Ellerston Global Mid Small Cap Fund

Ellerston Capital

Monthly Report as at 31 January 2025

APIR Code: ECL3306AU, ECL8388AU | ARSN 609 725 868

Concentrated portfolio of global mid small cap securities, built through a contrarian, high conviction, and benchmark independent approach.

Targets companies which the Portfolio Manager feels are in a period of "price discovery" and which offer an attractive risk/reward dynamic.

Aims to outperform the benchmark with a focus on risk management and capital growth.

Performance Summary - Class A

Performance	1 Month	3 Months	1 Year	3 Years (p.a.)	5 Years (p.a.)	Since Inception (p.a.)^^
Net^	1.00%	5.06%	14.91%	3.02%	7.68%	10.90%
Benchmark*	3.40%	9.33%	23.68%	9.50%	9.75%	11.29%
Alpha	-2.40%	-4.27%	-8.77%	-6.48%	-2.07%	-0.39%

Performance Summary - Class B

Performance	1 Month	3 Months	6 Months	1 Year	3 Years (p.a.)	Since Inception (p.a.)^^
Net^	1.00%	5.06%	6.88%	14.84%	3.13%	9.02%
Benchmark*	3.40%	9.33%	11.61%	23.68%	9.50%	12.86%
Alpha	-2.40%	-4.27%	-4.73%	-8.84%	-6.37%	-3.84%

*^ The net return figure is calculated after fees & expenses, assuming all distributions are reinvested. Past performance is not a reliable indication of future performance. * MSCI World Mid Cap Index NR (AUD). ^^Class A inception is 28 February 2017. Class B inception is 18 August 2020.*

Key Information

Portfolio Manager(s)	Bill Pridham & Nick Markiewicz
Investment Objective	To outperform the benchmark by 3% over a 5-year rolling period.
Benchmark	MSCI World Mid Cap Index NR (AUD)
Liquidity	Daily
Target Number of Holdings	20-40
Number of Holdings at Month End	33
Minimum Investment	Initial investment - $25,000 Additional investment - $10,000
Distribution Frequency	Half-Yearly (where available)
Management Fee	0.75% p.a.
Performance Fee[1]	10.00%
Buy/Sell Spread	0.25% / 0.25%
Class A Unit Prices & Fund Size	Application – 1.5981 Net Asset Value – 1.5941 Redemption –1.5901 Fund Size – $45,219,800
Class B Unit Prices & Fund Size	Net Asset Value –1.3509 Redemption –1.3475 Fund Size – $48,284,999

[1]Of the investment return above the benchmark, after recovering any underperformance in past periods.

Growth of $10,000 Investment

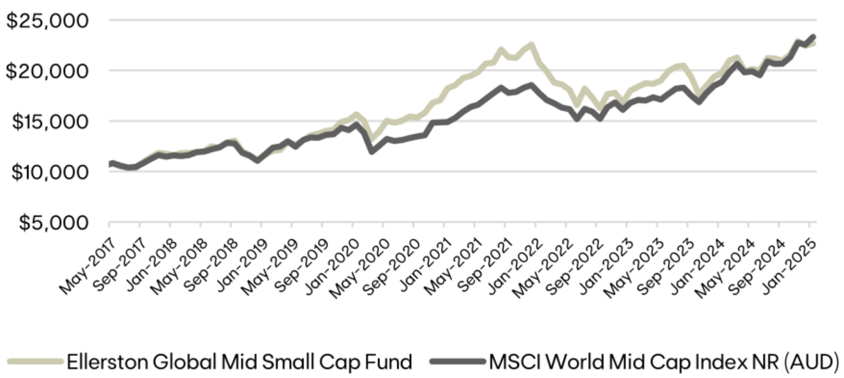

Source: Ellerston Capital.
Performance shown are for Class A Units and net of fees, assuming all distributions are reinvested.
Past performance is not a reliable indication of future performance.

