Axioma By SimCorp AXYON.AI

WHITE PAPER

FROM AI SIGNALS TO ACTIVE RETURNS: A PORTFOLIO MANAGER'S GUIDE TO CAPTURING CONSISTENT ALPHA

March 2026

CONTENTS

Introduction

An alpha signal is only as good as the portfolio it powers. While backtest results are no guarantee of future returns, they can guide us in how best to implement a strategy and demonstrating that a signal survives the rigors of real-world portfolio construction is crucial.

Axyon AI and SimCorp have teamed up to show that the combination of cutting-edge AI-driven signals and world-class portfolio construction tools is greater than the sum of its parts. For Axyon AI, this means proving its signals translate into genuine market-beating portfolios; for SimCorp, it is an opportunity to demonstrate how its Axioma risk and optimization tools[1] can unlock alpha sources.

In this paper, together we demonstrate how an active manager may want to think about building such a strategy. We set out the methodology, validate the signals and then apply them with a practical framework a long-only asset manager might use.

About the Signals

Axyon AI's signals are designed to identify stock-level out and underperformers within investment universes. Large datasets are ingested, cleaned, validated and centralized into Axyon AI's AI-ready features for model training. Across the full modelling of all the AI models, over 2 million rows of data are processed per day, supported by high-performance computing and scalable cloud infrastructure.

Model development is driven by an automated Learning-to-Rank (LTR) framework that trains between 10,000 and 100,000 candidate models per universe, selecting optimal combinations through an automated ensembling process. By combining heterogeneous models, the ensemble approach reduces overfitting, improves stability across market regimes and enhances predictive consistency.

Axyon AI's signals provide alpha sources over various forecast time horizons: 1 week or 5 trading days, 1 month or 20 trading days and 3 months or 60 trading days. Signals are delivered daily before the European market opens. Within each time horizon users can access the most recent signal, an average signal for the time horizon and a time weighted version where more recent signals receive progressively higher weight through a linear decay function, improving responsiveness while retaining the robustness of averaging.

1 Source: http://www.simcorp.com/axioma

The latter two are intended to smooth out the alpha and therefore reduce portfolio turnover. Although Axyon AI provides signals on over 6,000 symbols or over 98% of the world's market capitalization, in this use case we will concentrate on the US All Cap universe (Large-, Mid-, and Small-cap stocks in the US, representing the top 99% of the investable universe by float-adjusted market capitalization) and the 1-month (20 trading day) signals.

Translating Rankings into Alpha Forecasts

Axyon AI's signals are provided as rankings of the assets for each date. In order to translate rankings into alpha forecasts we follow the general procedure of (Grinold, 1994) to compute expected returns from the rankings based on z-scores. We compute the z-scores using the standard procedure of incorporating the probit function to convert rankings to z-scores. See (Cunnane, 1978) for a survey of such methods. Note that the probit function is the inverse of the cumulative distribution function associated with the standard normal distribution. As an example, if we have 1,500 assets ranked from 1 to 1,500, with 1 being best, then our z-scores for the best and worst assets are:

$$\alpha(1) = -\Phi^{-1}\left(\frac{1-0.5}{1500}\right) = 3.40293 \; and \; \alpha(1500) = -\Phi^{-1}\left(\frac{1500-0.5}{1500}\right) = -3.40293.$$

The strategies used in this research have an objective that maximizes expected return, so we do not scale by IC or volatility as this would not change our optimal portfolios[2].

Validating the Axyon AI Alphas

Our first step was to confirm that Axyon AI's signals behave as expected, and to see if one variant dominates the others. In order to determine if the signals have predictive power over and above whatever industry or style exposures it may have, we created factor-mimicking portfolios (FMPs)[3], similar to those used to build Axioma risk models. The FMPs were constructed to be minimum variance long-short portfolios, with unit exposure to the alphas and no net exposure to any style factors (e.g. Size, Momentum, Earnings Yield, etc.) or industry factors (using the latest GICS industries). The underlying risk model was the Axioma US5.1 medium-horizon model. These "test FMPs" were generated monthly at the end of each calendar month for the 10 years ended December 2025.

