7%	1.2%	20.0%
FY25	2.4%	-0.4%	-5.6%	-2.6%	-1.3%	0.5%	5.4%	2.3%	-3.0%	-1.7%	0.9%	-4.3%	-7.7%
FY26	-0.9%	-0.6%	-7.7%	-2.4%	5.8%	-2.2%	-3.5%	0.8%	-3.3%	0.2%	-1.5%	6.5%	-9.2%
FY27	6.5%												6.5%

Amalthea Cumulative Performance



In July, Amalthea gained 6.54% net whilst the MSCI ACWI (in \$A) fell by 1.55%. Performance in the long book was led by **Regeneron**, with **DSM Firmenich**, **Amazon**, **Visa** and **Philip Morris** also providing meaningful contributions. **Alnylam Pharmaceuticals** acted as the primary detractor while **Murata** and **5N Plus** also weighed on performance. On the short side, widespread market jitters worked strongly in our favour. As speculative excesses began to cool our short positions captured robust gains, concentrated heavily in the highly valued **Technology**, **Basic Materials**, and **Industrials** sectors. These positive shorting results were only modestly offset by minor headwinds in the **Consumer Non-Cyclical** sector

Fund Features		Metric	Amalthea	MSCI ACWI (in AUD)
Investment Objective	Maximise risk-adjusted returns over 3-years	Sharpe Ratio ¹	0.62	1.10
Min. initial investment	\$100,000 (for qualifying investors)	Sortino Ratio	1.08	1.91
Min additional investment	\$50,000	Annualised Standard Deviation	12.55%	10.60%
Applications/redemptions	Monthly	Largest Monthly Loss	-7.67%	-8.00%
Distribution	Annual	Largest Drawdown	-30.01%	-15.97%
Management fee	1.5%	Winning Month Ratio	0.56	0.65
Performance allocation	20%	Cumulative return ²	234.65%	467.55%
Administrator	Citco Fund Services	1-year annualised return	-2.38%	11.79%
Auditor	Ernst & Young	3-year annualised return	1.34%	16.59%
Custodians/PBs	Fidelity, Morgan Stanley, JP Morgan	5-year annualised return	6.22%	11.83%
		Annual return since inception	9.54%	14.00%

¹ Sharpe and Sortino ratios assume the Australian cash rate as the applicable risk-free rate

² Returns are net of all fees

Disclaimer: This report has been prepared by Bronte Capital Management Pty Limited. This report is for distribution only under such circumstances as may be permitted by applicable law. It has no regard to the specific investment objectives, financial situation or particular needs of any specific recipient. It is published solely for informational purposes and is not to be construed as a solicitation or an offer to buy or sell any securities or related financial instruments. No representation or warranty, either express or implied, is provided in relation to the accuracy, completeness or reliability of the information contained herein nor is it intended to be a complete statement or summary of the securities, markets or developments referred to in the report. The report should not be regarded by recipients as a substitute for the exercise of their own judgement. Any opinions expressed in this report are subject to change without notice. The analysis contained herein is based on numerous assumptions. Different assumptions could result in materially different results. Bronte Capital Management Pty Limited is under no obligation to update or keep current the information contained herein. Past performance is not necessarily a guide to future performance. Estimates of future performance are based on assumptions that may not be realised.

In normal times a 5.12% quarter should make us happy. We run with positive but low beta, and this performance is not something easily achievable.

This is not a normal time.

In the last quarter we have seen:

- SpaceX become the largest IPO in history,
- Elon Musk, the central figure in this bull market, become the world's first trillionaire, and
- AI capital expenditure (CapEx) needs become so vast that even Google (one of the most cash-generative companies in human history) has had to raise additional capital.

We have failed to keep up with a roaring bull market. If you run long-only money heavily exposed to the hottest of sectors (suppliers to the artificial intelligence CapEx boom) then you have made a lot more money than us.

The AI Investment Boom

The central thing happening in the world is the AI boom and we have no idea how it ends.

As with many periods of rapid technological progress, there is a market mania. Market manias happened with railroads (several times in the 19th century), with radio and automobiles (in the 1920s), and with the internet (late 1990s). Most manias ended in busts even when the technology was real and did change the world.

Just how good is AI?

Our readers vary from the most sophisticated in Silicon Valley to suburban dentists. At the risk of boring the sophisticated we should state just how far this has progressed.