Sector Allocation

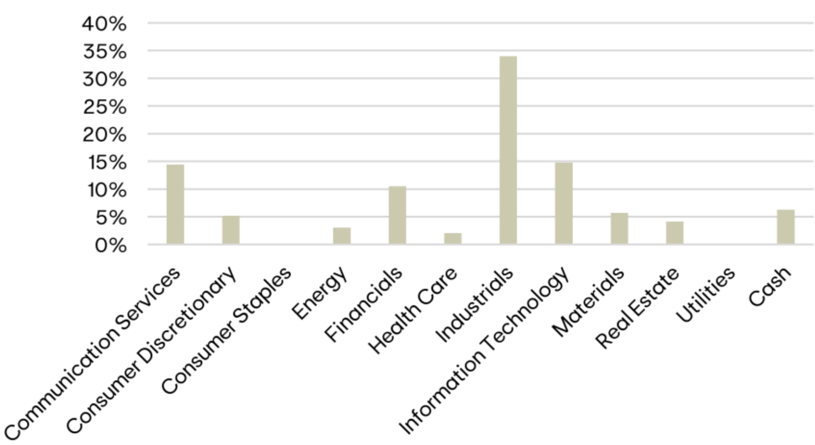

Market Capitalisation Exposure

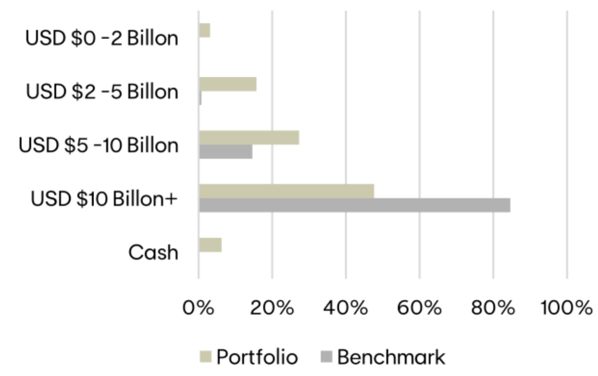

Source: Ellerston Capital.

Regional Exposure

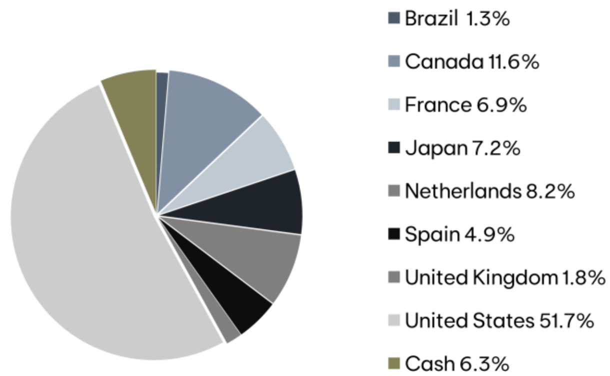

Source: Ellerston Capital.

Top Holdings

Company	Sector	Country	Portfolio Weight
GFL Environmental Inc	Industrials	United States	6.1%
TKO Group Holdings, Inc. Class A	Communication Services	United States	5.8%
Corpay, Inc.	Financials	United States	5.3%
AerCap Holdings NV	Financials	United States	5.0%
Cellnex Telecom S.A.	Communication Services	Spain	4.9%
GXO Logistics Inc	Industrials	United States	4.3%
Bureau Veritas SA	Industrials	France	4.1%
Core Scientific Inc	Information Technology	United States	3.7%
Galaxy Digital Holdings Ltd.	Financials	Canada	3.4.%
Nebius Group N.V. Class A	Information Technology	United States	3.3%

Source: Ellerston Capital

MARKET COMMENTARY

The month was capped by the launch of DeepSeek R1, an Artificial Intelligence (AI) model built by a team from a relatively unknown Chinese hedge fund called 'High-Flyer'. The model's new reasoning techniques and purported efficiency gains in compute sent shockwaves through the AI community, and subsequently, through prices of AI related equities. This event impacted our portfolio performance given exposures to AI 'picks and shovels' plays such as data centres and power.

The Deepseek R1 model grabbed the market's attention on account of being the first open-sourced reasoning model ever released, its high performance (matching some existing frontier models on certain outputs) and extremely low cost (high efficiency) – some ~90% cheaper than existing models to train and run queries.

We will discuss the validity and implications of this further below.

Leading figures in the AI community were uniformly excited by DeepSeek's achievements, hailing their reasoning techniques and open sourced nature as a genuine advancement in the field. The model was also a hit with users, with Deepseek shooting to #1 on Apple's AppStore, dethroning ChatGPT, and accounting for 23% of ChatGPT active users overnight – something other AI models haven't achieved. In response, ChatGPT CEO Sam Altman, somewhat defensively, scrambled to release new (also highly praised) features on ChatGPT in the subsequent days.

