

# Quarterly Manager Views

## FSSA Global Emerging Markets Focus Fund

April 2026

### The Friction Test: What the market gets wrong about AI disruption

*“The future is already here – it’s just not evenly distributed.”*

William Gibson

It has been a difficult period for the fund. At the time of writing, a number of our holdings are being challenged by an increasingly narrow view of where durable value creation will sit in an artificial intelligence (AI)-driven world. On one side sit the obvious beneficiaries of the current AI build-out – semiconductors, memory companies, selected hard assets tied to data-centre investment, parts of defence and other businesses linked to physical capacity. On the other side sit the digital and capital-light companies we tend to prefer, many of which have been treated as though disruption is not only coming, it is likely to impair their long-term economics in a meaningful way. In most of these businesses, underlying earnings have continued to grow, buybacks have generally increased, and competitive positions appear intact – in some cases, stronger. Yet share prices have de-rated materially.

In our view, part of the problem is that investors are applying a broad “AI discount” to businesses that may look similar at a distance, but whose economics, competitive positions and exposure to disruption differ markedly. We continue to see this as an area where a long-term investor can add value: not by denying the importance of AI, but by distinguishing between businesses where AI may genuinely alter long-term industry structure and cash-flow durability, and those where we believe the effect is being extrapolated too broadly.

In a companion piece, we discuss the key drivers of performance and current positioning across the portfolio. But here, we want to dive deeper into AI, to assess where it is already changing the economics of an industry, where we think the effect is being overstated, and what that means – specifically – for the businesses we own. We believe the real question is not whether earnings are holding up today, but whether AI is likely to weaken a business in a lasting way. In many cases, we think investors are conflating what large language models (LLMs) are already good at with a much broader set of capabilities that they do not reliably possess.

### Where AI clearly works today

The pattern across how AI is actually being used today is fairly clear. AI works where five traits are present together : 1) the input is structured – text, code or formal data; 2) the output can be verified, either deterministically by a test or a reconciliation, or quickly by a reviewer; 3) feedback loops are short; 4) the workflow is digital end-to-end; and 5) the task is bounded. Where all five are present, AI is delivering meaningful productivity gains today, and in some cases doing the work outright. Where even two or three are absent – for example, if the execution timeline has a long horizon; tacit judgement and real-work physical context retain primacy; or output must be completely trusted and reliable – current performance is far less consistent.

Coding is the clearest best-use case. The input is text, the output is text, and tests or compilers can verify the result deterministically. Feedback is immediate, the workflow is fully digital, and tasks are naturally bounded at the function level. Rather than a categorical exception, coding is the leading edge of a general pattern. The interesting question is which adjacent workflows share the same traits.

Within most white-collar industries there is a layer of work that is genuinely in the crosshairs today. First-pass legal document review, reconciliation and control work in accounting and audit, candidate screening, content moderation, first-level customer support scripting, basic sales development representative (SDR) and lead-qualification outreach, and a meaningful portion of routine coding, testing and migration work in IT services – all these share the structured, verifiable, bounded character described above. The common thread is not a sector; it is a type of task. Within any industry, the tasks that tick all five boxes are the ones most likely at risk.

Two further observations are worth making. First, most white-collar AI adoption today still looks more like assistance than replacement. Serious users tend to use these tools to draft, check, refine and accelerate work, not as a black-box substitute for judgement. That distinction matters. Many business models built around hosting or organising knowledge work – workflow platforms, data platforms, professional services and similar – are being discussed as though AI will simply automate the task away. But if AI mainly helps people do the work better, rather than removing the need for them altogether, the platform hosting that work may remain a critical part of the workflow.

Second, business-to-business (B2B) workflow AI and consumer AI are not at the same stage, even if they are often discussed as though they are. B2B workflow AI – coding assistants, programmatic agents inside production pipelines and copilots embedded in enterprise software – is already real, monetisable and scaling. Consumer autonomous agents are not. Eighteen months after capable agents became widely available, consumer transaction volumes running through agent interfaces – shopping, travel, bookings and similar activities – remain a very small part of the relevant end markets. Consider, for example, the unwinding of OpenAI's Instant Checkout programme, which involved partnerships with major retailers that enabled customers to make purchases directly through ChatGPT. Walmart withdrew from the programme in March 2026 after encountering various issues, including customer complaints that ChatGPT was putting the wrong items in their checkout carts. Etsy also pulled out due to a lack of sales through the platform.