The other day we were looking at a British gravel company. The gravel mostly goes into readymix concrete. The UK is a well-known economic basket case and construction (and hence this company's profits) are way below expectations. We have no idea whether the UK will turn around or whether it has been so mismanaged it is permanently impaired.

But we wanted an estimate just how bad it is right now versus history, and (by implication) how much more profitable the gravel company would be if UK concrete usage recovered to somewhere near the middle of its historic range.

We asked various AIs for a 30-year chart of readymix use in Britain.

There is no consistent series published. There are at least three historic series, not entirely compatible, but at least overlapping.

Making some assumptions about the overlapping period you can produce a *plausible* historic series.

Claude (Anthropic) unsurprisingly was the best: it produced an overlapping series, clearly stated its assumptions, and verified its assumptions about the overlapping period against other data sets. And it gave us all the caveats.

This took a few minutes. In the past it would have taken a good graduate employee a few days and we doubt the result would have been as accurate/realistic.

Als are rapidly improving on this

Joshua Gans (a respected economics/management professor) recently wrote that previous AI models (up to Claude Opus 4.8) were like enthusiastic but slightly incompetent PhD students, while Claude 5.0 “was something more; maybe along the lines of a future Nobel prize winner as a PhD student”.

This is a very high standard, and machines that think that well (and that fast) will change the world far more than we can imagine.

How much CapEx?

The highly compelling story about how AI will change the world (along with scaling models that posit a direct relationship between computing power and intelligence) is fueling monstrous amounts of capital expenditure.

AI CapEx estimates for 2026 alone range from US\$700 billion to US\$1 trillion. After increasing in 2027 and 2028, the total bill could exceed US\$3 trillion USD. Broader estimates (including implementation expenditures) for these three years top US\$6 trillion. Estimates by Gartner, a respected global advisory firm, are substantially larger.

As a percentage of GDP, these numbers are exceeded only by the 19th-century railway boom. However, the growth in spending (as a percentage of GDP) is faster than the railway boom.

How do you get a return on that capital?

As an investor, the obvious questions are:

- How does society get a return on all of that CapEx, and
- To whom does that return accrue?

This is far from obvious.

The last huge capital expenditure boom in technology was laying fibre-optic cables across the world. This globalised information and enabled much of the modern internet. Society enjoyed a substantial return on the investment.

The companies making the investments did not.

Global Crossing, Enron (broadband), and WorldCom all went to zero. Most of the other companies that spent heavily had woefully inadequate returns (at best) or took substantial write-offs.

The (giant) capital expenditure in laying out the global internet seems quaint next to the AI numbers. It is worth explaining how much return is required.

Let's guess that cumulative CapEx during the heat of this boom comes to US\$5 trillion (including implementation expenses). This is well below Gartner's estimate, but above estimates given to us by many AI models, and above plausible estimates for the infrastructure spend alone.

Because the spending is on tech, we know it's likely to be made obsolete soon (by even better technology). Suppose further you want — and this is perfectly reasonable in tech-land — a five-year payback.

That means you need to make \$1 trillion per year in gross margin. Because there are running costs (electricity, security, etc), it's reasonable to guess that the trillion in gross margin will require US\$1.8 trillion in revenue.

This revenue has to be paid by somebody. The jargon is that people will buy "tokens".

\$1.8 trillion won't be paid by consumers. Global streaming revenue is just over \$150 billion. The required revenue here is ten times as large. *The money therefore will need to come from businesses. To pick a really big number, global spending on food is under \$10 trillion. The required revenue for AI is nearly a fifth of global food spending.*

Bluntly, businesses will need to buy most of US\$1.8 trillion in "tokens" annually.

They will pay a lot for good reason...

It is attractive to have a research assistant on tap who is as good as a future-Nobel-Prize winning PhD student. Moreover, that "person" does not get bored, works 24/7, and can handle idiots with good grace.

They are amazing everywhere from customer service to senior management.

Still, they have to be paid in aggregate maybe \$1.8 trillion per year, and that has to come from *somewhere*.

There are two main possibilities:

- It comes from replacing labour or
- It comes from greater "factor productivity" (ie economic growth).

(Detailed accounting shows other possibilities but we think they are unlikely.)

In the long run, most of the returns from new technology will come from economic growth. The returns from the railways, for instance, came from opening new lands to higher uses such as growing food, cotton, or wool for export.

This is the economist's default position in the long run. Over the centuries technology has freed up labour and other resources and people have found new productive uses for those resources.

