AI IN GLOBAL FINANCIAL MARKETS

CAN AI EVER DELIVER ON ITS PROMISE?

A BRIEFING TO HELP YOU CUT THROUGH THE NOISE,
AND HAVE BETTER CONVERSATIONS ABOUT AI

FEBRUARY 2025

MODELOMNI

AI IN GLOBAL FINANCIAL MARKETS
M2502.01

SUMMARY

The surge in the trialing and application of AI-enabled systems, and the idea that there are very few areas of business that will remain untouched by AI, is having a real and tangible effect in global financial markets.

But AI is not easy to have a conversation about, for a number of reasons:

1. It raises different questions and challenges for different people, from those who regulate to those who work on the underlying science.
2. AI is not just one thing – it covers many different related ways of automating tasks, some old and some new.
3. Finally, not everything associated with AI will, or even could, deliver on the promises being made.

Modelomni is a leading practitioner of AI in global financial markets, and has prepared this AI book to cut through the noise about AI.

By doing so, we help prepare you for the conversations you will have as you adopt and master this potentially transformative technology.

We also share our opinion on why AI is not delivering the breakthroughs we need in global financial markets, and point out two things that we think institutions should do to realize the potential benefits.

CONTENTS

1. THE CHALLENGES THAT AI PRESENTS FOR ITS STAKEHOLDERS

AI is most simply defined as algorithms that can learn patterns from data at scale, and so mimic aspects of human intelligence. However, there are many ways of more precisely defining AI, and with good reason: it gets defined differently, by different people, for different purposes.

For many people, AI is synonymous with the relatively recent, but hugely influential, generative AI systems such as GPT: software applications, particularly in the form of interactive 'chatbots', that not only interpret data and give answers, but generate data in the form of text, images, videos and music.

This has secured a huge amount of industry investment and has brought AI to the attention of many more researchers, business makers and industries than at any point in the last 40 years of research and development. Consequently, the conversations being had about AI range across a huge number of issues and topics.

In this section, we look at what AI is from the perspective of three different types of people:

1. those who use and experience AI systems, and the regulators that look after their interests
2. those who build products and services with AI
3. and the scientists and engineers who work on the underlying science and algorithms

