

For personal use only

savana

# Letter to Investors

November 2025

In this month's Insights:

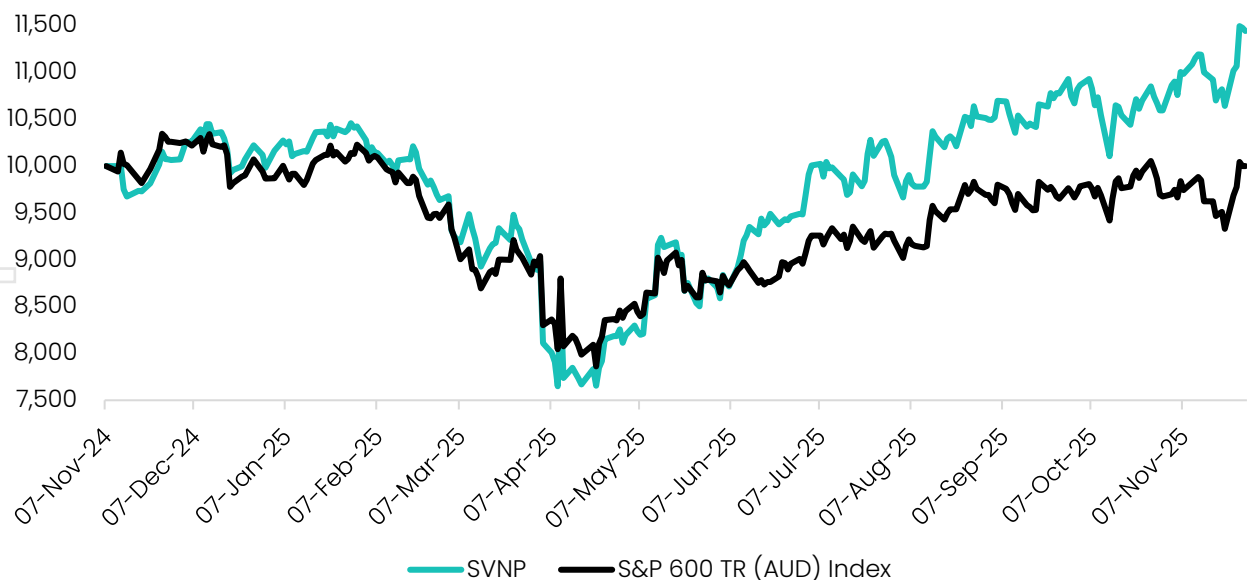
1. **Happy Birthday SVNP:** Reflections on our First Year of Performance
2. **The Art of the Sell:** The Payoffs – and Trade-Offs – of Strict Exit Discipline
3. **The AI Reality Check:** What Our LLM Stress Test Reveals About Reliability, Bias and Real-World Use

## Happy Birthday SVNP

### Reflections on our First Year of Performance

During November, Savana celebrated the one-year anniversary of its inaugural listed ETF, the Savana US Small Caps Active ETF (ASX:SVNP).

Despite a turbulent period in US markets, we're pleased to report that SVNP has delivered 13.62% p.a. since inception, representing 13.61% p.a. outperformance versus the benchmark. In a year where small caps broadly struggled, this result is both validating and encouraging.



