

AI and quant investing: a smarter collaboration

Artificial intelligence is reshaping quantitative investing. Systematic asset managers use AI to refine factors, decode Fed language, and enhance research efficiency. Yet the partnership between humans and machines remains key.

By Dr Maximilian Stroh

Artificial intelligence (AI) has made dramatic advances in recent years, driven by progress in machine deep learning and, more recently, large language models (LLMs). Systematic asset managers have been using AI-related technology for almost two decades and remain early adopters of LLMs. While the potential is significant, there are also many concerns regarding a 'black box' approach to investing.

This article discusses the opportunities and limitations of AI in quantitative investing, focusing on a few examples.

The 'small data problem' in finance

Unlike domains such as image recognition or natural language processing, financial markets provide little clean and stationary data. Equity return series, for instance, only span a few decades, with structural breaks, such as the dot-com bubble, the GFC, and COVID, limiting comparability.

The dimensionality of the problem is vast: thousands of

stocks, hundreds of factors, and multiple layers of interaction across sectors, geographies, and time horizons. Statisticians call this the 'curse of dimensionality'. There is simply not enough data to reliably estimate highly complex models without overfitting.

Finally, intense competition in (almost) efficient markets makes it even harder to distinguish meaningful signals from noise in return predictions. As arbitrageurs and systematic strategies rapidly exploit identifiable patterns, any persistent edges tend to shrink. In practice, even real effects are easily swamped by noise, transaction costs, capacity limits and crowding. Only signals that survive stringent out-of-sample tests, robustness checks and economic-rationale screens should be trusted.

This explains why AI cannot easily be adopted in finance the same way it is used in other applications, where insights from training on larger data sets can be applied to new problems.

Instead, models must be carefully designed, validated, and constrained to ensure economic plausibility. It takes human experts to apply the technology effectively. Let's look at a few examples:

Example 1: Machine learning for large factor models in equities

Factor investing – building portfolios based on characteristics such as value, momentum, quality, or low volatility – has become a mainstream approach for stock selection. Although hundreds of factors exist, the small data problem limits quants to a few well-proven ones with economic rationale, typically combined in stable linear models. Yet few factors capture everything, and relationships rarely stay linear or constant, creating a trade-off between overfitting (losing stability & rationale) and underfitting (missing hidden factors / relationships).

This is where machine learning (ML) comes into play. ML can analyse many more factors than a human analyst could and model



nonlinear interactions and conditional effects across many dimensions.

We applied machine learning to a stock selection problem. Starting with around 100 well-established investment characteristics, we generated thousands of variations to build a more flexible 'structured' large factor model and compared it to a traditional linear model. Both performed similarly under realistic conditions but combining them yielded the strongest results, lifting the information ratio from 0.94 to 1.11.

AI can enhance more traditional factor models for stock selection, but the improvements are more evolutionary than revolutionary, and careful model engineering is required.

Example 2: LLM's and 'central bank speak'

In contrast to the small data problem entire models and financial forecasts face, it is often far easier to summarise unstructured data into input signals for models. LLMs are already trained on a massive corpora of financial and general language. Therefore, they can also analyse financial texts.

For example, central banks communicate regularly through speeches, press conferences, and minutes. In comparison to traditional approaches such as sentiment dictionaries, large LLMs can better capture nuance in tone, context, and forward-looking statements. They can interpret the Fed's (and other central banks') 'hawkish' or 'dovish' stance and how it's evolving over time. Such a 'Fed Speech Signal' can serve as powerful input into forecasts for asset

classes, equity risk premia, or currencies.

We applied LLMs to analyse central bank communications and generate long/short investment signals for sovereign bond futures across six countries for the time period January 2011 to April 2025. The results show that the LLM-based approach delivered a clear edge over traditional machine learning, achieving an annual return of 4.4% versus 2.7% for the traditional model. Although the LLM strategy exhibited slightly higher volatility, its superior information ratio (0.82 versus 0.59) highlights a stronger risk–return profile.

When both models were combined, performance improved further, delivering a 5.0% return and the highest information ratio of 1.0, underscoring the value of model diversification.

Example 3: Agentic models supporting the quant analyst

Beyond alpha generation, AI can also transform the workflow of quantitative analysts such as coding, testing models, cleaning data, and documenting results. This is where agentic AI systems – models that can reason step-by-step, interact with data, and generate executable code – are particularly promising.

Consider a quantitative analyst tasked with designing a new factor model. The process begins with idea generation – formulating a hypothesis such as combining earnings reports with stock price reactions to capture momentum effects.

Traditionally, implementation requires writing Python or R code, building data

pipelines, and validating regressions. The analyst must then test numerous variations, a process that can take hours of manual adjustments and fine-tuning before reaching a robust result.

An AI coding assistant can accelerate this process by generating code snippets, debugging errors, suggesting statistical tests, and even automating backtests. Instead of implementing every step, the analyst focuses on the conceptual design of the model, the interpretation of results, and the communication of insights to portfolio managers or clients.

This shift highlights a broader trend: AI reduces the friction of implementation, freeing human experts to spend more time on higher-order reasoning, creativity, and oversight.

Conclusion: AI as a tool, not a replacement

Artificial intelligence is transforming quantitative investing. Machine learning enhances factor models by uncovering new relationships and nonlinear effects, while large language models make unstructured data, such as central bank communications, usable for investment insights. Agentic AI tools streamline workflows, freeing quantitative analysts to focus on design and interpretation.

Challenges such as limited data, regime shifts, and the need for transparency and accountability ensure human oversight remains essential. Instead, AI marks the start of a new partnership, where technology and human judgment combine to advance systematic investing. ■



Dr Maximilian Stroh

CFA, Head of Research,
Quoniam Asset Management

SUMMARY

Because financial markets provide only limited and often unstable data, human insight remains indispensable.

AI enables powerful new approaches to factor modelling, but its impact should be seen as an evolution rather than an instant revolution.

Large language models can analyse and decode the subtle tone of Federal Reserve communications, revealing signals that were previously hidden.

Generative AI is transforming the way quantitative researchers work by dramatically increasing coding speed and overall research productivity.

AI is a game-changing ally. Those who adapt their workflows will lead the future of investing.