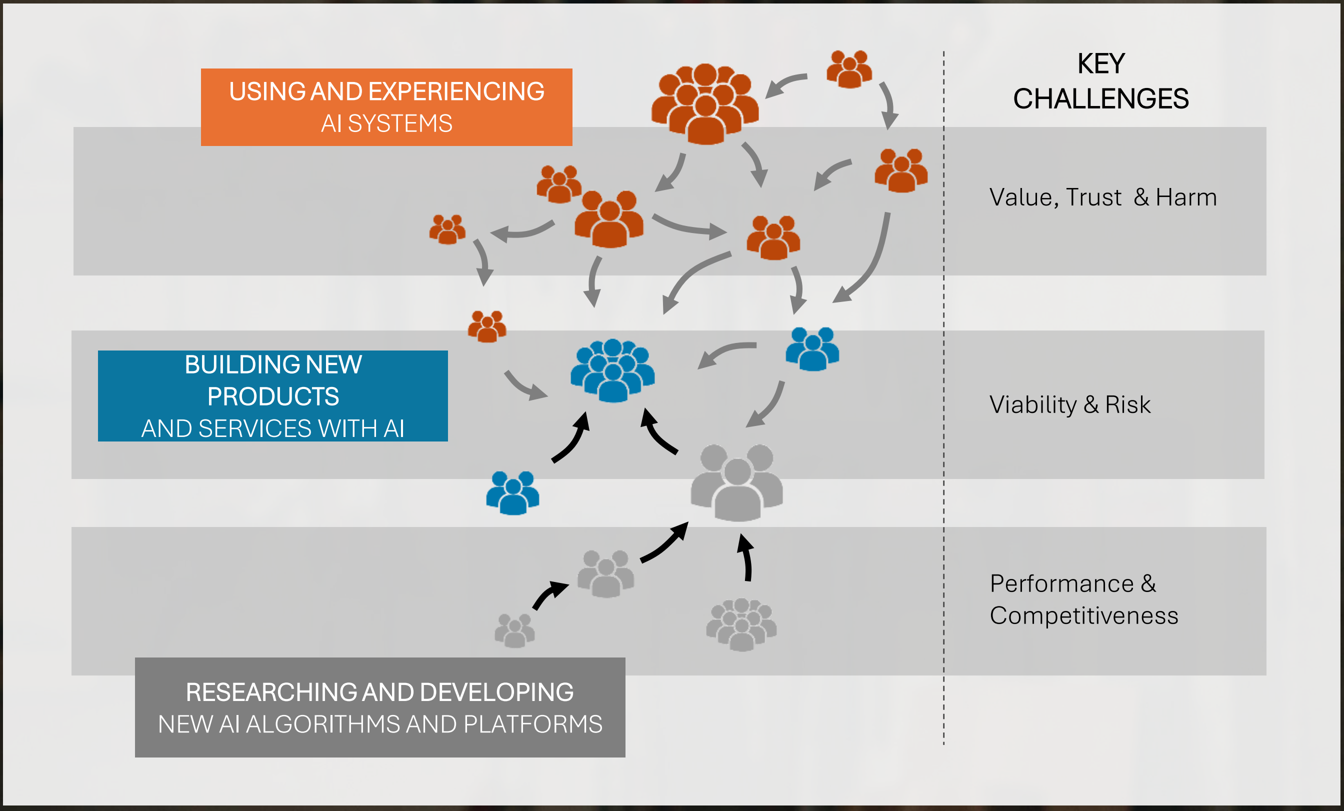

The resulting network of stakeholders with its flows of influence and resources goes a long way to illustrating why conversations about AI are not straightforward, involving as they do multiple stakeholders, motivations and perspectives.

USING AND EXPERIENCING AI SYSTEMS

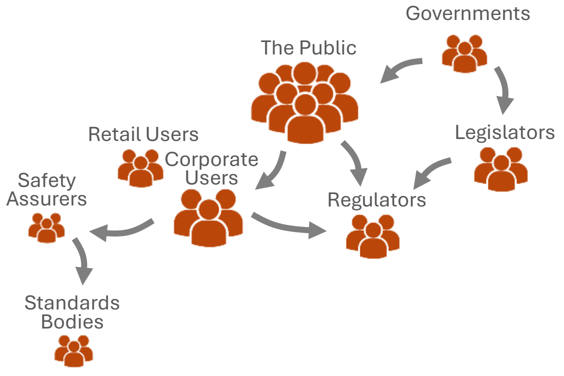

Does this AI-enabled product or service actually help me do something better, faster ,cheaper than other ways of doing the same thing, with no more risks?

There have never been more users of AI systems than today, whether a corporate procurement team using an AI agent to negotiate a long tail of lower-value contracts (Walmart using Pactum AI) or simply a plumber creating a quote for their next job more quickly than they could before (ChatGPT).

Predictions differ as to the extent that AI will support and replace processes and roles across industries, but one popular estimate is that it will impact or replace up to 40% of all jobs globally[1], with optimists suggesting that there will be a net increase of higher productivity work, rather than mass unemployment[2].

Whether a positive or negative scenario, the application of AI will be pervasive. Recommendation systems embedded in shopping web sites take natural language requests, and return preferred products and services. Algorithms embedded in medical devices constantly monitor a patient's health indicators, and decide on the amount and frequency of dosage of drugs. Workplace robots collaborate with humans – so-called 'co-bots' – on work that is too repetitive, or harmful to be done by humans alone.

In all of this work, users and buyers of AI products and services are simply looking for 'faster, better, cheaper', whether enabled by AI or not. They are not so interested in the technicalities of how an AI system is developed, but in the end effect of greater efficiency and productivity.

There are challenges, however. Regulators are concerned about the impact of services and products on people and society, and AI is no different. For these organizations, 'AI' is just another system that have the ability to cause harm through effects and actions in the real world. As these systems are being introduced, the U.S. Food and Drug Administration is assessing the reliability of AI-based image recognition in healthcare diagnosis, while the U.K.'s Financial Conduct Authority considers the role of AI in the processing of loan applications, and Ofcom, its regulator of communications services, considers discrimination and bias in the algorithms that drive hyperscale social media and news.

In a benign case, users of AI systems might simply be surprised by their interactions with AI-enabled systems, and suffer no harm other than an erosion of trust. Others, who are not even direct users or beneficiaries of a system, might discover more drastically that the self-driving features of cars, boats and drones can too easily make the wrong decisions, at the wrong time. And governments worry about the effect on the environment of AI energy consumption, the effect on manufacturers, on tax revenues and on communities and society.