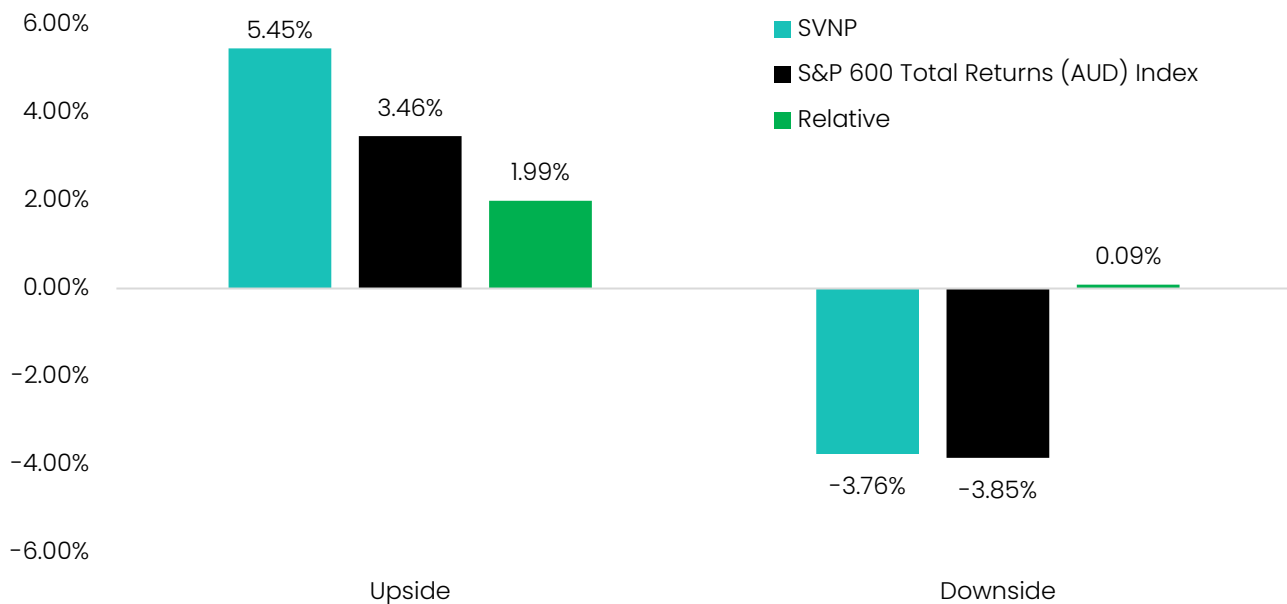
Source: Savana, S&P Global. Total returns are calculated in Australian dollars based on the close-of-day net asset value per unit. Returns are after fees and costs with dividends reinvested. Past performance is not a reliable indicator of future performance.

Importantly, the key attributes we observed during paper-trading continue to be validated in live performance: mitigated downside, strong upside capture, and

the ability to deliver genuinely differentiated, excess returns. These dynamics are clearly visible in SVNP’s total returns chart.

Through April–May, the portfolio navigated a sharp ~20% market drawdown triggered by heightened concerns over proposed US tariff policies. While the timing of such a broad sell-off so soon after listing was in some ways unfortunate, it has ultimately provided an invaluable stress test – reinforcing the resilience of our algorithms under the harshest conditions, and demonstrating its ability to capture significant upside on the rebound. This is reflected in our capture ratios: on average, in months where the market has gone up, SVNP has outperformed by 1.99%, while in months where the market has gone down, SVNP has outperformed narrowly by 0.09%.

## Monthly Upside / Downside Capture



Source: Savana, S&P Global. The chart displays the average monthly return for SVNP and the index since inception during months when the index increases (“Upside”) and decreases (“Downside”). Returns are after fees and costs with dividends reinvested. Past performance is not a reliable indicator of future performance.

Looking forward, despite continued uncertainty in the US market, we believe SVNP presents a compelling case under both Bull and Bear scenarios. If conditions improve and market breadth widens, SVNP stands to benefit from strong upside capture and a potential rotation into US small-to-mid caps. Conversely, if economic data and sentiment continue to weigh on returns, the need for high-performing, genuinely active management becomes even more important.

In our view, remaining invested and exposed to the US market – where returns have historically been higher and where innovation and productivity continues to

be world-leading – is an imperative for any investor. We see SVNP as a valuable complement to traditional core portfolios: diversifying US exposure away from the increasingly top-heavy S&P 500, while providing access to a differentiated, idiosyncratic return stream.