But in the short term the returns are almost certainly from labour replacement. The AI companies know this. Anthropic CEO Dario Amodei previously warned that AI could wipe out half of entry-level, white-collar jobs within a few years, potentially driving unemployment to 10% or 20%.

The AI companies have now walked away from these statements, focussing instead on the total productivity gains. But these gains — if they happen — will come from the labour savings. If AI is going to achieve the necessary revenue it will have to displace labour quickly. To the extent the labour finds new, productive uses, we will get total productivity gains.

Alas \$1.8 trillion per annum of token spend has to displace a lot of labour to be justified.

The total wage bill of the United States in 2025 was between \$12 trillion and \$13 trillion, based on the Bureau of Labor Statistics' employment and wages data set (with the range of uncertainty resulting from seasonal adjustments).

With the help of AI, we worked out the wage bill for the OECD as a whole is likely under \$40 trillion.

You are not going to sack a worker or change your process in the short term unless the AI saves you 50%. [Hiring workers is expensive, bringing in new processes is expensive, and you need proper savings to make big changes fast.]

That means we will need labour savings of maybe double the \$1.8 trillion in token spend in the next four or five years — nearly 10% of the entire OECD's wage bill.

An AI CEO like Dario Amodei would probably assume that AI companies will get a good return on all that capital they are spending. He really does believe in that much AI job displacement. And he wasn't pulling the number of 10-20% unemployment out of thin air.

Given who he is, he likely had early access to Claude 5.0 — and probably made that estimate with the help of AI.

There will be a bust, but it is not clear who will suffer it

The conclusion here is that either

- A. The AI companies will not get a return on the trillions of capital they are spending, or
- B. The AI companies will get a return but at the expense of very large short-term unemployment, or
- C. (Most likely) a combination of the above.

At the end of this, once we, as a society, find new uses for all that displaced labour, there should be a massive rise of total factor productivity. That has happened with every other new technology in history, but this time there may be a caveat: maybe the AI is so much better than any labour, no matter how skilled; and the AI is so cheap; that humans become redundant.

We are not going to explore that nightmare here. If you look online you will find plenty of people who have. And if you ask Claude really nicely, he (/she/it) will do it better than we can.

Doesn't the AI spend just stop

There is another possibility of course: the AI spend dramatically slows as Wall Street and the AI funders realise they simply can't get a return. How likely that is remains uncertain to us.

The AI effect on our own stocks

With all those caveats about what we do and do not know about AI, we should discuss our own portfolio.

We will start with the stock that worries us most and is nearest the centre of the heat: Alphabet/Google.

To explain the Google's challenge, we offer a cautionary tale about management failures that destroyed a pillar of the Australian financial system — and explain how Google has so far avoided that trap.

The decline of Australians Mutual Provident (The AMP)

Five large institutions dominated the Australian financial system in the mid-1980s: the four big banks and The AMP society. The government refused to allow any of these institutions to merge; it was called the “five pillars policy”.

Of these, AMP was the most powerful institution.

This strikes young Australians as ridiculous, because the banks are still enormous and younger Australians think of AMP as a basket case (if they have heard of it at all).

Go back 40 years, though, and AMP was a giant.

AMP had a network of tied financial planners that dominated the Australian savings/investment market: most Australians would go to these financial planners and be offered a menu of AMP funds. There was an AMP-branded Australian equity fund, a foreign equity fund, a bond fund, and maybe some infrastructure assets — all produced by AMP.

The AMP's positions in many Australian companies were so large that AMP management were the most powerful people in the country. They could ratify or destabilise the board of almost any company.

Outside of AMP, there were small networks of financial planners associated with then-smaller life insurance companies. (MLC was one.) These financial planning networks opened their platforms — so if you went to an MLC planner you could get an MLC product, but you could also get products from Platinum Asset Management, or Fidelity, or one of hundreds of other providers.

Because they offered a variety of products, the open networks were fundamentally better than closed networks.

So the open networks gained share each and every year. This continued for 40 years, and eventually the AMP network became irrelevant.

AMP management knew that this was coming — and that the right thing to do was to open up the AMP network. But the moment they did so, 30-50% of the money would leave to other providers.

No AMP CEO wanted to make the decision that would halve profits so they all deferred the decision. For 40 years. Until they were irrelevant.

Bad short-term incentives led to atrocious long-term outcomes.

Google faces a similar threat — but its management is superior to AMP's.