[1] An International Monetary Fund article from January 2024
[2] A World Economic Forum article from February 2024

KEY CHALLENGES

- Understanding harm and risk arising from inconsistent, unpredictable systems
- The complexity, opaqueness and trustworthiness of algorithms
- The ethics and legality of the transfer of decision-making from people to machines
- The cost-benefit trade-offs for society that are being made by the providers of AI-based systems
- The creation of monopolies and competition in the development of foundation models
- The use of data, whether personal or copyrighted, and associated privacy concerns

BUILDING PRODUCTS AND SERVICES WITH AI

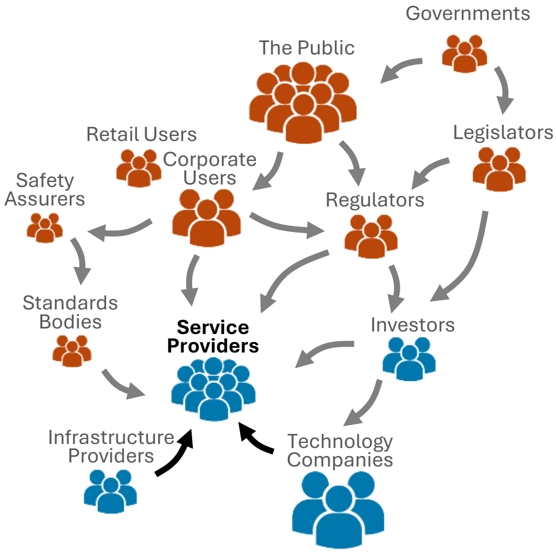

Do my new models perform in the real world, and is their effect acceptable to the public?

Is the business process they are wrapped in commercially viable ?

The builders of AI-enabled products and services, whether a start-up targeting a market need, or a corporate team tackling a specific use need, need to master the two fundamental reasons why AI systems are different, if they are to effectively take advantage of the opportunities afforded by AI technologies:

First, they are developed not by software engineers writing rules and instructions, but in part from algorithms that learn by example from data. Consequently the selection of data, which is necessarily imprecise, becomes a key activity, affecting the ultimate behaviour of the system. This means that systems can be developed, in principle, faster and more efficiently than before, with more flexibility than 'coded' systems.

Second, the data is used to create a model that can then respond to new data, meaning that the system can appear to take on intelligent decision-making roles, and be deployed into situations that shift the balance of human and machine. This means that AI systems can take on more complex tasks that augment or replace people, doing these tasks more consistently, more accurately, and faster, than any human could. Ultimately, different business models and opportunities become possible.

Done well, this can result in an 'intelligence capability' autonomous behaviour with apparently 'intelligent' decision making and reasoning, that can generate results based on the evaluation of huge amounts of information, including novel outputs generated given never-seen-before data, and the ability to operate, with varying degrees, independently of people.

Done not so well, this can deliver systems that are unpredictable in their performance, expensive to maintain, and a liability when used in anything other than a close approximation to their original proof of concept demonstrations.

Meanwhile, the underlying platforms for building AI systems, increasingly provided by hyperscale technology companies, and the infrastructure providers that provide the computational power to run processing-intensive software, have their own competitive pressures to establish dominance, consuming huge amounts of capital, but also attracting huge interest from the investor community, with the associated pressure on return from investment.

- Demonstrating value through the consistency, efficiency of products and services
- Establishing usability and acceptability
- Scaling and commercial risk & viability

RESEARCHING AND DEVELOPING NEW AI ALGORITHMS AND PLATFORMS

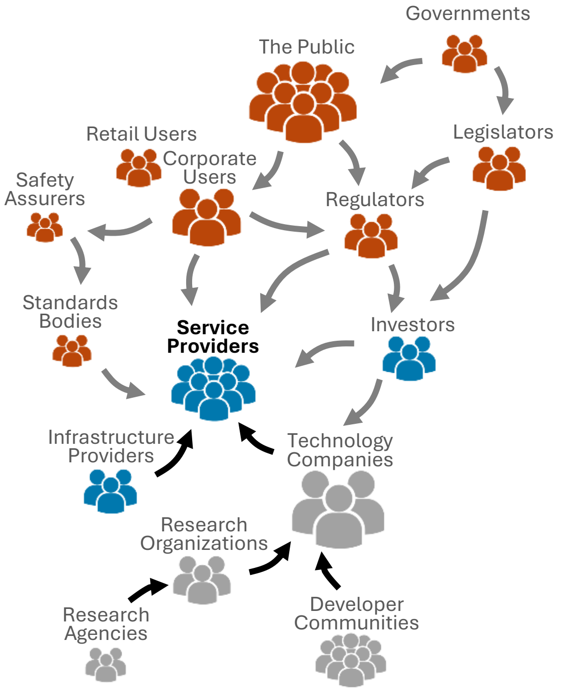

Is this new way of developing functionality better than others?
If so, do I understand why?
Is the development method feasible in terms of data volume and computational cost?