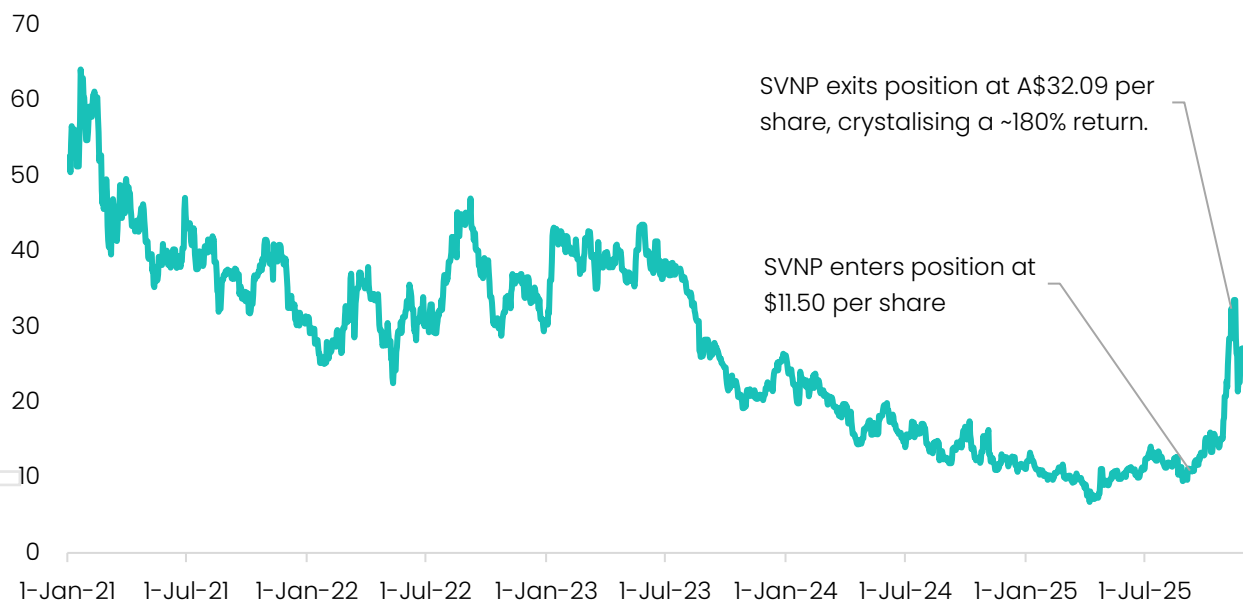
With a solid foundation established, we are incredibly excited to see what the second year of SVNP brings!

## The Art of the Sell

### The Payoffs – and Trade-Offs – of Strict Exit Discipline

Last newsletter, we spoke about the exceptional two-month performance of Canadian Solar (NASDAQ: CSIQ). We initiated our position on 8 September at \$11.50 per share, and exited 63 days later for a ~180% gain. Since our exit, the share price has retraced, falling ~16% to \$27.15.

#### Canadian Solar Share Price



Source: Savana, S&P Global.

CSIQ is a 'perfect' example of our disciplined, algorithmic strategy in action. Our models identified an opportunity in CSIQ after a five-year price slide – precisely the kind of “falling knife” situation that many investors avoid. And after a 180% rebound, when many would chase the momentum further, our signals told us the opportunity had played out. So we sold.

But it doesn't always land this perfectly. So, to keep thing real, we also wanted to share a few cases where our disciplined selling strategy didn't go quite to plan...