2 We also considered computing the expected returns via a stacked isotonic regression of the rankings versus subsequent 20-day and 60-day returns (for the respective signal horizon) over many historical periods to compute the mapping of ranking to expected return empirically. The mapping of rankings to expected returns resulting from this regression generally had the same shape as the probit function, so we opted for the much simpler probit function mapping.

3 Source: https://www.simcorp.com/resources/insights/industry-articles/2025/five-ways-factor-mimicking-portfolios-reveal-hidden-insights

We found strong evidence of consistent positive returns for each variant of the alpha model from 2015 through August 2024 (Exhibit 1). After that there was a substantial drawdown through the rest of 2024, a short recovery period and then another smaller shortfall through the end of the test period. While the version that used the month-end forecast ("Last") fared better than the others for most of the test period they all experienced similar drawdowns and ended the 10 years with roughly the same return. In addition, the "Last" version required higher turnover, which was not surprising given smoothing of alphas in the others.

Despite the drawdown in the last 1.5 years of the test we found the case for using these alphas quite compelling. Over the full period they produced admirable information ratios, even accounting for modest transaction costs (Exhibit 2). Results were consistent as well, with the FMP generating a positive return in 78-80 of the 120 months of the study, or roughly two-thirds of the time.

Exhibit 1 Cumulative return, 20-day alpha factor-mimicking portfolios

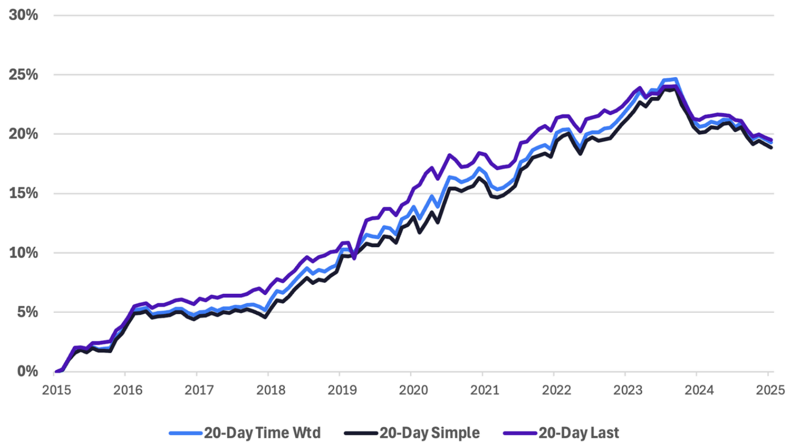

Exhibit 2 Summary Statistics

	Time Weighted	Simple Average	Last
Annualized Return	1.77%	1.74%	1.78%
Average Annualized Risk	1.69%	1.68%	1.60%
Average 2-Way Monthly Turnover	133%	128%	150%
Information Ratio	1.05	1.03	1.11
T cost (10 bps) * Turnover	0.0133%	0.0128%	0.0150%
T cost per year	0.16%	0.15%	0.18%
Return adjusted for t cost	1.61%	1.59%	1.60%
Information Ratio adjusted for t cost	0.95	0.94	1.00

Note: tests run from January 2016 through December 2025. Portfolios were rebalanced monthly.

Real-world Testing: Long-Only Portfolios

Having validated the efficacy of the signals, we then moved on to creating more realistic portfolios. We used the Axioma Portfolio Optimizer for portfolio construction, along with the US5.1 medium-horizon risk model to assess the risk-return tradeoff.

There are so many different strategy variations used by portfolio managers. We aim to consider some of the variations we see to get a good cross-sectional representation of long-only managers. For many of these choices, using the “frontier” feature in the optimizer allowed us to easily assess the impact of these various choices. We hope to find that the long-only constraint does not have an unduly negative impact and that results – active returns – are consistent through time without huge drawdowns.

For the tests described below we chose to use the 20-day time weighted signal. Since all three variants produced similar long-term results, any would have likely told the story. However, we believe the choice of time-weighted strikes a balance between the recency of the month-end (“Last”) forecast and the lower turnover generated by the simple average of the alphas over the month. We also tested both Russell 1000 and Russell 3000 benchmarked portfolios, and found we got slightly higher information ratios using the Russell 3000, so the rest of the paper will focus there. We invite readers interested in results for the other alphas or benchmark to contact us for those results.