The development of AI has its roots in academia reaching back decades, and research agencies of national governments, and the academic organizations that they support, still drive much of the development of the fundamental data science and algorithms behind AI technologies.

However, the field is being increasingly dominated by the technology companies, with the cash and the motivation to establish their particular approach to machine learning as being the key to this next industrial and societal revolution, justifying investment against a potentially dominant market position in the future. Since the current generation of technologies is largely predicated on scale in order to achieve its effects, the barrier to entry in terms of developing the fundamental platforms and foundation models is currently high, with these companies spending amounts beyond the reach of many governments.

Academics, meantime, in response to needs from those who use and experience these systems, are also studying ways to evaluate and benchmark AI technologies. What are the definitive tasks and scenarios in which both performance and safety can be assessed? What guidance can be given to those who build with these technologies such that they don't expose unacceptable risk And moving from computer science into philosophy, jurisprudence and other disciplines, what constitutes acceptability?

- Ways to improve performance with lower cost in terms of data and resources
- Methods for evaluating and benchmarking performance
- Recruiting and retaining talent at the cutting edge
- Securing competitive advantage and sustaining corporate valuation

2. THE DIFFERENT TYPES OF AI SYSTEM

What all stakeholders seem to have in common is a belief that AI systems have huge potential to transform work, re-shape current boundaries between technology and humans, and change the way that businesses and markets operate.   Progress towards realizing this goal is being realized through the application of a diverse number of approaches and technologies. This adds another layer of complication to the conversations about AI.

There are conventionally two different ways that people think about this diversity – one is to do with the capability the system exhibits, based on how it is built, and the other is to do with how it is being deployed.

IS AI ABOUT PATTERNS IN DATA, OR ABOUT REASONING?

Earliest attempts to build AI, from the 1950s onwards, have followed two different paths, and have had differing success as underlying implementations improved:

1. all about patterns: mimicking how a brain works, with a neural network made up of many individual neurons, processing signals, generating results, and developing 'weights' that reflect patterns in the signals. This goes along with a tendency to be complex, opaque and appear to be non-deterministic in its behaviour
2. all about the meaning: mimicking how human reasoning works, with symbols representing meaningful concepts, connected together to represent some area of understanding. These are more likely to be understandable and explainable, but are harder to assemble and maintain

The first of these has experienced huge success in the last 20 years, in particular with the breakthrough of the layered deep learning approach that enabled a step-change in image recognition. The second is experiencing a resurgence as some foresee the limits of what is possible through patterns in data alone, with further variants such as causal modelling, and hybrid approaches, showing promise amidst the debate on where exactly the line lies between the two approaches.

HOW MUCH AUTONOMY IS BEING GRANTED TO THE SYSTEM?

The majority of AI systems, once packaged up as a product or service, provide a useful service, with caveats and guardrails around their usage – image recognition, or recommendations. The same underlying mechanisms can be used, however, to drive a decision loop in which the recommendation is connected to action, giving autonomous decision-making behaviour.

This commonality means we can be talking about AI as the algorithm behind the promotion of news articles in social media, at the same time as we are considering a system that is sent out into underwater operations to seek out and destroy what it thinks might be mines.

In each case there can be debate over responsibility and liabilities, and the transfer of decision-making from humans to machines. However, in the former case there is direct control over the actions of the AI system by people – the developers of the social media algorithm who are monitoring their system and adjusting how it behaves. In the second case, the system may well be 'beyond line of sight', and empowered to take actions with a much more obvious and immediate effect in the physical world.

A further source of confusion is that not all autonomous systems need to be AI-enabled. Models and functions created by hand or with other tools can be granted, once deployed, as much autonomy as those built using a learning algorithm, though likely to be more understandable, and less capable.

In this section we consider the common process that underlies machine learning, and give some examples from other industries.