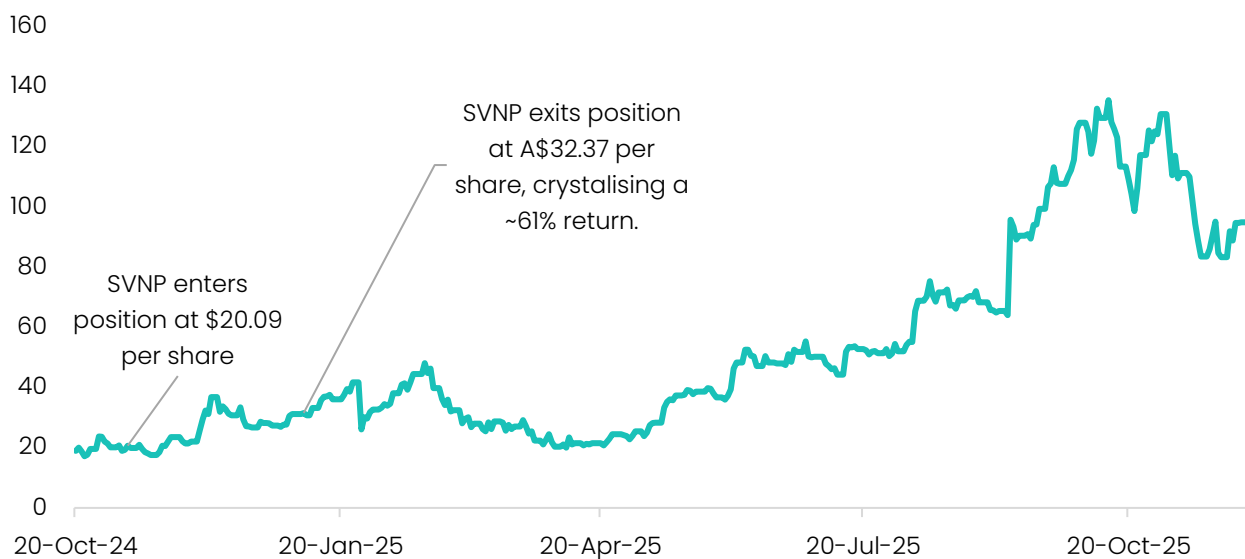


## Nebius Group NV (Nasdaq: NBIS)

Digital infrastructure provider Nebius was SVNP's first major winner. After quietly resuming Nasdaq trading in October 2024 (following a ~2.5-year suspension), SVNP initiated a position in November at \$20.09 per share and exited just two months later at \$32.37, realising a 61% gain. By March, the stock had slipped back to around \$20, which appeared to validate our decision to sell.

Since then, however, the share price has demonstrated that validation can sometimes be short-lived. Nebius now trades around \$94, after reaching \$135 in October. *Ouch*. The resurgence has been supported by a handful of developments - most notably surging demand for the company's AI-infrastructure offering and accelerating revenue growth - contributing to a substantial re-rating.

### Nebius Share Price



Source: Savana, S&P Global.

## SanDisk Corp (Nasdaq: SNDK)

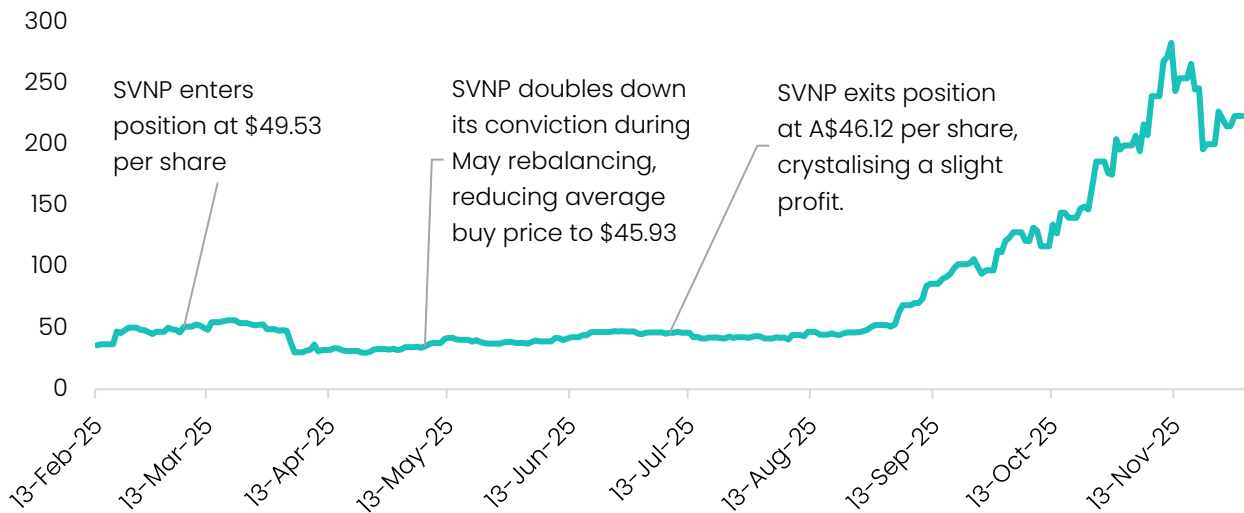
This one still hurts.