This first set of tests was run with the objective of maximizing expected active return, given active risk targets of 1%, 2%, 3% and 4%. The only other constraint was that the portfolio could not hold short positions and had to be fully invested. While this scenario was more “real-life” than the FMPs discussed above, it still involved a good deal of turnover and may have taken other factor or industry bets, which we will discuss shortly.

Not surprisingly, we see increasing active returns (Exhibit 3) as our target active risk increases, although at 1.26 the information ratio is highest at the 1% active risk target for the Russell 3000 portfolio (realized active return of 1.85% versus realized active risk of 1.47%). At the 2% TE target the information ratio falls to 1.20, with a return of 3.18% and realized risk of 2.65%.

Although the Axyon AI signals do not cover all of the names in the Russell 3000 universe, as of December 31, 2025, roughly 98% of the *market value* was captured. We chose the Russell 3000 as the benchmark for the rest of our tests because it is well-known and commonly used and produced better information ratios than we saw for the Russell 1000. In the optimization, the whole Russell 3000 index is considered for possible investment. If there is no signal the optimizer may take a small position if it is a good enough hedge but these positions are likely to be small.

Exhibit 3 Active risk frontier, Russell 1000 and Russell 3000

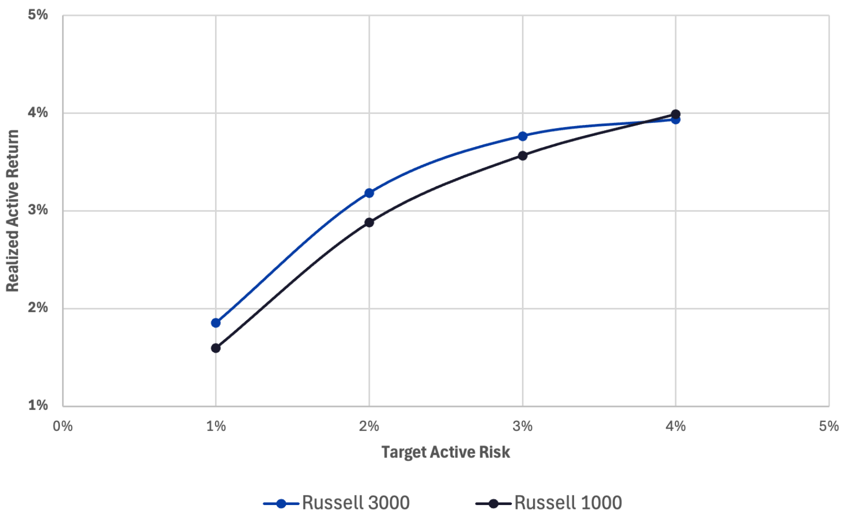

Note: Based on monthly rebalancing, 2016 - 2025

Source: FTSE Russell, Axioma US5.1 risk model, Axyon AI 20-day time weighted alpha forecast

While – like every investment process – the strategy does not outperform the benchmark in every period, it is clear that the cumulative return over time is not dominated by just a few good months. In fact, the active return was positive in 75, 77, 78 and 74 of the 120 months for the 1%, 2%, 3% and 4% targets, respectively, so 61% - 65% of the time. Drawdowns were clearly bigger the higher the tracking error, and the post-Covid era was difficult, but this unconstrained strategy was profitable at all levels of active risk (Exhibit 4).

Risk Characteristics of the Unconstrained Portfolio

Since the drop-off in information ratio was small as we went from 1% to 2% active risk, and the return was substantially higher, we chose to use 2% for our next set of tests.

Portfolio returns are ultimately tied to the risks taken. We analyzed the components of risk in the unconstrained portfolio with 2% active risk, and found what we hoped to find: that the majority of risk is stock specific, with less than 10% each coming from style factors and industries. As for individual style factors, exposures varied through time, but on average all accounted for less than 5% of the variance (Exhibit 5).