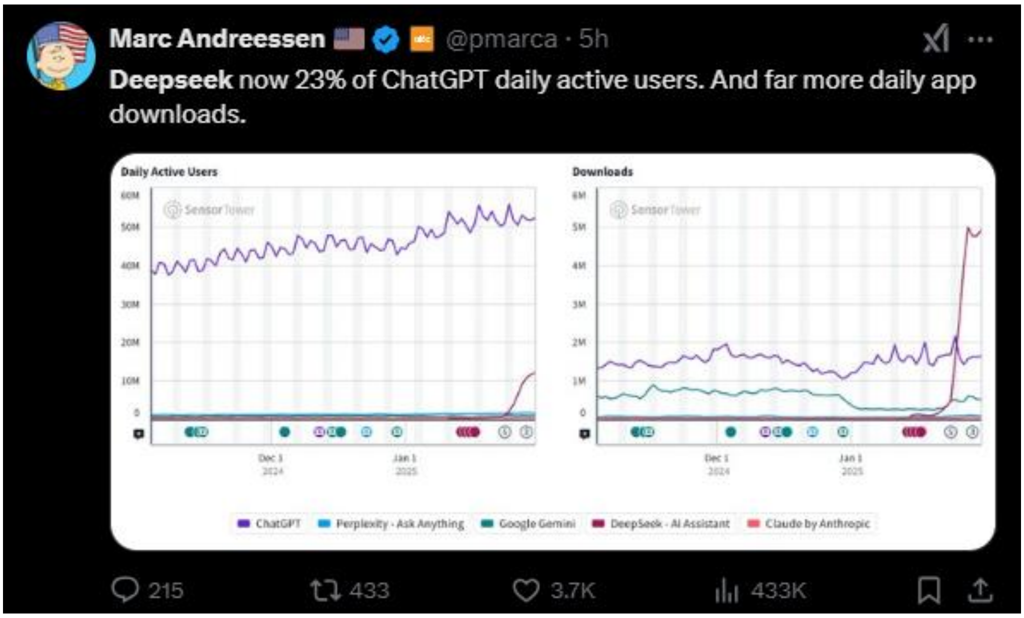

Source: X.com

Why was it so impactful?

Following DeepSeek's R1 release, AI related equities fell sharply lower, with Nvidia, the world's second most valuable company by capitalisation, falling 17% in a single trading session, wiping nearly USD$600b in value, the most ever for a company. Other AI related equities had similar, if not bigger falls.

In our view, stocks fell for many reasons:

Firstly, there was the raw shock that an unknown Chinese company with limited resources (fewer and older GPUs) had released a model that seemingly matched (or at least gotten close to) the outputs and performance of the world's best Large Language Models (LLMs) developed largely by the US tech giants, raising questions around the $100s of billions of dollars being spent to build, train and run these models.

Secondly, the purported efficiency gains of the DeepSeek model have raised questions around the future need for compute power, with a basic interpretation being that a 90-95% improvement in inference and training efficiency may result in significantly lower compute needs in the future, sharply reducing required capital expenditure (very bad news for semiconductors, data centres, hardware and power suppliers).

Finally, from a more technical standpoint, the sell-off and subsequent volatility was a natural outcome of the market's positioning, with the investors holding large absolute amounts of the magnificent-7 (a function of their record weights in the index) as well as hedge fund crowding in the momentum factor (the best 12 month factor by a very wide margin), which comprises many AI names.

While past minor setbacks to the AI narrative have resulted in stocks snapping back quickly, most AI equities are still yet to reclaim their 'pre-DeepSeek' levels. This is due to the still wide range of views from pundits as to the direction of AI more broadly in a 'post DeepSeek' world. An added challenge in assessing DeepSeek is the technical nature of artificial intelligence itself, making it hard for traditional investors to express a confident view given AI sits at the forefront of math/physics, and is evolving at breakneck speed.

What are the key debates?

We are definitely not experts in Artificial Intelligence, though based on our still evolving understanding (having spoken with experts), it would appear that DeepSeek's achievements, while significant, are more likely 'evolutionary' for the industry rather than 'revolutionary'. Furthermore, if one is constructive on the path of AI applications (ie real world use cases), the DeepSeek news could be broadly positive for the development and acceleration of AI.