Following its spin-off in February, we initiated a position in SanDisk three weeks later at \$49.53 per share. As the market sold off through April–May, SNDK fell with it, retracing to around \$30 at the trough. By July, the stock had recovered to roughly \$46 per share. At that point, our algorithms identified more compelling mispriced opportunities elsewhere, and SNDK was narrowly excluded from the portfolio during the July rebalance.

If only we had held on.

Today, SNDK trades near \$220 per share - approximately 4.5x our original entry point. SanDisk is another beneficiary of the broader AI-infrastructure boom, with its flash-memory and data-centre storage products experiencing a surge in demand. The company's price-to-book ratio has expanded from 0.58 in March to over 3x today, reflecting this rapid re-rating.

## SanDisk Share Price



Source: Savana, S&P Global.

## Note on Process and Discipline

Episodes like this offer a useful reminder of what our system is designed to do - and what it is not.

Our mandate, built on more than a decade of R&D, is to allocate capital to the 30 most undervalued opportunities in the addressable market at any point in time, based on a consistent, valuation-driven framework. That discipline is what underpins the long-term performance of the strategy.

However, this approach also entails trade-offs. A model that systematically rotates towards the most undervalued names will, by design, sometimes exit positions before their full upside is realised. The alternative - holding onto past winners in the hope of further gains - introduces discretion, path-dependency, and behavioural bias, all of which our strategy is deliberately built to avoid.

In short: our edge comes from valuing stocks better than the market and acting with unwavering consistency. While this occasionally means leaving some upside on the table, it is this discipline - not speculation - that drives the robustness of long-term returns.

## The AI Reality Check

### What Our LLM Stress Test Reveals About Reliability, Bias and Real-World Use

With the rapid rise of Artificial Intelligence, an existential question now confronts the industry:

#### *What does AI mean for the future of investment management?*

Amidst the global 'AI race', a growing number of new entrants are promoting AI-driven research engines and decision tools that promise sharper insights, faster analysis, and a step-change in active management capability.

But beneath the froth and excitement lies a fundamental question: *are today's LLMs reliable enough to support - or meaningfully influence - live portfolio decisions?* Their fluency is undeniable, but fluency does not guarantee the coherence, logic, or lateral reasoning required for real-world investing.

To explore this, we ran a controlled experiment. We asked Claude Sonnet 4.5 to generate Buy/Hold/Sell recommendations for 30 stocks using impartial, long-form research reports produced independently by ChatGPT. We then stress-tested those recommendations by:

1. Issuing repeated queries with identical input data
2. Introducing controlled lexical and structural rephrasings of the same reports
3. Applying positive and negative framing biases while preserving all underlying facts.