©SimCorp 2026

Exhibit 4 Cumulative active return by target tracking error, unconstrained portfolios

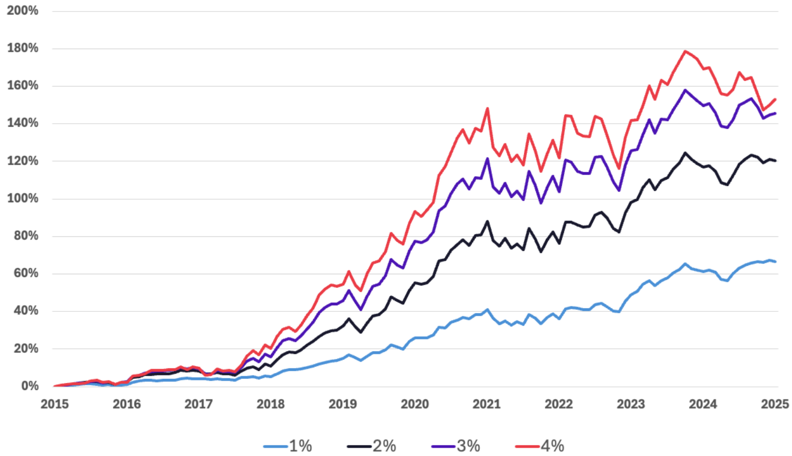

Note: Based on 20-day weighted average alpha forecast, benchmark Russell 3000

Exhibit 5 Percent of variance by factor block (left) and by individual style factors (right)

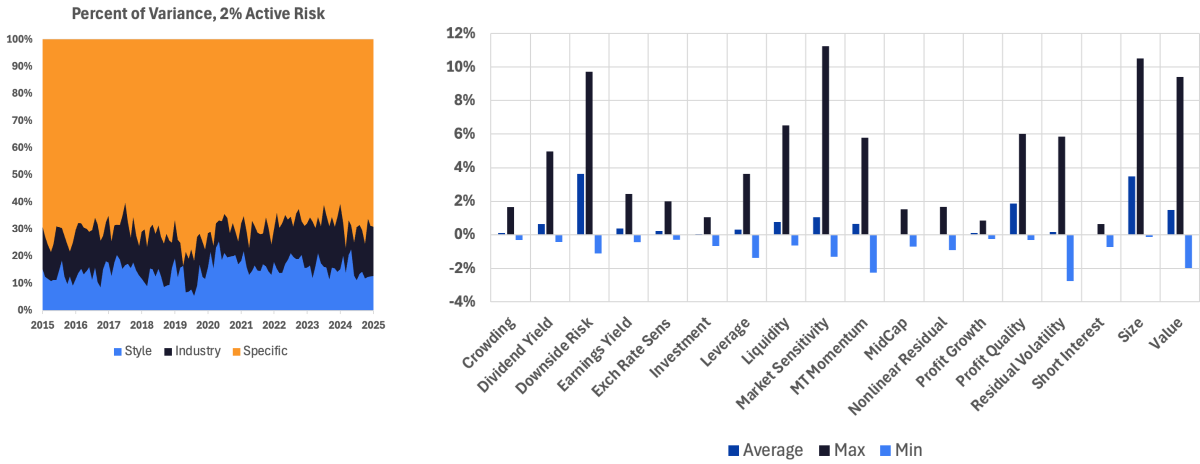

Deeper dive into attribution

We ran factor attribution for the unconstrained 2% portfolio, and confirmed that as expected given the risk profile, almost all the active return was stock specific, rather than coming from industry or style bets (Exhibit 6). This confirms that the Axyon AI signal is finding stock characteristics that are not easily replicated with a factor model, and that the performance is not the result of some big industry bets.

Of course, the small total factor contribution could be the result of factors cancelling each other out, even though their risk contributions were small on average.

The results do not bear out that possibility, however, as we see in Exhibit 7. Most factor contributions were small compared with the overall active return, with only low Downside Risk exposure and high Profit Quality exposure contributing meaningfully, and no style factors producing a big drag on returns. Similarly, an overweight in the Industrials sector was additive, while only Materials detracted from returns.