That is not to say that such initiatives will not be more successful in the future. It is simply a reminder that behaviour, trust and habit tend to move more slowly, and less in straight lines, than the underlying code.

## The friction test

A simple test we have found useful when thinking about our digital holdings is what we call the friction test. It is, in our view, one of the clearest ways to distinguish between businesses that are genuinely at risk and those that are not.

The question is not whether an AI tool can imitate part of the user experience. The question is whether taking the user out of the incumbent platform actually makes the task easier, cheaper or more reliable. In many cases, it does not. It simply moves the friction somewhere else. This matters particularly in consumer internet. Even if an agent can technically navigate a browser or complete part of a workflow, that does not mean it is offering a better service. In many categories, user intent is not fixed at the outset. Discovery is part of the product. A user looking for a hotel, a flight, a home or a restaurant is often refining preferences as they go – comparing options, changing dates, weighing trade-offs, reacting to price, location, reviews, availability and presentation. That process is not just a hurdle on the way to the transaction. In many cases, it is the transaction journey itself.

Even if agents become better at execution, removing the user from the ecosystem can also mean losing access to what makes the platform valuable in the first place – live inventory, merchant density, room availability, pricing, fulfilment, trust, service and integration with payments. In that sense, the friction is not eliminated. It is relocated. The user may save one step but lose quality, choice, reliability or price. There is also a behavioural point here. Friction in many of these categories is already low. The existing product works well, is familiar and is increasingly optimised for convenience. It is therefore not obvious that users will switch in meaningful numbers to a new interface that, at least today, is likely to offer a worse experience, weaker economics and greater security risk. That does not mean the top of the funnel is immune. Discovery is clearly an area to watch. But it is much less clear that a new winner is emerging, and the degree of change required to dislodge incumbent behaviour should not be understated.

Finally, our digital holdings are not standing still while this develops. These are not businesses waking up late to a new technology cycle. They are already close to the frontier, understand the direction of travel and are incorporating these tools into their own products and workflows. In many cases, the more likely outcome is not disintermediation, but adaptation – stronger products, better service and deeper customer relationships. That was part of the initial thesis we laid out in our October 2025 letter "[On AI: the Age of Extremes](#)", and recent product developments across the portfolio have only reinforced that view. Tencent's agents and Trip.com's Genie are good examples of this. In both cases, AI is being used to improve the product from within – making search, planning, service and execution more efficient – while leaving intact the underlying advantages of the platform: user relationships, trust, inventory, distribution and workflow ownership.

Against that backdrop, the recent weakness in a number of these businesses has, in our view, created a more compelling opportunity for long-term investors. We have therefore used this period to add to a number of our holdings, focusing on companies where we believe the underlying economics remain robust and, in some cases, are likely to strengthen.

## The businesses we own

One area where this distinction matters is enterprise software. TOTVS has been weak partly because the market has become concerned that AI will weaken the role of software vendors by making it easier for customers to build, adapt or orchestrate workflows themselves. We think that risk is being framed too broadly.

The reason is straightforward. TOTVS is not a generic software interface sitting on top of a workflow. It is deeply embedded in the workflow itself. Its products sit inside accounting, payroll, tax, inventory, compliance and other day-to-day operating processes, where the data is structured, the actions have consequences, and reliability matters more than elegance. In that environment, AI may change how users interact with the system, but it does not remove the need for a governed system of record. That is still where permissions sit, where audit trails live, where compliance logic is embedded and where the underlying process is executed. This is precisely why management has been clear that the right way to think about AI is not as a threat to enterprise resource planning (ERP), but as a way of extending what the platform can do.

What is more interesting, in our view, is how the model evolves from here. The key point is that AI does not remove the need for what TOTVS does; it changes how value is delivered and, over time, how it is priced. Today, most of TOTVS' revenue is still linked, directly or indirectly, to seats and modules – clients pay for access to software that allows them to run their operations. AI challenges that model at the margin. If a customer can automate part of the workflow, the number of users required to operate the system may fall. Taken in isolation, that looks like a headwind. But that framing misses the more important dynamic. The relevant question is not how many users sit in front of the system, but how much work flows through it. If AI allows a customer to run the same business with fewer people, the value of the system increases, not decreases. The system is doing more, not less. In that context, pricing is likely to move over time from seats towards usage, outcomes or tasks – what management describes as “Task as a Service”.