Google's problem

Google has a search business that produces vast profits by displaying ads. It is threatened by AI. There is a *possible* future emerging where people query computers via AI, bypassing the search engine. There is a bleaker (for Google) future where people interact with computers as you might in [*Star Trek: The Next Generation*](#): you ask the computer something with your voice, and the computer gives a nuanced answer and acts as your agent.

In a voice-controlled AI world there is no space for advertisements — and no place for traditional Google Search.

Google was not going to sit on their thumbs and milk their old business until it died like AMP did. Google leant into the disruption. If someone is going to disrupt your business it might as well be you.

It is hard to express just how forward-thinking Google's decisions were. But let's lay a few out:

In voice search

- In 2007, Google launched a toll-free number you could call that allowed them to start building a human voice search database for training software.
- In 2008 they launched actual voice search, first for the iPhone and later Android
- In 2012 they introduced voice features to partially control android phones (things like settling alarms)

In AI hardware

- Google started chip design as early as 2007
- By 2016 they had designed and built working tensor processing units (TPUs) that made them a genuine competitor to Nvidia.
- By 2021 they had built large numbers of TPUs that they could start putting in data centres.
- Today, Google is selling TPU capacity to Anthropic. The latest generation of TPUs has two versions, one optimised for "training" AI models and the other for using AI models ("inference").

In AI models and software

- In January 2014 (more than twelve years ago) Google acquired DeepMind. This was really a deal to acquire a team: Demis Hassabis, Shane Legg, and Mustafa Suleyman and about 50 to 75 researchers specializing in machine learning.
 - Hassabis remains the key developer of AI at Alphabet. Shane Legg remains at Alphabet, but Mustafa Suleyman is the head of AI at Microsoft.

- Alas a few other key staff have defected to Anthropic.
- In 2016 DeepMind took on Lee Sedol, the world's best Go player. No computer before DeepMind played Go to a level beyond a reasonable club player.
 - This was an astonishing feat. People expected it would be decades before Go computers got to "9 Dan" level.
 - There is a wonderful documentary on this which you can find [on YouTube](#). Much of what became commonplace knowledge about AI is hinted at in that documentary, which John watched at the time — though he wasn't smart enough to recognise its significance.
 - This documentary skirts anthropomorphizing the AI. Lee Sedol concludes that the program is "creative".
- In late 2017, almost as an afterthought, DeepMind released its famous chess engine, AlphaZero. This was a simple AI program that trained itself by playing millions of games against itself for just four hours, learning what does and does not work in chess.
 - Alphazero crushed Stockfish - the seemingly unbeatable and dominant chess engine of the time. The phrase "Stockfish resigns" was uttered with amazement.
 - AlphaZero played chess differently to other chess engines: it was aggressive, and the moves looked more "human" than the (dull) positional play of other chess engines. These games are fun to watch. [In the words of Gotham Chess](#), "AlphaZero will sacrifice material for time, it will sacrifice material for damage, and it will sacrifice material to take your soul".
 - Google also purchased a substantial US based AI Lab which eventually became Google Brain. This was merged with DeepMind. This produced the seminal Transformers idea (2017) which massively improved AI language models.
- There are now many Google AI products, led by Gemini and NotebookLM. If you have not used any, just drop a large number of complicated documents in NotebookLM and ask it to create a podcast explaining the issues to you. It is very good. For instance, it will give you a decent explanation in English of Japanese technical manuals.

More to the point, Google's stated goal from inception was to "organize the world's information and make it universally accessible and useful". Now their stated goal is to take that information and make the world's greatest superintelligence.

This is an amazing pivot. It is the opposite of what AMP did.

We do not think we have ever seen a better managed company.

[But the existential threat is still there](#)

Google has dealt with the existential threat brilliantly, by pivoting to meet it head-on.

Alas, the threat remains. John's son is a computer science graduate and more sophisticated than the average person. He hardly uses Google search any more. His default search is Claude, and a typical question might be to work out different ways of getting from Lyon to Geneva (on holidays) when he needs to be there around 6pm tomorrow. This is much faster than Google search, produces a nuanced answer (giving options at various price points), and denies Google a high-value revenue search.

How much is Google spending?

Google's spending is truly breathtaking: they announced \$180-190 billion for 2026 and guided for a substantial increase in 2027.

The 2026 spend exceeds their cumulative profits over the entire 2010s. Their total AI spend will be larger than their cumulative profits throughout history.

The whole application of Google's achievements, at least in a financial sense, has been to fund a push into AI. Google has even raised cash in 2026.

We repeat: one of the most cash-generative businesses in history has had to raise cash.

