

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

20th July, 2026

Dear Partners,

The Longriver Partners Fund's net return this quarter was 8.5%. We are flat for the year to date and up 60.5% since inception, equivalent to a 14.5% annualised compound return.

Our investment strategy is to ride the long-term value created by our companies, so I encourage you to take a similarly long-term view of your investment in the Fund. Consider it as you would an investment in real estate or a private company, assets for which you would not receive a volatile quarterly valuation.

FROM TOKENS TO WORK

I never planned to write an annual update on artificial intelligence, but here I am again.

Two years ago, I asked whether the customers would come and where the value of AI would accrue. Big Tech had embarked on one of the largest private investment programmes in history, but the commercial evidence was thin. "If you build it, will they come?" was the question ([link](#)).

Last year, I wrote that AI was moving from hype to habit. It had become central to my own work, as better models and better tools made it useful for research, writing and thought ([link](#)).

The important change this year is the arrival of coding agents, and agents more generally. They are semi-autonomous, persistent systems which can plan and execute work. Rather than merely returning answers like a chatbot, agents get things done.

The question is no longer whether people will talk to a chatbot, or how AI can augment the way we work. It is how work can be rebuilt around a theoretically infinite number of agents working continuously and in parallel. That has the potential to radically increase AI's economic value and the magnitude of demand.

It also raises harder questions. Who will control the work? Who will get paid? And will the massive infrastructure build-out required to support this demand earn an adequate return?

The answers will shape the future of business, both large and small, including those in our portfolio, such as Amazon, KLA, Meta, Nvidia and TSMC.

MY SUMMER OF CODEX

'Tech debt' is the hidden cost of not maintaining systems. Paying it down is disruptive and unglamorous, but it becomes more difficult the longer it's deferred. Put it this way: if you had to fix an aeroplane's engines mid-flight, would you rather a small, regular tune-up while things were going well, or a complete overhaul when things were going wrong?

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

Longriver is a solo practice and time is my biggest constraint. Time paying down tech debt is time not spent elsewhere. I have lots of ideas to improve the way I work, but they've often been unfeasible because they would have required an engineer, an external developer or months of my own time.

OpenAI's Codex coding agent changed that calculation. I had been steadily chipping away at my tech debt whenever I could snatch the time. Over the last few weeks, I took a sledgehammer to it.

What began as 'vibe-coding' (describing what I wanted in plain English and watching software appear) has become a functioning, cloud-based, AI-integrated research and knowledge management system tailored exactly to how I think and invest.

It is vertical market software for a market of one.

Building it was an education, too. 'Vibe-coding' sounds whimsically easy. But I soon discovered that I'd built little more than a sprawling, unstructured, brittle and ultimately unusable prototype. AI could produce code faster than I could judge whether it was any good. And it offered no guidance on product design. It just kept saying yes.

It was only after a number of expensive dead ends that I started again and learned through trial and error how to use the tools properly. I now plan before I begin, demand to know the trade-offs, scope work within tight constraints, and include processes to check the results. I get much more done for every dollar I spend.

It may be counterintuitive, but that actually makes me want to spend more, not less. As I've become more productive, more projects have become worth attempting.

Consider this. I asked AI how much it would cost to add a feature to my research system. It replied: *"A production-safe implementation is more likely one to two engineering days. In Codex terms, I would expect roughly two focused prompts, perhaps 50,000–100,000 tokens."*

In other words, it could do the job for a few dollars and less than an hour of my time.

Work that would previously have been too expensive to contemplate has become cheap enough to attempt almost casually.

In fact, the returns are so high at this stage that it almost feels irresponsible not to do more. This is why you hear stories of people walking around with open laptops so they don't interrupt their agents.

I can empathise. I started working one night on an improvement to my retrieval system, and before I knew it, it was 04:30 and I was still happily running A/B tests.

People fancy that AI will help them work less. My experience has been almost the opposite. Whether writing code or researching companies, I have been working harder than ever.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

WHO CONTROLS THE WORK?