THE COMMON PROCESS FOR BUILDING AI

Despite the diversity, there is fundamentally a common process that needs to happen for any AI system, even of the 'reasoning rather than patterns' type, to be successfully put into operation.

1. PROBLEM

There's a problem – a task that needs done, actions that need taken, decisions to be made, and faster, more consistently, at greater scale than before. Clarity on this is essential in order to effectively direct the AI building process.

2. DATA

The task is defined in terms of goals that can be quantified – what does success look like? – and data is sourced that gives many examples of things that are related to achieving, or not achieving, that goal.

3. MODEL LEARNING

A model is developed by tools and algorithm that can assemble source data into functions that will take new data as an input and generate appropriate results or actions that align to the goal.

4. DEPLOYMENT

The model is deployed as part of a software application, consuming data and generate useful results faster, more consistently, more accurately and at greater scale than people could, or than systems programmed in a traditional way could.

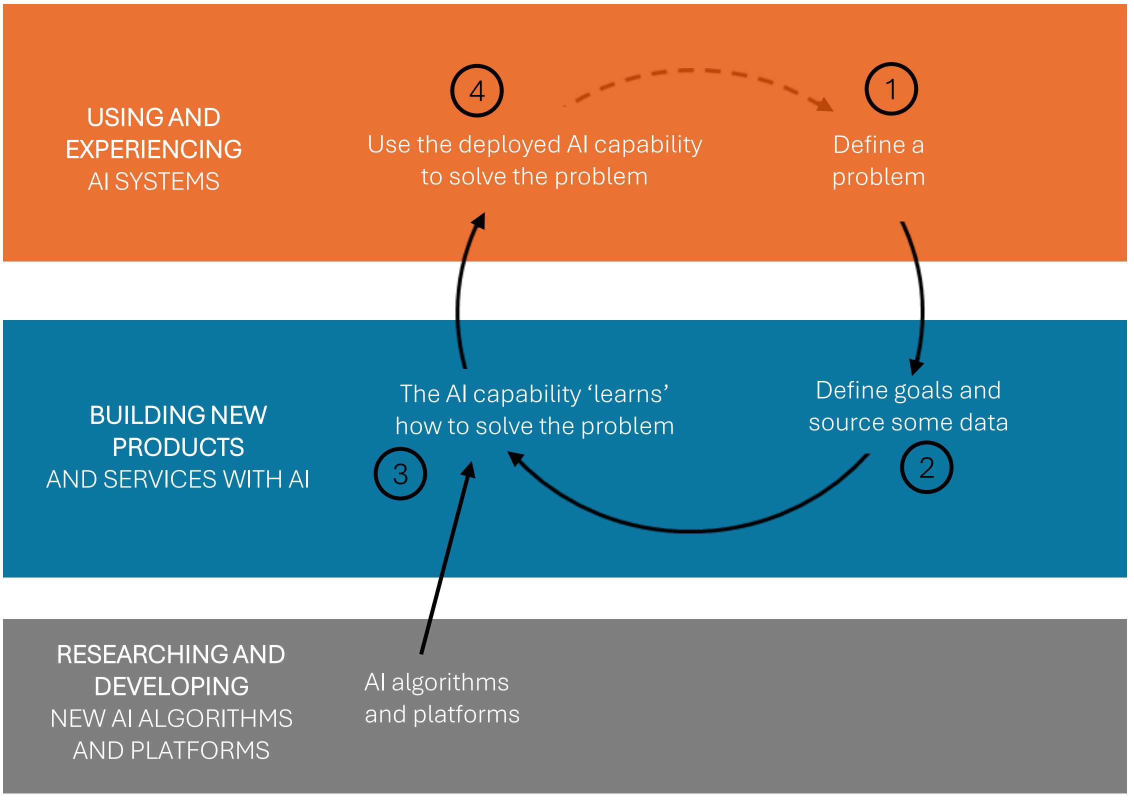

EXAMPLES FROM DIFFERENT DOMAINS

	An example of a service provider	1 PROBLEM	2 DATA
Cancer diagnosis cancer	Paige, developing diagnostic products and services for clinical pathologists	Identify cancer more reliably, with greater capacity	Chest X-ray images as data
Fraud detection	Swift, a provider of anomaly detection services to banks	Financial loss through fraud going undetected and unprevented	Banking transaction data
E-commerce recommendations	clerk.io, developing a search and recommendation engine that retailers can integrate into their shoppers' online experience	Customers visiting a site but not making a purchase	Online shopping data
Customer service	Intercom, a provider of AI agents that support people using a web site	Web site users becoming disengaged and leaving a site	Data from previous conversations, and everything on the internet
Weather forecasting	DeepMind, the AI research division of Alphabet, building models to solve large-scale challenges, including weather forecasting	Inaccurate wind forecasts leading to sub-optimal wind turbine energy generation	Recorded and modelled weather data
Embedded robotics	Boston Dynamics, a developer of industrial mobile robots	Getting a robot with limbs to walk	Data on the position of a robot's limbs following an action