What do we think about this? A summary

This presents us with a difficult decision. We love Google as a company and think its management team may be amongst the best assembled in the history of capitalism.

But the stock is abnormally risky. The core business is being disrupted, and that disruption may be total.

The company's reaction is to spend its entire accumulated profits on disrupting itself with no certainty about the outcome, or even whether there is any profit at the end.

We keep trimming the stock. And we know in the end we will be squeamish about the outcome. At the one end Google is *the* provider of the world's finest superintelligence, and is the central company in displacing a good slice of the world's labor (while privatising the profits). At the other end, the company will have its core business destroyed and will have wasted all of the money it made during the good years.

We own it. We genuinely have no idea what it is worth. If somebody tells you that they do, they are pretending to know more about the future than they do about the past or present.

State of the market

For years we have lived in a market dominated by "the Magnificent Seven": Google, Apple, Amazon, Facebook/Meta, Microsoft, Nvidia and Tesla.

That era appears to be over.

The Mag 7 have been mixed this year: Google is up 14% year to date, and Apple is up over 6%, but Facebook/Meta is down, and Microsoft has dropped nearly 20%.

Leadership, however, has been ceded. Google and the others are buyers of AI equipment. The companies that have performed best are sellers of AI equipment, above all the memory (chip) companies. Year to date Micron has tripled. Sandisk is up 850%.

These companies' profits are staggering.

Annualising their most recent quarterly filings, the three major memory companies, Samsung, SK Hynix and Micron, have combined operating income of over US\$380 billion. That exceeds last year's combined operating income for Microsoft, Amazon and Alphabet.

These three memory companies, loss-making in 2023, are now amongst the most profitable companies in history.

An accounting aside: are the earnings of the S&P misstated?

S&P per-share earnings were \$243 and \$271 in 2024 and 2025 respectively. Yardeni Research (an admittedly bullish source) is estimating \$331 and \$375 for 2026 and 2027. These numbers almost always come in lower than estimates as companies on average miss, but the growth rate appears approximately right.

Part of this is an accounting quirk. The profits of these memory sellers are real. The cashflow out of the hyperscalers and into the AI equipment sellers is real. But there is a timing mismatch: cash inflow is recognised as revenue immediately, while the cash outflow is capitalised and depreciated over time.

In Q1 2026, Meta, Alphabet, Microsoft, Amazon, and Oracle spent over \$100B more in CapEx than they recognised in depreciation. If CapEx projections hold true, this difference is likely to exceed \$500B in 2026.

This accounting quirk covers almost a quarter of S&P earnings. It only matters if the CapEx fails to generate a decent return.

Much of the investing world comes down to the same question: what will the returns of this (huge) spending spree be and to whom will it accrue?

Our positions in the AI chain

Alas, we have not owned the memory companies. Had we owned them your results would be much better. But we are long many small positions in the AI chain (mostly small-cap stocks in Japan). Many of these have been multi-baggers. Each individual position is small but the aggregate has been a large win.

We have been trimming these positions. The best have run far harder than Google/Alphabet but they scare us less than Google/Alphabet. Each individually is small, and the threat to their business is not existential. The threat to Google's core search business is existential.

Shorts not claiming to be AI stocks

Lots of companies now claim to be AI stocks. If you make concrete, you are an AI stock — because concrete usage in data centres is way up.

Sometimes the market even believes it.

But there are plenty of stocks that are not AI-adjacent, and junk stocks that are not "semiconductor-adjacent" are declining.

In June we made money on both longs and shorts. It has been a while since that has happened. There were many days on which we made money on both sides of the book (and alas a few days on which we lost money on both sides of the book).

At least recently this has become a good market for us. Our high-quality, non-AI-supply chain stocks are hanging in there, not shooting the lights out, but okay. And our junky shorts (provided they are nowhere near the heat of the bubble) are going down.

We have lost some money shorting junky stocks that purport to be AI adjacent. That will always happen: promoters will always attach themselves to the boom. But the hot part of this market is narrowing. And that suits us.

The next few years

Over the next few years, we will all have super-intelligent “people” we can slot into our business, and our competitors will also have super-intelligent “people” slotted into their business.

You should wonder how it changes your world. And we should wonder how it changes ours. For the moment, though, it is *increasing our opportunities*.

The scammers are getting better and they are using AI. This should be good for a short-seller.

Like all of this, we can’t be sure. But we will be ready.

Yours Sincerely,

The Bronte Team