My experience points to a broader change in AI demand. When a chatbot gives you an answer, consumption is constrained by the number of questions you can think to ask. When an agent does work, it is constrained by the number of jobs you can find for it. That is a much larger number, at least where the work can be specified and checked cheaply enough.

Work also changes how AI can be priced. Tokens are not equal. A cheap model is not cheap if it misunderstands the task, requires repeated prompting or creates errors that take longer to find than the original work itself.

Nor are two identical models equally useful if one has access to better tools and richer context. This is called the model's 'harness'. The harness allows the model to inspect files, call other software, test its output and recover from errors. The context includes history, permissions, data, rules and previous decisions. The model supplies general intelligence. The harness makes that intelligence useful. Context makes it useful to you.

The relevant measure, therefore, is not the cost of a token but the total cost of successfully completing a job.

That gives OpenAI and Anthropic, the leading frontier model companies, a possible route to pricing power among enterprise customers that they never found among consumers. A casual question about where to go on holiday probably isn't worth paying for. But a superior model and harness that save a senior employee a week can be cheap, even if they're priced higher than inferior alternatives.

I learned that, too, when a friendly notification popped up in Codex informing me that if I'd used a faster model, I could have saved eight hours over the past week. At one and a half times the base price, the faster model surely would have been worth it. The question is whether the model company can retain enough of that value.

Agents are not yet general-purpose substitutes for labour. They work best where tasks are digital, tool-based and verifiable. Coding is the clearest proof point because it gives AI a scoreboard. The software works, or it does not. Tests pass or fail. A bug is fixed or remains. The agent can observe the result of its actions and adjust its next attempt.

The same pattern can spread into many other forms of knowledge work, though the scoreboard becomes less clean. A legal document can be checked against precedents. A financial reconciliation can be compared with source data. A customer ticket can be judged against the eventual outcome. Research can be retrieved, challenged and updated against defined standards.

The commercial evidence suggests customers are willing to pay for this. Revenue growth at OpenAI and Anthropic has accelerated sharply. Anthropic, founded in 2021, has reportedly seen its annualised revenue run-rate increase from USD 1 billion at the end of 2024 to USD 9 billion at the end of 2025 to almost USD 50 billion in May this year. OpenAI's trajectory is similar. OpenAI reported that by June this year, more than five million people were using Codex each week, more than six times as many as when its desktop application launched in February.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

That proves demand, not durable economics. Value created is not the same as value retained. The model companies must still fund inference, research, distribution, infrastructure and financing, while repeatedly paying to remain at the frontier. There are reports that Anthropic and OpenAI are achieving very high gross margins on inference, but the picture is blurred by other costs. Eventual public disclosure will tell us more.

Competition has not disappeared among models, either. Customers are incentivised to route simpler work to cheaper models to manage their growing spend. Chinese laboratories and open models can already provide much of the useful capability at far lower prices. DeepSeek, Moonshot and MiniMax show how quickly a broad tier of model capabilities can be commoditised, even as frontier and difficult production workloads remain harder to reproduce.

OpenAI and Anthropic recognise this and hope to differentiate themselves by creating a learning advantage. Where users permit it, the more work completed directly in Claude or Codex, the more evidence the model companies have on which approaches succeed, where agents fail and how users correct them.

Those lessons can improve the model and harness. The better the system understands the recurring context of your work, the more likely it is to complete work successfully. And the more successful it is, the more work you are likely to do within it, and the less likely you are to leave. But that only becomes a moat if the model company controls enough of the interface, context, billing and feedback to influence the relationship.

Of course, this puts OpenAI and Anthropic in direct competition with the application layer. Claude or Codex may become the place where work begins, while Microsoft Office, Salesforce or Adobe remains the system of record, execution layer or home of the customer's data. Or those applications may embed multiple models, making the intelligence beneath them interchangeable.

The contest is therefore not simply over which model is best. It is over who controls the work: the interface, tools, context, execution, billing and customer relationship. And that contest does not stop with models and software. It is spreading across the entire stack.