In each case, the AI approach to solving the problem augments or replaces either a manual process, or one supported by models that are in large part created and maintained by hand.

3 LEARNING	4 DEPLOYMENT	The type of machine learning used
Learn which patterns are associated with a cancer diagnosis	Give diagnoses for thousands of new X-rays	Supervised learning, based on 'labelled' examples of cancerous or benign images
Learn which patterns are associated with fraudulent activity	Prevent or allow specific transactions	Unsupervised learning, grouping behaviours into 'normal' and 'not normal'
Learn which patterns are associated with a purchase	Recommend a product to a particular customer	Supervised and unsupervised learning, matching customer characteristics to product characteristics
Learn which patterns are associated with a target word	Output the next best word, over and over, to generate responses to a customer	Self-supervised learning, creating large language models (LLMs) powering AI chatbots
Learn which patterns are associated with the next time-period wind conditions	Output a wind condition forecast based on current conditions	Self-supervised learning, creating a 'large weather model'
Learn which patterns are associated with walking vs falling over	Take actions that result in walking rather than falling over	Reinforcement learning, using the effect of actions as its 'training' data

Although there is a predominant type of machine learning for each problem, AI techniques are increasingly used together.

In particular self- and semi-supervised learning, and the now almost ubiquitous generative AI / general purpose AI / large language models associate with these types of machine learning, can be adapted to many problems, and used to augment other techniques.

3. WHAT IT MEANS FOR A MACHINE TO LEARN

The model learning process is what most clearly distinguishes AI from other software systems – the way that the functionality of the system is developed.

Traditionally, software systems are developed by programming – setting out the instructions and rules that determine what response is given when the system is given some data. AI systems are developed differently. The end product is still a function – given data, it returns an answer. But the development technique is based on an algorithm that builds that function - given a target or a goal, the algorithm takes in data at scale and builds (or 'learns') what is typically called a model, rather than a function, that will give a useful answer.

Many of these development techniques have long histories, but have only achieved success, from the early 2000s, given sufficient data and sufficient processing power. This success has led to a huge variety in the detail of the techniques, drawing on statistics, mathematics and computer science, with different approaches to hyperparameters, optimizations and technical architectures, all from different camps or movements, with claimed advantages in terms of accuracy, reliability or performance.

In this section, we look at the four main types of development techniques for AI, categorized according to the degree of support, or 'supervision', that is given to the learning algorithm that is used.

SUPERVISED LEARNING

A landmark technique popularized in the early 2010s by the step-change in performance that it could achieve in image recognition tasks.

The idea of supervised learning is that the learning process is 'supervised' through the provision of training data – data that has been 'labelled' to show the relationship between the data and the kind of answer that is associated with that data. The model-building algorithm takes this data, at scale, and attempts to generalize the relationship, such that new data, not seen before, can produce the right sort of answer.

In our 'diagnosing cancer' example, a training set of X-rays is prepared, each associated with the result that an experienced clinician would have given based on the image.

In its simpler forms, the resulting model is similar to traditional statistical techniques such as linear regression – a mapping between the 'answer', and the values of another data set. This leads some saying that simple AI is just 'automated statistics'.

In its more complex forms, such as 'deep learning', something magical happens – as parameters are set through different layers of the model, small features seem to be recognised, and assembled into more significant features to give 'intelligent' results.

