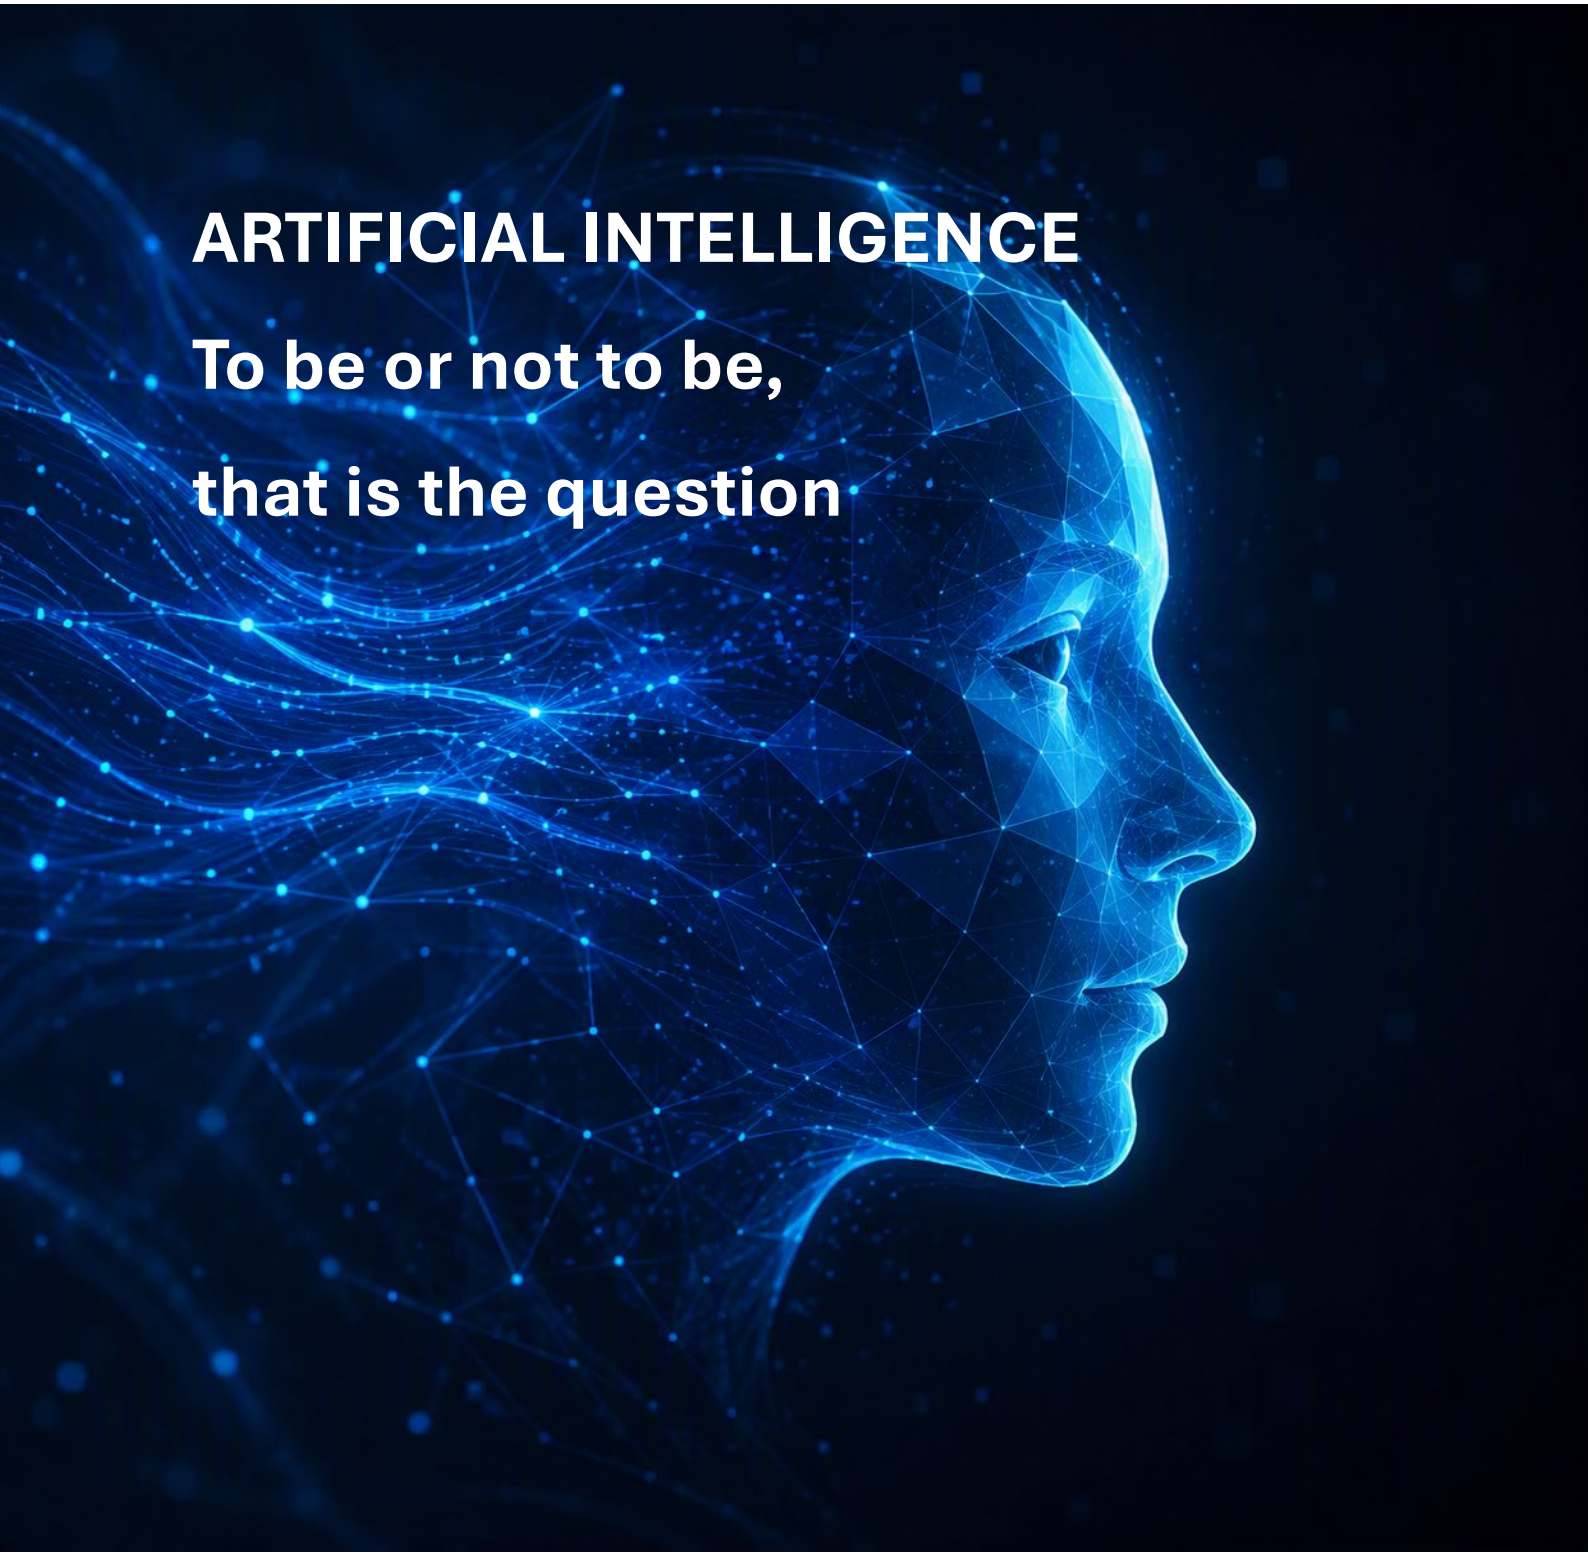
QUINTAINANALYTICS



ARTIFICIAL INTELLIGENCE

To be or not to be,

that is the question

September 2026

The term AI (Artificial Intelligence) is ubiquitous, with marketing departments finding new ways to make every product have it (my fridge has AI?), and investors also jumping on the bandwagon as this is the next new buzzword after dot com, block chain and crypto. In Capital Markets, it is appearing in everything from reconciliation software through to portfolio and research content applications. But is it such a great advance and how does it apply to this industry, or more specifically to behavioral analytics in Capital Markets?

This article summarizes some of the findings from our R&D that looks at a series of deep use cases in behavioral analytics and assesses a range of different AIs (including different model versions of the same large language model, "LLM"). Note that data security is also an important consideration, however this is not discussed here.

In short

AI has potential, but not applied to everything. Like most things in life, if it is too good to be true then it probably is. AI won't magically fix, or replace, everything and everyone but with very specific knowledge based application it can have a noticeable impact on work productivity and influence both revenue and costs.