EVERYONE IS COMPETING WITH EVERYBODY, IN EVERYTHING

Let me step back for a moment. I began investing in technology a decade ago when I realised that certain companies offered something the market failed to value: profitable growth with unusual certainty thanks to network effects, bundling or economies of scale. The competitive logic was simple: own an indispensable control point and make everything around it interchangeable.

Until recently, the US Big Tech companies' most profitable control points remained largely distinct. Competition existed, but rarely threatened all of them at once: Apple did smartphones; Alphabet did search; Amazon did retail; Meta did entertainment; Microsoft did enterprise. They may have overlapped in providing cloud services, but the market was big enough for everyone. And their profits could sustain the ecosystem around them.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

Artificial intelligence complicates that framework. It is a general-purpose capability that allows companies to attack one another's control points. It has also created two new competitors, OpenAI and Anthropic, with the scale and capital necessary to go after other layers of the stack.

Suddenly, it is beginning to feel like everyone is competing with everybody, in everything.

OpenAI and Anthropic want to own how we work with artificial intelligence. But they also want freedom from the companies supplying their compute. Anthropic, for example, is working across cloud providers and chip architectures. It wants enough portability to prevent any one cloud or chip supplier from controlling its economics.

The incumbent platforms are responding in kind. They do not necessarily need to own the entire stack. They need to control the customer relationship while making models and chips interchangeable.

Amazon and Microsoft are both trying to own the platform while making the intelligence beneath it interchangeable. Trainium reduces Amazon's dependence on Nvidia, while Bedrock makes OpenAI, Anthropic and other model providers compete inside a common platform. Microsoft is building its own models and infrastructure, while designing Copilot so that one model can be swapped for another. Both want the models and chips to remain suppliers inside their systems, while the workloads, data and customer relationships remain their own.

Alphabet has gone furthest. It owns the chips, cloud, models, consumer distribution and monetisation. This is not just about efficiency. It is intended to prevent Nvidia or OpenAI from inserting themselves between Google and its users, while making every layer of Alphabet's stack reinforce the others.

Meta and Apple are protecting different control points. Meta is investing vertically in models and custom silicon because it cannot allow an outside supplier to control the intelligence behind its feeds, advertising and future devices. Apple may rent frontier intelligence from others, but it intends to keep personal context, permissions and action inside the iPhone. Meta wants to own the learning loop. Apple wants to own the interface.

The suppliers are defending themselves too. Nvidia is funding neo-clouds that can compete with the hyperscalers and publishing open models that give customers alternatives to OpenAI and Anthropic. Enterprise software companies are building their own agents before the model companies can displace them as the place where work begins. Everyone is trying to create more competition for what they buy and less competition for what they sell.

But escaping dependence is expensive. Every layer brought in-house adds capital, complexity and the risk of competing with the partners that helped build the original business. Some spending may create growth and enhance control. Some may be defensive: strategically necessary to avoid surrendering an existing profit pool, but not necessarily capable of creating a new one. A company may preserve its position while losing some of the financial qualities that made it attractive in the first place.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

Thus, the profit pools that attracted me to invest in technology are blurring. Some will migrate to new control points. Others may be competed away entirely. The investment question is no longer simply who will grow, but which layer can still earn excess returns when everyone is competing with everyone else.

AI'S 'CHINA MOMENT'

Regardless of how much can ultimately be earned from AI, accelerating demand is putting pressure on an already strained data centre supply chain.

The first wave of AI infrastructure investment was driven by training. Inference has landed on top of that base. Reasoning models spend more compute producing an answer, while agents repeatedly call models, software and databases to complete a task. Serving intelligence is becoming a continuous industrial workload.

The situation reminds me of when China joined the World Trade Organization in 2001 and entered a period of hypergrowth. Its voracious demand for resources could not be met by a commodities complex built for a slower world. Inelastic demand collided with inelastic supply. Mines, ports and mills could not be built fast enough. Prices surged.