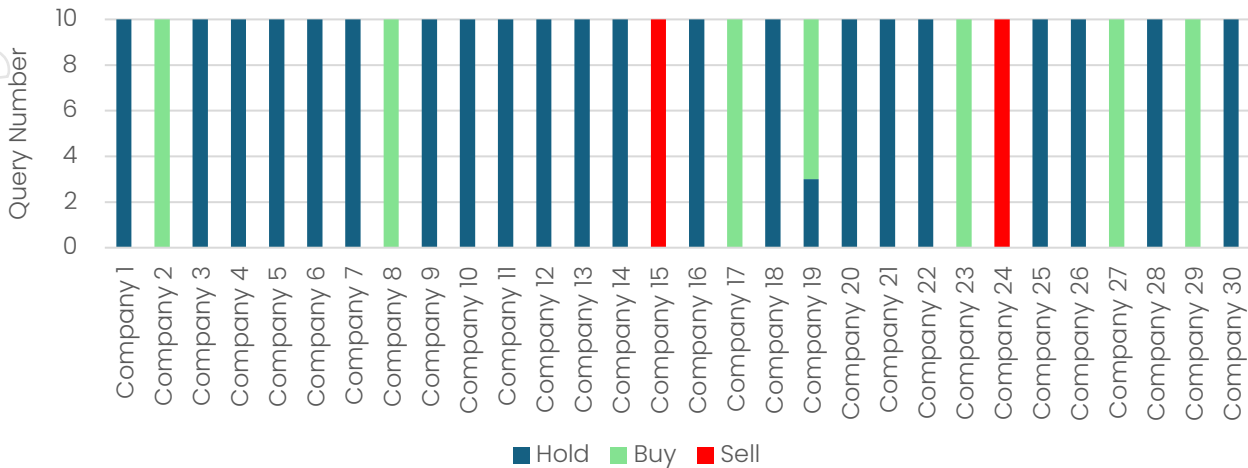
The outcomes were striking.

#### **Test 1: Stability Under Identical Inputs**

Our first test asked a basic question: does the model behave consistently when nothing changes?

The answer was mostly yes. Across 30 companies and repeated identical prompts, the model produced the same recommendation 97% of the time.

## Recommendation When Repeating Queries



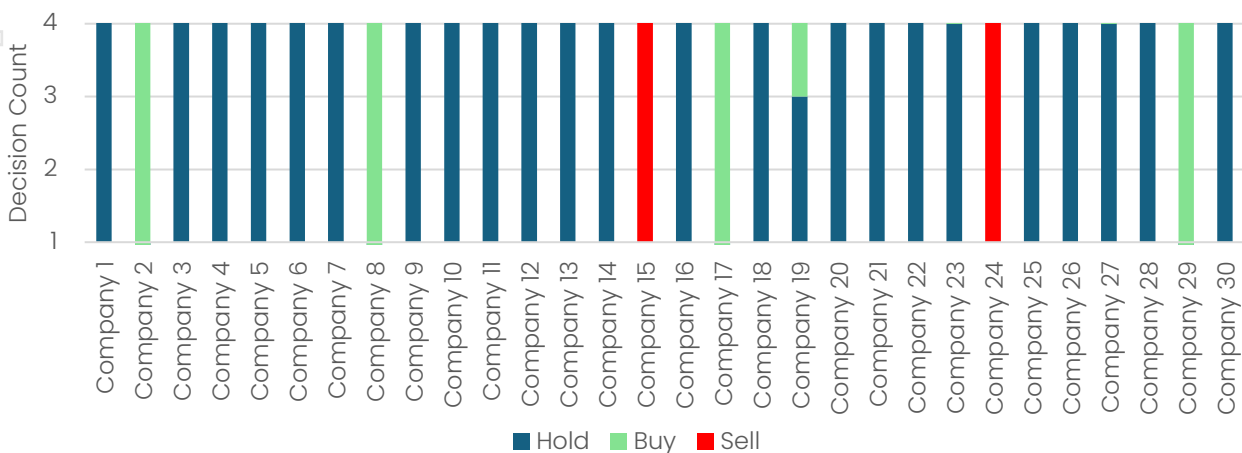
Only one company (Company 19) shifted (from “Hold” to “Buy”), showing that while LLMs are trained via stochastic processes, their inference-time behaviour is generally deterministic.