Exhibit 6 2% target active risk, cumulative active return attribution

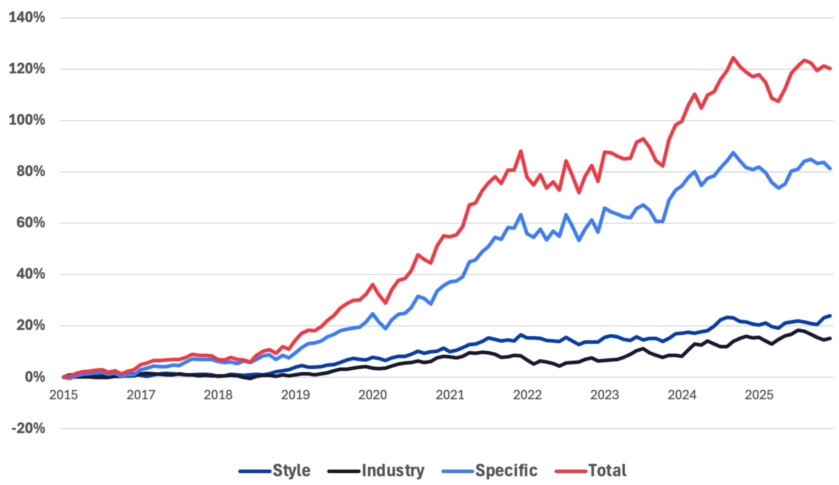

Adding Constraints

Many managers are required to constrain certain risk exposures such as industry weights as a “belt and suspenders” approach to maintaining levels of active risk.

Constraints on optimization typically lead to less-optimal solutions, and it has been said that if you trust your risk model (which of course we do) you shouldn’t need constraints to ensure ex-post risk is within expectations. Still, we ran one version of our tests constraining exposures to all risk factors to be equal to those of the benchmark while maintaining the same 2% tracking error target. We could have singled out the exposures that would have been expected to hurt performance over time because of their long-term risk premia and constrained only those, but if the tightly constrained version produces acceptable results, then we would expect the selectively constrained version to do even better.

Exhibit 7 2% target active risk, individual style and sector return contributions

Source of Return	Contribution	Avg Exposure
Portfolio	**17.48%**	
Benchmark	**14.30%**	
Active	**3.18%**	
Specific Return	**2.15%**	
Factor Contribution	**1.03%**	
Style	**0.63%**	
Crowding	0.17%	0.03
Dividend Yield	0.08%	0.02
Downside Risk	0.38%	-0.10
Earnings Yield	-0.04%	-0.00
Exchange Rate Sensitivity	0.03%	-0.01
Investment	-0.01%	0.03
Leverage	0.00%	0.16
Liquidity	-0.07%	-0.04
Market Sensitivity	0.05%	0.01
Medium-Term Momentum	-0.04%	-0.04
MidCap	0.00%	0.04
Nonlinear Residual	0.05%	0.03
Profit Growth	0.00%	0.03
Profit Quality	0.31%	0.13
Residual Volatility	-0.05%	-0.01
Short Interest	-0.04%	-0.01
Size	-0.17%	-0.06
Value	-0.01%	-0.13
Sectors	**0.40%**	**0.00%**
Communication Services	0.03%	-1.39%
Consumer Discretionary	0.00%	-0.50%
Consumer Staples	-0.08%	-0.85%
Energy	0.04%	-0.70%
Financials	-0.06%	0.14%
Health Care	0.05%	-0.22%
Industrials	0.44%	2.94%
Information Technology	0.14%	1.23%
Materials	-0.24%	0.80%
Real Estate	0.02%	-0.64%
Utilities	0.07%	-0.80%

Note: Based on 20-day weighted average alpha forecast, benchmark Russell 3000

Over the full period of the test the unconstrained portfolio far outperformed the constrained version (Exhibit 8). However, most of that outperformance came in the first five years of the 10-year test period. After that, the return difference remained stable. And in fact, the constraints helped during the drawdown period in 2024 mentioned earlier.

We noted above that while the impact of incidental style and industry exposures was small, it was nonetheless positive. So, it is not surprising that when we constrain those exposures it has a major negative impact on return.

The constrained portfolio not only lost the benefit of the positive contribution from styles and industries, but we also note a small reduction in the stock-specific contribution (Exhibit 9). Adding constraints seems to detract from the stock selection benefits of the alphas.