Artificial intelligence is experiencing a similar 'China Moment' in 2026. Fresh from the COVID boom-and-bust, many suppliers cut capacity or invested cautiously, while AI demand accelerated beyond their forecasts. Capacity is now scarce at many critical points, from advanced packaging and memory to power and skilled labour. Prices are rising from the substrates that circuits are built on to the wages of the electricians who ultimately install them.

It also changes what counts as supply. The number that matters is not how many chips, factories or megawatts have been announced. It is how much working computing capacity can be delivered on time, with the memory, networking, power, cooling and software required to use it. A warehouse full of GPUs is useless without the rest of the system.

Rising capex budgets are translating into more working compute, but less quickly than the headline numbers imply. Much of the increase is absorbed by higher prices, duplicated sites and development options, and the cost of coordinating the whole system.

Nvidia anticipated this earlier than most, using its scale and balance sheet to secure capacity at TSMC and supplies of critical components like memory and optical connectivity. Buyers may not want to pay Nvidia's margins, but they may still do so if Nvidia offers the fastest and lowest-risk path to working compute. Supply-chain coordination has become part of Nvidia's product.

The bears argue that time will resolve these bottlenecks as profits encourage new supply. But the bulls point to a new twist: the extreme performance requirements of AI computing systems require extreme engineering. Suppliers that once appeared commoditised are becoming indispensable partners. Their value lies not only in adding capacity but also in solving the thermal, power, bandwidth, yield and reliability problems that determine whether a system works at all.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

Again, Nvidia was ahead of the pack in its evolution from designing individual circuits to complete 'AI Factories' that not only work, but work efficiently. By integrating across chips, networking, memory, software, power and cooling, Nvidia is able to coordinate design to achieve superior performance, avoid commoditisation and capture more of the system's margin.

TSMC's expansion into advanced packaging is a similar effort to squeeze more performance out of the chips it manufactures. Packaging is no longer merely back-end assembly; it is part of the system-scaling problem. The need to design, qualify and coordinate these steps together increases complexity and, therefore, customers' dependence on TSMC to manage them, even if other qualified providers can perform some of the work.

High-bandwidth memory is another example. HBM stacks DRAM to move data to an accelerator much faster than conventional memory, but a usable stack must meet strict heat, power, bandwidth and packaging requirements. It is qualified bandwidth capacity for a particular accelerator, not merely a collection of memory dies. That makes it less interchangeable than ordinary DRAM, but not immune to a capacity cycle.

Networking and optical connectivity provide another example. As AI clusters grow, switches, optical modules, retimers and cables increasingly determine how efficiently expensive accelerators can work together. But technical importance does not guarantee durable pricing power: customers will standardise, dual-source and bargain wherever they can.

The market's strongest performers this year have often sat at the tightest bottlenecks, where rising prices have produced the most operating leverage and earnings upside. But the question now is which advantages survive when supply catches up? A bottleneck can be real and still be a bad investment if Mr Market prices it as permanent.

AI'S CAPITAL CYCLE

The scale of AI investment has become large enough to define Big Tech's future. The build-out is no longer being financed solely from operating cash flow. Leases, purchase commitments, and debt are increasingly involved, and, in Alphabet's case, there is a recent equity raise.

CEOs point to attractive early returns, but they have also described an arms race in which failing to invest may threaten their existing businesses, even if the standalone returns on some of the new capacity prove modest. This is a prisoner's dilemma: each company may act rationally, yet the industry collectively commits too much capital.

At the same time that the AI complex has become more competitive, it has also become narrower and more interdependent. Much of the incremental demand is concentrated in two model companies and a small group of hyperscalers and neo-clouds scrambling to expand capacity. They are buying accelerators and custom silicon from Nvidia, AMD, Broadcom and others. Most leading-edge logic is manufactured by TSMC, while SK hynix, Samsung and Micron supply the memory. These manufacturers are raising their own capex and buying from a concentrated group of equipment suppliers.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

The suppliers are being paid today. But the economic burden of proof still sits with the customer at the far end of the chain.

This leaves me with a paradox.