Costs: DeepSeek excited markets by sharply reducing the cost to both train and run queries on AI models. On the inferencing (query) side, while the DeepSeek model has a much lower advertised price per token than peers (the price to process a unit of text/data), the model conversely also has very slow response times/output, which increases the required GPU infrastructure (i.e cost) to achieve output performance similar to other models. Said differently, the cost of actually operating the DeepSeek model for real time applications is potentially much higher than stated – though still likely competitive in absolute terms vs other models. Reflecting this, DeepSeek has had to limit queries for users, perhaps because it is losing too much money per query, but more likely because it doesn't have the necessary GPU infrastructure to support its model - a core reason that hyper scalers are spending $100's of billions of dollars on GPUs and datacentre infrastructure!

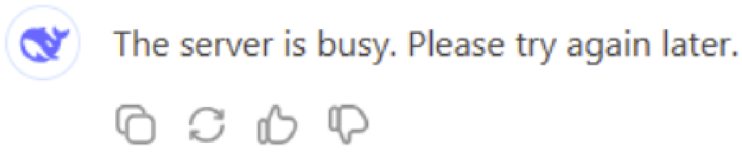

Source: Deep Seek

Model architecture: DeepSeek has been credited as the first open-source reasoning model to use new methods such as 'Mixture of Experts' architecture, which only use small parts of the parameter set of a model to answer to a query, thus saving considerable compute. This is considered a breakthrough of sorts and has AI experts excited about its use and applications. These techniques however were already being experimented with and used by non-open source AI models, though it remains to be seen how far advanced they are. One area that DeepSeek is getting praise for is that it has open sourced these techniques, meaning that they are likely to be copied and integrated into other models very quickly, continuing the current pace of declines in compute cost (GPT-4 token costs fell 150x over 18 months from early 2023 to late 2024).

Model performance: DeepSeek's outputs are similar (though still not superior) to other frontier LLMs. This is a wonderful achievement for a bootstrapped product. However, it should be noted that Deepseek is a distilled model, which itself was trained off other Open-source models. As a result, logically, a model that is distilled from other existing models, is unlikely to exceed the original model's outputs. This would suggest that the large GPU clusters and training time required for LLMs is still a fundamental requirement to push the boundaries of AI, with DeepSeek having not changed this. However, DeepSeek has perhaps raised questions around the long-term 'moats' around some LLMs if they can be distilled and improved upon quickly.

Compute demand: The major fear factor for AI capex is the suggestion that more efficient AI models will result in less compute demand. This has not been the experience in the AI industry thus far, and the long term history in other analogous fields (data, semiconductors) suggests that abundant, cheap compute will likely spur more demand – something called 'Jevon's Paradox'. Growing demand for AI compute appears to have already been happening, even before DeepSeek, with a proliferation of AI applications breaking records for the shortest time for a SAAS company to reach $100m in ARR – this includes a business called Cursor – which has no marketing and team of fewer than 20 people, and has grown from $1m to $100m ARR in 12 months! Feedback from other AI startups and hyperscalers is that they still remain compute constrained, with demand greatly exceeding supply. This is without the broad-based adoption of these applications, nor advances in agentic AI, which could spur exponential demand again.