Cost saver or revenue generator?

AI can be used to commoditize low value, often mundane and repetitive tasks, allowing savings to be made or funds re-allocated to alpha accretive initiatives. There is much in the press on job redundancies and purported cost savings due to use of AI (and in some cases this have been reversed due to going to far, highlighting that it isn't necessarily suitable for everything) which highlights the need for appropriate governance and clearly defined objectives.

AI can also be used to generate revenue, both directly and indirectly with specific targeting it can sift through large amounts of data that traditionally would take a significant number of human hours to do so. The smartly used AI can find opportunities, arbitrage similarly to some algorithms, and identify and highlight trends that can better optimize opportunities.

Both cases highlight the need for expert human oversight as AI is not (yet) the full package.

User driven usage (Generative AI)

A prompt on a website, or an AI assistant, is probably where most people reading this article will have seen or used AI (outside the use case of your home fridge!). Many of these are simple bots with no real AI, and react with limited, pre-programmed, information. Where AI is implemented, results can be mixed.

When AI is left to its own devices, it can head off in an unpredictable direction and end up somewhere far from where the user intended. It also can create information, cite imaginary sources or cases, and can lead the user astray. This is not the AI's fault, as it will seek the answer to the question in the order most statistically likely in any given LLM. An AI answer is a best guess based on the model, and what it determines to be the likely response based on input, as it does not inherently have any of the domain expertise. Focusing queries through parameterizing (via prompts) is key to more optimized responses, particularly if you can control the content of the data lake being analyzed. For example, viewing key engagement behavior, and associated trends over a period, or determining which city is best for roadshows with certain investors or clients, without the need to analyze multiple charts and pivot data.

Behind the scenes generation

Parameterization extends its usefulness outside of direct interaction with humans. For tasks where findings need to be generated on a large scale behind the scenes, or "server side", this can be used to sift through large amounts of data to produce targeted and relevant summaries and opportunities for inclusion in systems or reports. Traditionally this has required specific programming to produce code that can analyze the data either at point of generation or by running batch jobs, often taking many hours. Another approach is Agentic AI, which uses multiple AI agents and exhibits goal driven behaviors and adaptability to make decisions by performing multiple step tasks with minimal human oversight. AI can offer targeted analysis, faster, and without the need to invest large numbers of programming hours for specific analysis. There is still effort upfront through parameterization but this is factors less than traditional coding.

...but it is included in the current applications so doesn't cost anything?

AI costs can vary hugely, and are initially expressed in tokens, with one token being roughly four characters with consumers billed both for inputs (prompts) and outputs (responses), thus it can be quite hard to gauge likely costs when usage is unknown. Additionally the type of task can greatly affect token use (and thus cost); tasks such as analyzing or generating complex files, for example, large Excel documents or images, can

significantly increase the number of tokens used, and therefore incur costs more rapidly than simple chat-style text exchanges.

There have been several headlines recently highlighting where budgets have been consumed prematurely as usage of AI has spiraled. Longer questions, more analysis across a larger dataset and an expansive response will all cost more. Understanding and controlling cost is key, target the desired dataset and parameterize the input and output, otherwise costs will balloon. There are also engineering considerations, such as tenanting a model, that are not covered in this article.

No two AIs are the same

In our analysis we performed several tests using a fixed set of data against a baseline of expected output that had been generated via a coded approach. We judged the results on accuracy, speed and cost with performance varying substantially. Time and accuracy varied hugely between our test models and the cost was no guarantee of acceptable results. The range of accuracy was most surprising with some performing very poorly whilst others were extremely accurate, showing some considerable difference in how the LLMs process data, but also that some models can err due to being overpowered for what they are being asked to do.

Going forward

There is no one size fits all, and our tests show that the right model and targeting are key to reaping the full benefits of AI. Depending on what outcome needs to be achieved will determine, to a large extent, the best model for any given use case. By configuring how AI is used will drive the overall experience, and the associated cost, which is becoming increasingly important to the overall return on investment of implementing AI.

As to whether AI is to be or not to be, it is very much here and in being, but perhaps the more relevant question is not whether AI is used but how is it best used. Here at Quintain Analytics this is something we continue to research, and explore in how to use across our products and services, to be able to continue to provide the most value and best experience to our clients.

Next Steps

If you are interested in this article or our firm and want to know more about our products please drop us a line on ai@quintainanalytics.com or view on our website quintainanalytics.com.