As someone who uses AI, I should hope that it stays cheap. But as an investor, I should hope that it doesn't. The system ultimately needs to produce enough cash to support its massive investments in training, user acquisition, infrastructure and capacity, or someone must remain willing to subsidise the gap. I need enough competition that no single layer captures too much value, but not so much that it destroys profits for everyone.

Here the other side of the China analogy becomes uncomfortable. China's growth seemed to justify years of breathless investment, but building eventually became detached from the return on what was built. Excess capacity lowered prices and wrecked balance sheets, from the miners to the mills, to the home builders, to the banks and the homeowners.

AI need not follow the same path. The largest investors have profitable businesses, strong balance sheets and demand that exceeds capacity. Some investments are already improving advertising, cloud growth and software revenue. But revenue can arrive well before owner earnings if depreciation understates economic obsolescence, or if leases, purchase commitments and the replacement capital required to remain competitive rise faster than the cash generated.

It is possible that AI will remain a profitable and enduring investment for all parties concerned. It is also possible that AI transforms business and society but does not yield an adequate return to the companies that built it. Just because it is different does not mean that we can ignore the same capital-cycle risk.

WHAT THIS MEANS FOR OUR PORTFOLIO

Our Fund has been invested in TSMC and Meta since 2023, and Nvidia since 2025. I invested in semiconductor equipment supplier KLA Corp earlier this year and reinvested in Amazon, too. Together, as I write this, these investments comprised 35% of our portfolio, having grown larger due to their outperformance relative to other holdings, before our reduced their size.

We are not a technology fund. As I explained earlier, I became attracted to these companies because they offered large and growing profit pools, high returns on capital, strong balance sheets, and earnings whose durability appeared underappreciated.

Amazon is becoming the neutral market for artificial intelligence. Bedrock lets customers use competing models without caring which chips sit underneath, while Trainium can lower Amazon's cost of serving them. If inference becomes the larger market, AWS can grow by selling the abstraction rather than betting on a single model.

KLA supplies the process-control tools that advanced chipmaking cannot do without. As wafers become more valuable and tolerances tighter, inspection intensity rises. Recurring service revenues and modest capital requirements allow much of that growth to reach shareholders.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

Meta is applying AI to one of the world's most productive advertising systems, improving recommendations and advertiser returns across billions of users. Its compute capacity largely serves its own products, so it is not relying on OpenAI or Anthropic to justify the investment.

Nvidia sells working computing systems, not merely chips. Its software, scale and supply-chain execution remain difficult to replace, while inference is broadening demand beyond training.

TSMC is the manufacturing choke point for leading-edge semiconductors. Few rivals can reproduce its yields, scale and execution, while AI provides another growth engine.

The aim is not to own every winner in the AI complex. In fact, we are not obliged to own any company in the AI complex at all. Our goal is to own a small number of businesses that can retain value and reinvest their earnings at high rates across a range of scenarios, and I want to buy them at cheap prices. I have therefore preferred businesses whose role should remain valuable after today's bottleneck moves and the immediate scarcity begins to ease.

I remain comfortable owning our AI-exposed holdings, but less comfortable assuming that their different positions provide meaningful diversification. I have therefore begun reducing their aggregate weight and raising cash, giving us greater flexibility as the capital cycle advances and the evidence becomes clearer.

WHAT REMAINS SCARCE?

Finally, let's spare a moment to ask: what happens to those of us living downstream of AI? Last year, I wrote that AI was turning the fixed cost of knowledge work into a variable one, and that the only sensible response was to lean in. I still think that is right.

But it's not enough. A business can use AI to do more with less, yet retain little of the gain if suppliers keep prices high, competitors adopt the same tools, or customers adopt the new standard as the status quo.

The technology can create value without allowing the company adopting it to capture that value.

I spent my summer investing time and money to improve the way I work at Longriver (and to do my bit to keep the AI complex afloat!). But anyone with internet access, a USD 20 per month subscription to ChatGPT and the initiative could do the same. Will this work help me pull ahead or just keep me in the game?