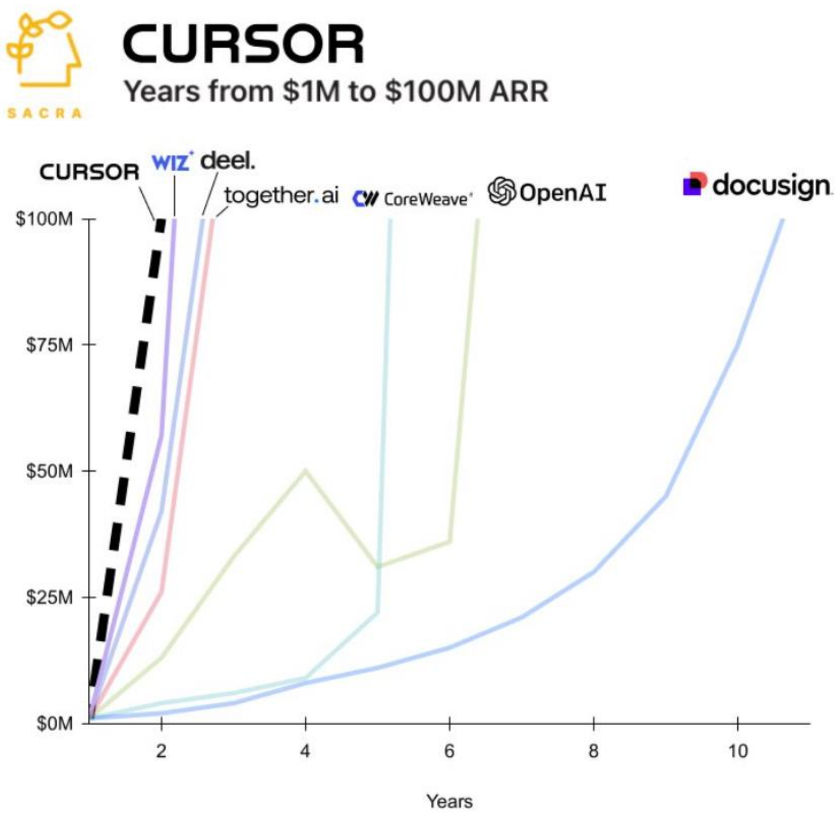

Source: Sacra research

Rather than rely on our opinions, or perhaps superficial understanding of artificial intelligence, we are looking at what the leading AI businesses (the magnificent-7) are doing with their money. Over the last 2 weeks as they have reported 4Q24 earnings and offered 2025 guidance for the first time, instead of cutting capex plans, they have been accelerating them – again, this is after the DeepSeek news. The table below doesn't even include the ~$500b announcement from Stargate, announced just three weeks ago.

Company	FY24 Actual Capex (B$)	FY25 Capex Guidance (B$)	YoY Growth Rate (%)
Microsoft	55.7	80	43
Meta	39.23	62.5	59
Amazon	78	100	28
Alphabet	52.5	75	42

Source: Deepseek

Finally, DeepSeek's own assessment is also instructive – they have demonstrated that smaller distilled models produce excellent results (which is good for AI inference demand more broadly), and importantly, that more powerful base models (that use more power/GPUs) are still required to push the boundaries of AI.

> Therefore, we can draw two conclusions: First, distilling more powerful models into smaller ones yields excellent results, whereas smaller models relying on the large-scale RL mentioned in this paper require enormous computational power and may not even achieve the performance of distillation. Second, while distillation strategies are both economical and effective, advancing beyond the boundaries of intelligence may still require more powerful base models and larger-scale reinforcement learning.

Source: Deepseek

Putting everything together, we remain constructive on the AI outlook, and largely see recent movements as possibly an overreaction. Our core AI holdings own real assets, are largely agnostic to whichever LLM model leads the race and are levered to the bottlenecks in the AI industry, i.e power and data centre capacity – factors that we believe are likely to continue to remain constrained. Furthermore, they remain modestly priced and continue to offer strong upside relative to their underlying risk.

PORTFOLIO COMMENTARY

The Ellerston Global Mid Small Cap Fund returned 1.00% net during the month, compared to the MSCI World Mid Cap (AUD) Index, which returned 3.40% over the same period.

The portfolio's top contributors Aritzia, TKO and Corpay added 190 bps to performance while Core Scientific, QXO and Nexans detracted 103 bps during the month.

Aritzia (+0.79%) is a consumer apparel business offering everyday luxury items to customers in over 180 countries both online and through more than 125 boutiques across North America. The stock price rose in mid-January after the company reported another strong quarter, achieving record earnings.

TKO (+0.56%) is a sports entertainment company, with key assets being the UFC, the world's premier mixed martial arts organization, and WWE, the global leader in sports entertainment. The stock rose following the strong launch of WWE on Netflix, coupled with a high chance that Netflix may bid for the upcoming UFC rights alongside ESPN.

Corpay (+0.55%) provides commercial cards (e.g., business cards, fleet cards, virtual cards) and accounts payable solutions (e.g., invoice digitisation, payments automation, cross-border payments) to businesses worldwide. The company currently has strong momentum, and is likely to be a primary beneficiary of strong nominal US GDP growth.

Core Scientific (-0.56%) is a leader in digital infrastructure for bitcoin mining and high-performance computing. A significant portion of its eight operational data centres is now allocated to supporting AI-related workloads. The stock price fell following the DeepSeek news and concerns around future demand for compute. We remain positive about global data centre demand and believe Core Scientific remains well placed to benefit from this, with new contracts likely to be announced in the coming months.

QXO (-0.25%) is a listed cashbox that CEO Brad Jacobs will use to acquire building products distributors. The company made an all-cash offer to acquire outstanding shares of Beacon Roofing Supply, one of the largest distributors of roofing, waterproofing, and related exterior products in North America. QXO's share price declined after Beacon's Board took actions to block the offer, suggesting a protracted and potentially more expensive battle for control.