Exhibit 8 2% target active risk, cumulative active return of constrained versus unconstrained optimization and the difference between them

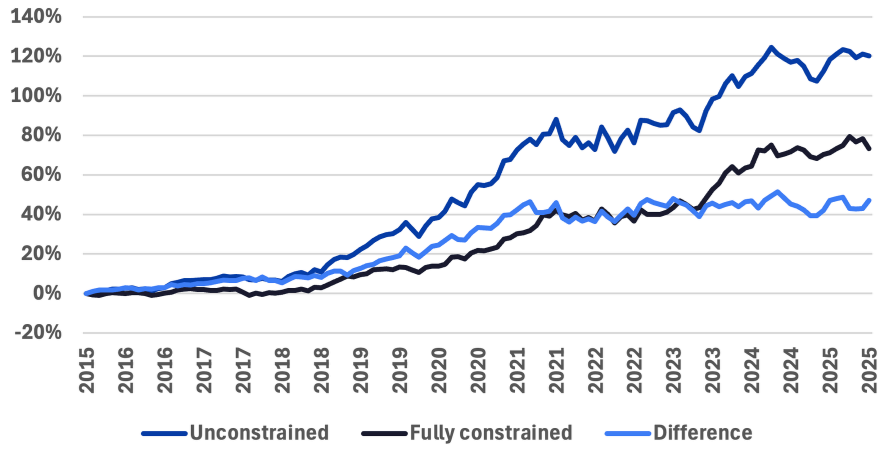

Exhibit 9 2% target active risk, comparison of performance attribution, constrained versus unconstrained optimization

Source of Return	Unconstrained	Constrained
Portfolio	17.48%	16.33%
Benchmark	14.30%	14.30%
Active	3.18%	2.03%
Specific Return	2.15%	2.03%
Factor Contribution	1.03%	0.00%
Style	0.63%	0.00%
Industry	0.40%	0.00%

There is one constraint that is usually unavoidable, however: turnover. With no constraint our 2% active risk portfolio averages 100% two-way turnover per monthly rebalance. This likely adds significant transaction costs, not to mention scaring off asset owners.

Adding any constraint that restricts trading introduces path dependency to the outcome, making conclusions more difficult by bringing up the question of how much of the performance is related to the legacy names in the portfolio that could not be traded because of the constraint. To try to mitigate this issue we ran several turnover frontier tests with different starting points (although the same end point). While this may not fully address the path dependency, it can provide some guidance to how big the problem is.

First, the good news is that, at least for our four test periods, there was not a meaningful difference in active return or information ratio under any of the turnover levels. In other words, path dependency does not seem to be a problem for these tests (Exhibit 10).