Warren Buffett reached a similar conclusion about Berkshire's textile mills. A new machine might lower one mill's costs, making its purchase appear rational. But once every competitor installed the same machine, the advantage disappeared. The industry committed more capital, while the benefit flowed to customers through lower prices.

Sport offers the same lesson. A marathon time that won major races a hundred years ago would no longer be competitive at the highest level. Runners have become dramatically faster, but the number of gold medals has not increased.

LONGRIVER PARTNERS FUND

2ND QUARTER 2026

The changing nature of work has also led the market to reassess many enduring franchises. Some businesses may possess genuine control of a workflow. Others may merely have controlled scarcity in software development, information, distribution or labour under the old regime. If AI removes that scarcity, value will shift elsewhere. What is at risk for these companies, therefore, is not only current cash flow. It is terminal value.

The burning question is what remains scarce once competent knowledge work becomes abundant?

To put it in the parlance of our times, what is AI slop and what is not?

For the uninitiated, Google Gemini tells me that “AI slop is low-quality, mass-produced digital content created using generative AI tools like ChatGPT or Midjourney. It often lacks human effort, accuracy, or genuine meaning.”

Let me share a live example. My wife Lizzie is writing a book on a niche area of the law. The ideas are all hers: she framed the questions, checked the authorities and decided which arguments matter. Critically, her name will go on the cover, and it will be her reputation on the line. But AI did the grunt work, reducing the time required for research and drafting from months or years to a few weeks.

Will her book be just another example of AI slop? Unequivocally not.

But what's stopping a thousand other smart lawyers from writing a thousand other books that they have long dreamed of publishing? With AI tools, nothing except their imagination and inertia. And if they do, will quantity overwhelm quality? If one good book is not AI slop, how about a thousand good books? Ten thousand? Or a million?

AI may well deliver abundance.

But are we ready for it?

CLOSING

Thank you for your continued trust and partnership. I do not take it for granted. If you would like to discuss the year, the portfolio, or this note in more detail, please reply and I will make time.

Sincerely,

Graham F. Rhodes
Founder & CIO
Longriver Investment Partners Limited

LONGRIVER PARTNERS FUND

1ST QUARTER 2026

Longriver is an Asia-informed global equity strategy for long-term investors who think like business owners. I invest from Hong Kong in a focused group of public companies that I can understand, trust and buy with a margin of safety. The aim is differentiated exposure to businesses worldwide that can compound over time.

WHAT THE FUND OWNS:

- Global, long-only best ideas equity portfolio.
- Businesses selected for resilience, re-investment potential and stewardship.
- Benchmark-agnostic portfolio construction.
- Purchases made with a margin of safety.

TOP 10 HOLDINGS (ALPHABETICAL):

AMAZON
APPLOVIN
CONSTELLATION SOFTWARE
GAMES WORKSHOP GROUP
HIKARI TSUSHIN
META PLATFORMS
NVIDIA
PLOVER BAY TECHNOLOGIES
TAIWAN SEMICONDUCTOR MANUFACTURING CO.
WISE

PORTFOLIO SNAPSHOT:

TOP 10 CONCENTRATION = 89%

2027E: 24x PE, 52% ROE, 20% EPS GROWTH

Notes: Benchmark is the MSCI AC World USD Net Return Index. Longriver net performance figures are unaudited. Portfolio characteristics are as at 31 March, 2026.

KEY INVESTMENT TERMS:

Manager: Longriver Investment Partners Limited
Mandate: Long-only, global, primarily equities
Launch: January 3rd, 2023
Minimum: USD 100,000
Fees: 1.00% p.a. + 10% over 6% hurdle + High Water Mark
Liquidity: Monthly subscriptions; quarterly redemptions with sixty days written notice


WHAT CLIENTS GET:

- Focus, patience and accountability.
- Fair fees with a hurdle rate.
- Personal capital invested alongside clients.
- Candid communication and direct access to the decision-maker.