**Takeaway:** LLMs can be stable - but isolated “flips” remind us that determinism should not be assumed to be perfect.

### Test 2: Sensitivity to Neutral Rephrasing

Next, we tested whether the model’s decisions were robust to benign changes in wording. We re-wrote each research report four ways - reordered, rephrased, expanded, and condensed - while keeping all facts neutral and identical.

### Recommendations With Semantical Perturbations



Here the cracks began to show.



About 17% of stocks experienced at least one change in recommendation due solely to neutral rewording. One company (Company 27) showed no robustness at all, flipping from its original “Buy” baseline to “Hold” under every version.

**Takeaway:** Current LLMs remain surprisingly sensitive to presentation rather than substance. Even subtle linguistic shifts – all factually identical – can meaningfully influence their decisions.

In real-world investing, information rarely arrives in a perfectly standardised format. The fact that small shifts in wording cause non-trivial decision flips suggest that LLMs do not yet possess the abstraction or invariance required for consistent investment decision-making.

### Test 3: Sensitivity to Framing Bias

Finally, we tested whether tone influenced the outcome. Each research note was regenerated twice: once with positively framed language and once with negatively framed language, while preserving all underlying facts.

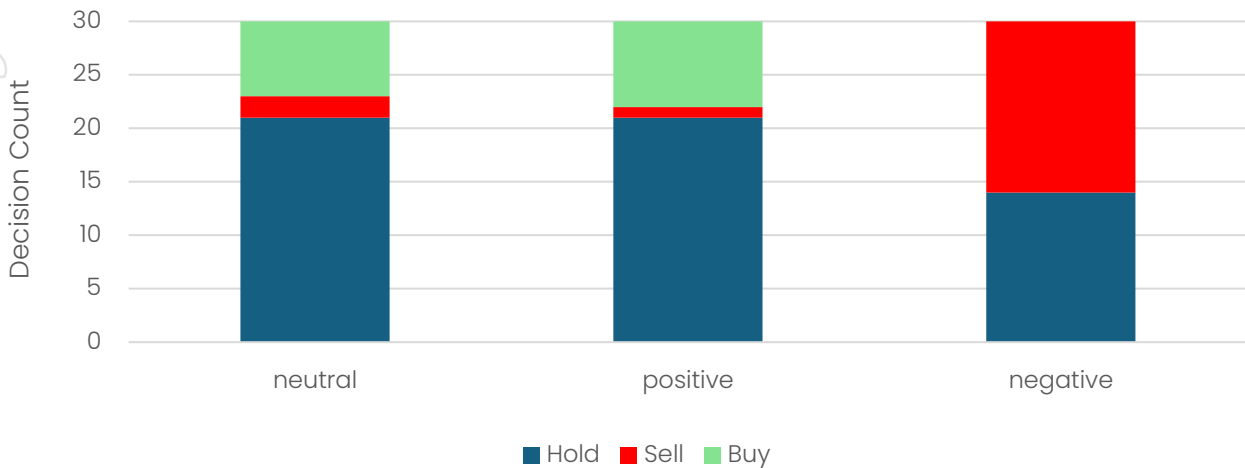
#### Example Reframing for Kohls (NYSE: KSS)

Framing	Excerpt
Positive	<i>“In November 2025, Kohl’s appointed Michael Bender as its permanent CEO, following his successful interim leadership since May 2025.”</i>
Neutral	<i>“In November 2025, Kohl’s appointed Michael Bender as its permanent CEO, following his interim leadership since May 2025.”</i>
Negative	<i>“In November 2025, Kohl’s appointed Michael Bender as its permanent CEO, following his interim leadership since May 2025—making him the fourth CEO in four years, a troubling pattern that raises concerns about strategic continuity.”</i>

The impact was material.