Exhibit 10 2% target active risk, Turnover frontiers, various starting years

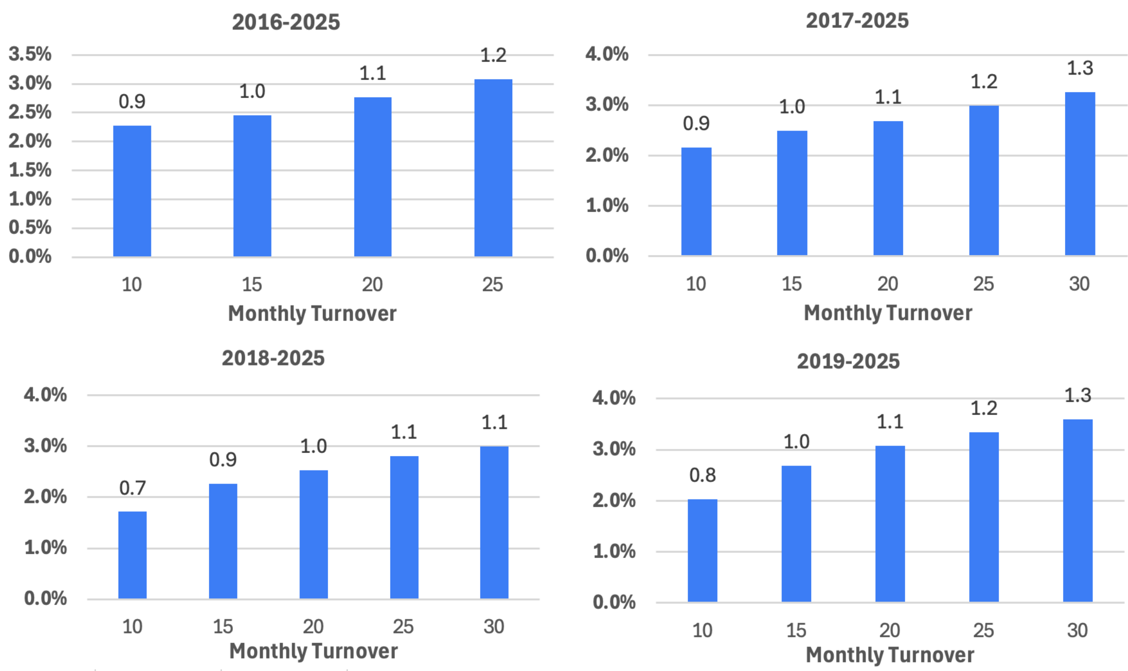

It is not surprising that higher turnover leads to higher return and higher information ratios As noted, when turnover is not constrained it averages about 100% per month, so these levels of turnover are restrictive. However, in that “best case” unconstrained scenario for our 2% active risk portfolio the annualized active return over the 10-year period was 3.2% (versus 3.1% at 30% turnover) and the information ratio was the same, 1.2, with just 30% of the turnover.

We conclude that the benefits of the Axyon AI alpha can be gleaned without taking on the excessive trading from having no turnover constraint. Users would need to determine the level of turnover they can bear and see if the return level is acceptable. Alternatively, one could choose an acceptable level of return and information ratio, and limit turnover to the associated level. For example, if an acceptable target information ratio is 1.0 at a target active risk of 2%, then choosing 15-20% turnover per rebalance should create the desired result.

Conclusion - Unlock Differentiated Alpha

The results presented in this paper demonstrate that the US All Cap Axyon AI’s 20-day signals not only have informative power beyond traditional factors, but also translate into tangible implementable alpha in real-world portfolios. Across both factor-mimicking portfolios and practical long-only constructions, the signals generated consistent active returns with compelling information ratios over a full market cycle.

This paper reveals that Axyon AI's portfolio signals serve as a consistent source of active return within a broader mandate over US Stocks. Notably, the majority of the alpha generated is stock-specific, proving that the AI successfully identifies sources of return that go beyond traditional factors.

Even good alpha forecasts can lose their power if the portfolio construction process is not robust, recognizing that expected return must be traded off against risk to ensure the best possible performance that is consistent and repeatable. We see the strong, effective signals from Axyon AI and the optimization and risk modeling capabilities of SimCorp as a winning combination to help investors produce consistent risk-adjusted returns.

Reach out to your SimCorp Axioma representative to explore how Axyon AI’s signals can enhance your existing investment process, whether as a standalone alpha source or integrated into your current investment management process. Discover how AI-driven signals can provide a differentiated edge in increasingly competitive markets.

Cunnane, C. (1978, May). Unbiases plotting positions - A review. Journal of Hydrology, 37(3--4), 205-222.

Grinold, R. C. (1994). Alpha is Volatility Times IC Times Score. The Journal of Portfolio Management, 20(4), 9--16. doi:10.3905/jpm.1994.409482

ABOUT SIMCORP

SimCorp provides a comprehensive suite of factor risk models, multi-asset enterprise risk management and portfolio construction tools through its Axioma analytics offering. Its client base includes global asset managers, asset owners, hedge funds and wealth managers.
For more information, please visit: https://www.simcorp.com.

ABOUT AXYON AI

Axyon AI is transforming how investment managers harness AI for a competitive edge. It develops advanced AI solutions to enhance decision-making, generate predictive and thematic insights and build transparent, explainable strategies. Its platform combines proprietary AI models with a human-in-the-loop framework to ensure actionable outputs and control.
For more information, please visit https://axyon.ai.

LEGAL NOTICE

The contents of this publication are for general information and illustrative purposes only and are used at the reader's own risk. SimCorp uses all reasonable endeavors to ensure the accuracy of the information. However, SimCorp does not guarantee or warrant the accuracy, completeness, factual correctness, or reliability of any information in this publication and does not accept liability for errors, omissions, inaccuracies, or typographical errors. The views and opinions expressed in this publication are not necessarily those of SimCorp. © 2026 SimCorp A/S. All rights reserved. Without limiting rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form, by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose without the express written permission of SimCorp A/S. SimCorp, the SimCorp logo, SimCorp Dimension, and SimCorp Services are either registered trademarks or trademarks of SimCorp A/S in Denmark and/or other countries. Refer to www.simcorp.com/trademarks for a full list of SimCorp A/S trademarks. Other trademarks referred to in this document are the property of their respective owners.