This is not theoretical. Even today, the cost of the software is a very small part of the customer's overall cost base, often a fraction of a percent of revenue. If AI allows a customer to reduce headcount or improve productivity, the economic value created sits primarily at the workflow level, not at the interface. Capturing even a small share of that value pool would more than offset any pressure from fewer seats.

That is why we think the debate around “AI killing software” is somewhat misplaced in this context. The system of record remains where execution happens. AI sits on top of it, makes it more useful, and over time allows the vendor to move up the stack – from enabling the work to doing part of the work itself. The risk is not that demand disappears, but that the monetisation model evolves. That transition will not be frictionless, and there will be pressure in parts of the model, but the direction of travel is not obviously negative.

Our base case is therefore not one of structural decline, but continued growth, supported by three drivers: ongoing penetration of ERP and cloud in Brazil, deeper embedding within existing clients as workflows expand, and a gradual contribution from AI-enabled products and services over time. The exact path will depend on execution, but the underlying demand for organising, governing and increasingly executing business processes appears intact.

## E-commerce – not just a checkout button

E-commerce is the other area where we think the current debate has become too simplistic. The concern is that agentic AI will sit between the customer and the platform, reduce discovery to a prompt, and turn the incumbent marketplace into little more than a fulfilment layer. We do not dismiss that risk entirely, particularly at the top of the funnel. But we think that conclusion is running ahead of the mechanism.

The key point is that purchasing is rarely a pure execution task. In many categories, the user does not begin with a fixed instruction and simply need an agent to carry it out. Discovery is part of the product. Users compare, refine, browse, change their mind, react to price, reviews, delivery times and availability, and often do not know exactly what they want until they are inside the platform. In that context, the marketplace is not just an interface sitting in the way of the transaction. It is a large part of the transaction itself. That is why our working view remains that discovery-led marketplaces are more likely to use AI to improve recommendation and search than to be disintermediated by it. Where they are being de-rated on AI fears alone, that looks more like candidate mispricing than confirmed structural damage.

This matters for both MercadoLibre and Sea. What is often described simply as “e-commerce” is, in reality, an integrated ecosystem of marketplace, logistics, payments, credit, advertising and merchant tools. These are not thin shopping interfaces. They are operating systems built around merchant density, fulfilment, trust and increasingly financial services. The marketplace may be where the customer starts, but a large part of the value sits underneath it. In MercadoLibre's case, Mercado Envios is integral to the proposition, reducing friction, improving the user experience and enabling faster

deliveries at competitive cost, while Mercado Pago deepens trust and engagement across both on- and off-platform transactions. In Sea's case, the same pattern can be seen in the combination of Shopee's logistics network and SeaMoney, which supports payments, credit and broader financial services across the ecosystem.

That is also why we would not reduce MercadoLibre's or Sea's economics to the checkout button. In both cases, payments, credit and merchant services are not side businesses; they deepen engagement, improve retention and create additional monetisation opportunities across the ecosystem. Proprietary data and machine-learning models also support underwriting and risk management in ways that traditional financial institutions often struggle to match. Even if discovery evolves at the margin, a large part of the value still sits lower in the stack.

That said, the recent weakness in these businesses has not been about AI alone. A more conventional concern has also weighed on sentiment: that competition will intensify, margins will stay lower for longer, and the economics will prove less durable than investors once assumed. China has conditioned investors to expect brutal convergence towards one or two dominant models, but Latin America and Southeast Asia are structurally different. Geography is harder, logistics are more local, merchant bases are more fragmented, and execution matters more. That does not mean competition is benign. It does mean that lower short-term margins do not automatically imply weaker long-term economics.

Our own view remains that MercadoLibre and Sea can coexist with other serious players in their respective regions while still earning attractive returns over time. In Sea's case, recent weakness (in addition to the AI fear) reflects a familiar concern that competition in ASEAN will keep profitability lower for longer. That may prove true near term. But the more relevant question is whether the current investment cycle is defensive or offensive. If the spending is deepening logistics, improving service levels and reinforcing habit while weaker competitors struggle, then lower margins today may be the cost of a

stronger platform tomorrow. We do not think investors are making that distinction clearly enough.

## Outlook

The past nine months have been sobering. But the underlying drivers of value in the GEM Focus portfolio have been more stable than recent price moves would suggest. Free cash flow has continued to grow across much of the portfolio and capital has generally been allocated with discipline. In other words, the businesses we own have, in aggregate, continued to increase their intrinsic value.