Nexans (-0.22%): Nexans is a French company specializing in advanced cable and connectivity solutions, designing and installing cables that support the renewable energy transition and electrification more broadly. The stock fell in sympathy with other 'electrification' plays after Deepseek. Nexans secured a new €1 billion project for the offshore grid connection program in the German North Sea.

In January, we initiated a new position in Kodiak Gas Services, the largest contract compression services provider in the United States, playing a critical role in the infrastructure that enables the safe and reliable production and transportation of natural gas and oil.

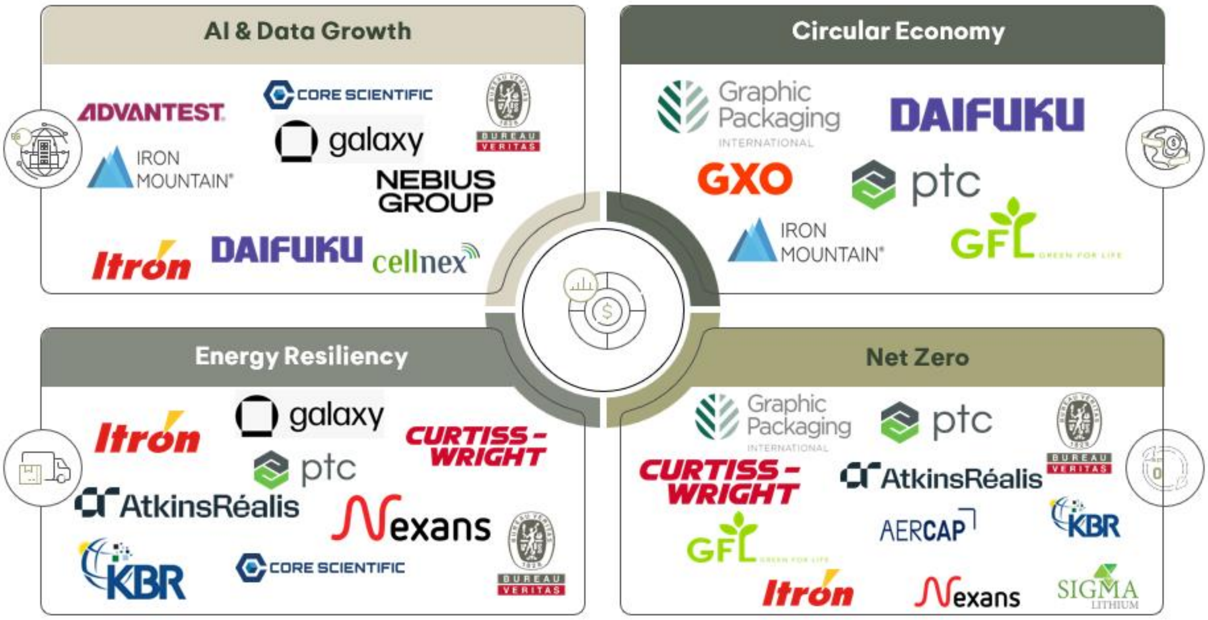

Source: Ellerston Capital

These businesses as well as idiosyncratic opportunities in the Fund should provide solid absolute and relative returns over the long term as secular and structural business drivers help mitigate earnings risk in times of economic uncertainty. As always, we thank you for your continued support and look forward to providing further updates in the future.

Regulatory Guide (RG240) Fund Disclosure Benchmark – Periodic Reporting (monthly)

- **Net Asset Value of the Fund and Redemption Price of Units**

Please refer to details on page one.

- **Any changes to key service providers including any change in related party status.**

There have been no changes to key service providers, including any change in related party status.

- **Net returns after fees, costs and relevant taxes**

Please refer to details on page one.

- **Any material changes to the Fund's risk profile and strategy**

There have been no changes to the Fund's risk profile and strategy.

- **Any material changes related to the primary investment personnel responsible for managing the Fund.**

There have been no changes to the primary investment personnel responsible for managing the Fund.

Contact Us

Sydney

Level 11, 179 Elizabeth Street,
Sydney, NSW 2000
+612 9021 7701
info@ellerstoncapital.com

Find out more:

Should investors have any questions or queries regarding the Fund, please contact our Investor Relations team on 02 9021 7701 or info@ellerstoncapital.com or visit us at ellerstoncapital.com.

All holding enquiries should be directed to our registry, Automic Group on 1300 101 595 or ellerstonfunds@automicgroup.com.au.

Ellerston Capital