NET RETURNS:

Since Inception	Fund	Benchmark	Delta
Cumulative Return	60.5%	95.4%	(34.9%)
Per Annum Return	14.5%	21.1%	(6.6%)

	Q1	Q2	Q3	Q4	YTD	BMK	DELTA
2026	(7.8%)	8.5%			0.0%	11.3%	(11.3%)
2025	(0.3%)	11.7%	5.8%	(2.6%)	14.8%	22.3%	(7.5%)
2024	13.9%	1.3%	6.6%	2.5%	26.0%	17.5%	8.5%
2023	10.1%	3.1%	(10.7%)	9.3%	10.8%	22.2%	(11.4%)



OPERATIONS:

Structure: Hong Kong Open-ended Fund Company
Regulator: Hong Kong Securities & Futures Commission
Admin.: Sinopac Services & Solutions Limited
Custodian: Interactive Brokers Hong Kong Limited
Auditor: Emerald CPA
Legal: Charltons
Contact: grahamfrhodes@longriverinv.com

LONGRIVER PARTNERS FUND

1ST QUARTER 2026

DISCLOSURE:

This document is issued by Longriver Investment Partners Limited, a Licensed Corporation (CE Number: BRV905) regulated by the Hong Kong Securities and Futures Commission ("SFC") to conduct Type 4 (Advising on Securities) and Type 9 (Asset Management) regulated activities in Hong Kong. Our registered address is Room 8, Unit 707, 7/F New East Ocean Centre, 9 Science Museum Road, Tsim Sha Tsui, Hong Kong.

This document is provided for informational purposes only. It does not constitute investment, legal, tax or other advice, a recommendation, an offer or invitation to the public, or an offer to sell, or a solicitation of an offer to subscribe for or purchase, any securities, interests in any collective investment scheme, fund interests, investment management services or investment advisory services in any jurisdiction. This document must not be relied upon as the basis for any investment decision. Any investment may be made only on the basis of the relevant offering memorandum, supplement, subscription documents and constitutional documents, as amended from time to time.

This document is presented on a confidential basis only to recipients who have been classified by Longriver Investment Partners Limited, or in respect of whom Longriver Investment Partners Limited has taken reasonable steps to verify, as Professional Investors as defined in the Securities and Futures Ordinance (Cap. 571) and its subsidiary legislation and/or, where applicable, "accredited investors" within the meaning of Rule 501 of Regulation D under the U.S. Securities Act of 1933, as amended. By accepting this document, each recipient represents that it is such an investor, is not receiving this document as nominee or agent for any person who is not such an investor, and will notify Longriver Investment Partners Limited promptly if that status changes.

This document, and any fund or service referred to in it, is directed only at persons to whom it may lawfully be distributed. No action has been taken to permit a public offering of any fund interest or service in any jurisdiction. The distribution of this document and any investment referred to in it may be restricted by law, and each recipient is responsible for observing all applicable laws and regulations in any relevant jurisdiction.

This document has not been reviewed by the SFC or by any other regulatory authority in Hong Kong or elsewhere. Any fund referred to in this document has not been authorised by the SFC under section 104 of the Securities and Futures Ordinance, and its offering documents have not been authorised by the SFC under section 105 of the Securities and Futures Ordinance. SFC registration, licensing or regulation does not imply any recommendation or endorsement by the SFC, nor does it guarantee the commercial merits, performance or suitability of any fund, product or service.

Past performance is not indicative of future results. The value of investments may rise or fall, and investors may lose a substantial portion or all of their investment. Any forward-looking statements, estimates or opinions are made only as of the date of this document and may change without notice. No representation or warranty is made as to the accuracy, completeness or fairness of the information, and Longriver Investment Partners Limited accepts no obligation to update it except as required by law or regulation.

Any product or service referred to in this document may not be suitable for any or all persons. Recipients should consult their own legal, tax, financial and other professional advisers before making any investment decision. This document is intended exclusively for the use of the intended recipient and must not be reproduced or redistributed to any other person without the prior consent of Longriver Investment Partners Limited, except that a recipient may provide a copy to its professional advisers on a confidential basis. If the reader of this document is not the intended recipient, you are hereby notified that you have received this document in error and that any review, dissemination, distribution or copying of this document is strictly prohibited.