Positive and neutral framing produced similar distributions of recommendations. But negative framing caused a heavy skew toward Sell, despite the factual content being unchanged.

## Recommendations By Framing Strategy



**Takeaway:** LLMs respond strongly to sentiment cues. When language becomes even slightly cautious or risk-emphasised, their recommendations shift accordingly. In real markets - where tone, emphasis, and sentiment vary constantly - this creates a material vulnerability.

### The Bottom Line

Across the three tests, a clear conclusion emerges: LLMs are generally stable when fed identical inputs, yet they remain highly sensitive to how information is presented. Even small, neutral shifts in wording caused decision flips in a meaningful number of cases, while negative framing pushed many recommendations toward Sell despite the underlying facts being unchanged. Taken together, this suggests that today's models are not robust enough to act as independent investment decision-makers without very careful input control.

So, while LLMs certainly add genuine value as research accelerators - helping synthesise information, surface signals, and streamline workflows - as engines for live portfolio allocation, they still lack the consistency, abstraction, and invariance required for dependable use.

## More Information

If you would like more information about this report or Savana, please contact [enquiries@savana.ai](mailto:enquiries@savana.ai). You can also speak to a member of the team below:



Marc Maasdorp  
Chief Executive Officer  
[marc.maasdorp@savana.ai](mailto:marc.maasdorp@savana.ai)



Samuel Atkinson  
Associate Director  
[samuel.atkinson@savana.ai](mailto:samuel.atkinson@savana.ai)

### DISCLAIMER:

This document has been prepared by Savana Asset Management Pty Ltd (ABN 79 662 088 904) (Savana). Savana is a corporate authorised representative of Fat Prophets Pty Ltd (ABN 62 094 448 549 AFS Licence No. 229183) (Fat Prophets), CAR Auth No. 1308949. The Savana US Small Caps Active ETF (ASX: SVNP) (ARSN 649 028 722) is issued by K2 Asset Management Limited (K2) ABN 95 085 445 094, AFS Licence No 244393, a wholly owned subsidiary of K2 Asset Management Holdings Limited (ABN 59 124 636 782). The information contained in this document is produced in good faith and does not constitute any representation or offer by K2, Savana or Fat Prophets.

This material has been prepared for both retail and wholesale investors and is for information purposes only. It is not an offer or a recommendation to invest and it should not be relied upon by investors in making an investment decision. Offers to invest will only be made in the product disclosure statement ("PDS") available from [www.savana.ai](http://www.savana.ai) and this material is not intended to substitute the PDS which outlines the risks involved and other relevant information. Any investment carries potential risks and fees which are described in the PDS. A Target Market Determination has been prepared for this product and is available from the same website. An investor should, before deciding whether to invest, consider the appropriateness of the investment, having regard to the PDS in its entirety. This information has not been prepared taking into account your objectives, financial situation or needs. Past investment performance is not a reliable indicator of future investment performance. No representation is made as to future performance or volatility of the investment. In particular, there is no guarantee that the investment objectives and investment strategy set out in this presentation may be successful. Any forward-looking statements, opinions and estimates provided in this material are based on assumptions and contingencies which are subject to change without notice and should not be relied upon as an indication of the future performance. Persons should rely solely upon their own investigations in respect of the subject matter discussed in this material. No representations or warranties, expressed or implied, are made as to the accuracy or completeness of the information, opinions and conclusions contained in this material. In preparing these materials, we have relied upon and assumed, without independent verification, the accuracy and completeness of all information available to Savana. To the maximum extent permitted by law, all liability in reliance on this material is expressly disclaimed.

This document is strictly confidential and is intended solely for the use of the person to whom it has been delivered. It may not be reproduced, distributed or published, in whole or in part, without the prior approval of Savana.

For personal use only