Looking forward, we expect the portfolio to compound earnings and free cash flow at mid-teens annual rates over the medium term. At today's prices, we are paying around a 5% free cash flow yield for that growth, against a return on capital in the low twenties.<sup>1</sup> We think that combination is genuinely attractive on an absolute basis, and it is a meaningful discount to the long-term average at which the portfolio has typically traded. The setup does not require any re-rating to deliver attractive returns from here. It needs our businesses to keep doing what they have been doing. And if valuations do eventually catch up with the earnings power of the businesses we own, the return profile from here improves further.

It is hard to shake off a sense of unease during periods of rapid technological change, particularly when the change is as fast-moving, and as prone to over-extrapolation, as AI has been. Yet for long-term investors, it is often in precisely such periods that the most attractive opportunities present themselves. When we look at the businesses we own, at the managers who run them, and at the prices at which they trade today, we remain confident in their long-term prospects. As the popular saying goes: tough times never last, but tough companies do.

We thank you for your continued patience and support.

The FSSA team

<sup>1</sup> FSSA Investment Managers, April 22, 2026.

## Important Information

This material has been prepared and issued by First Sentier Investors (Australia) IM Ltd (ABN 89 114 194 311, AFSL 289017) (FSI AIM), which forms part of First Sentier Group, a global asset management business. First Sentier Group is ultimately owned by Mitsubishi UFJ Financial Group, Inc (MUFG), a global financial group. References to “we” or “us” are references to First Sentier Group. Some of our investment teams use the trading names First Sentier Investors, FSSA Investment Managers, Albacore Capital, Igneo Infrastructure Partners and RQI Investors. A copy of the Financial Services Guide for FSI AIM is available from First Sentier Investors on its Australian website.

This material is directed at persons who are wholesale investors or clients and is not intended for persons who are retail investors or clients. Any advice in this material is general advice only. It does not take into account your objectives, financial situation or needs. Before making an investment decision you should consider, with a financial advisor, whether this information is appropriate in light of your investment needs, objectives and financial situation. The information in the material does not constitute an offer of, or an invitation to purchase or subscribe for any securities.

The product disclosure statement (PDS) or Information Memorandum (IM) (as applicable) for the FSSA Global Emerging Markets Focus Fund and 610 729 005 (Fund), issued by The Trust Company (RE Services) Limited (ABN 45 003 278 831, AFSL 235150) (Perpetual), should be considered before deciding whether to acquire or hold units in the Fund(s). The PDS or IM are available from First Sentier Investors. The target market determination (TMD) for the Fund is available from First Sentier Investors on its website and should be considered by prospective investors before any investment decision to ensure that investors form part of the target market.

MUFG, FSI AIM, their respective affiliates and any service provider to the Fund do not guarantee the performance of the Fund or the repayment of capital by the Fund. Investments in the Fund are not deposits or other liabilities of MUFG, FSI AIM, their respective affiliates or any service providers to the Fund and investment-type products are subject to investment risk including loss of income and capital invested.

Any opinions expressed in this material are the opinions of the individual author at the time of publication only and are subject to change without notice. Such opinions: (i) are not a recommendation to hold, purchase or sell a particular financial product; (ii) may not include all of the information needed to make an investment decision in relation to such a financial product; and (iii) may substantially differ from other individual authors within First Sentier Group.

Reference to specific securities (if any) is included for the purpose of illustration only and should not be construed as a recommendation to buy or sell the same. All securities mentioned herein may or may not form part of the holdings of a First Sentier Group portfolio at a certain point in time, and the holdings may change over time.

We have taken reasonable care to ensure that this material is accurate, current, complete and fit for its intended purpose and audience as at the date of publication. No assurance is given or liability accepted regarding the accuracy, validity or completeness of this material and we do not undertake to update it in future if circumstances change. No part of this material may be reproduced or transmitted in any form or by any means without the prior written consent of FSI AIM.

Past performance is not a reliable indicator of future performance. Unless otherwise stated, performance returns for periods greater than one year are annualised. Net performance returns are calculated assuming reinvestment of distributions and use exit prices which are net of management fees and performance fees (where applicable). Gross performance returns are calculated using a time-weighted return methodology. No allowance has been made for an investor's own brokerage when they buy their units on a securities exchange.

Copyright © First Sentier Group, 2026

All rights reserved.