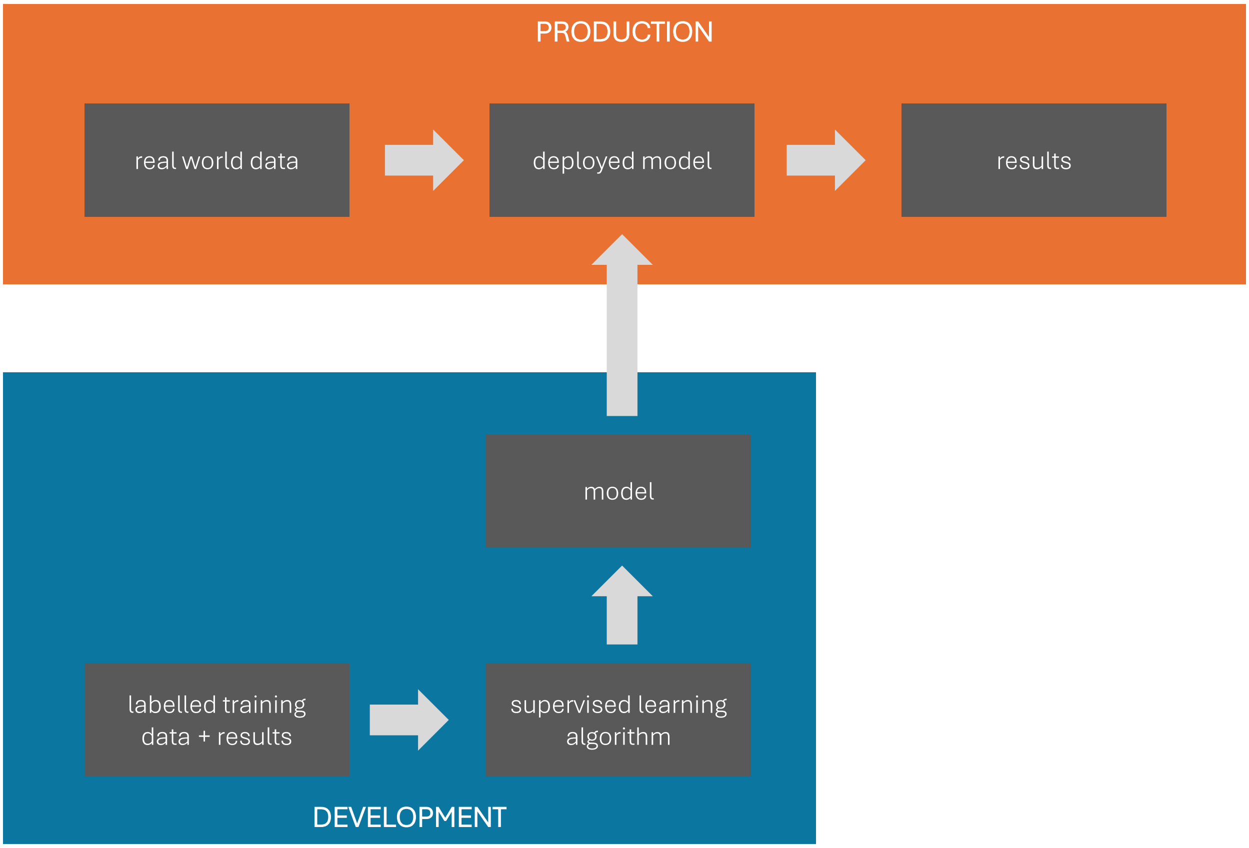

KEY CHALLENGES	• There is a cost to labelling enough data • A drive for accuracy can lead to 'overfitting', in which a model works well but only if the real world data is very similar to the training data • There can be a fragility to the resulting models – a small change to data can lead to wrong answers • The data can contain 'hidden' biases and issues, which when combined with a model that is almost impossible to understand, leads to trust and dependability issues

UNSUPERVISED LEARNING

The term ‘unsupervised learning’ typically refers to a range of more mature data science techniques that don’t require data to be labelled in order to come up with a useful model, and you don’t necessarily have to define what an ‘answer’ might look like, though these techniques perform better given some hints about what this might be. These algorithms are designed to simply take in data, and return useful insights.

The most commonly used unsupervised technique is ‘clustering and classification’. In our fraud detection example, data points associated with each transaction – the amount, the vendor, the account, the account balance, the customer’s location, and so on – are given to an algorithm that assesses how similar these transactions are, clustering them into groups. The intention in this case is to end up with a group that is ‘not normal’ in some way. This analysis, repeated regularly, might be all that is required to create results that are useful. Equally the clustering model might be deployed into production to classify new transactions in real-time.

There are many other techniques, all best seen as tools that tell you something you don’t know about your data – discovering correlations in data that might be too hard for people or conventional statistical techniques to discover.

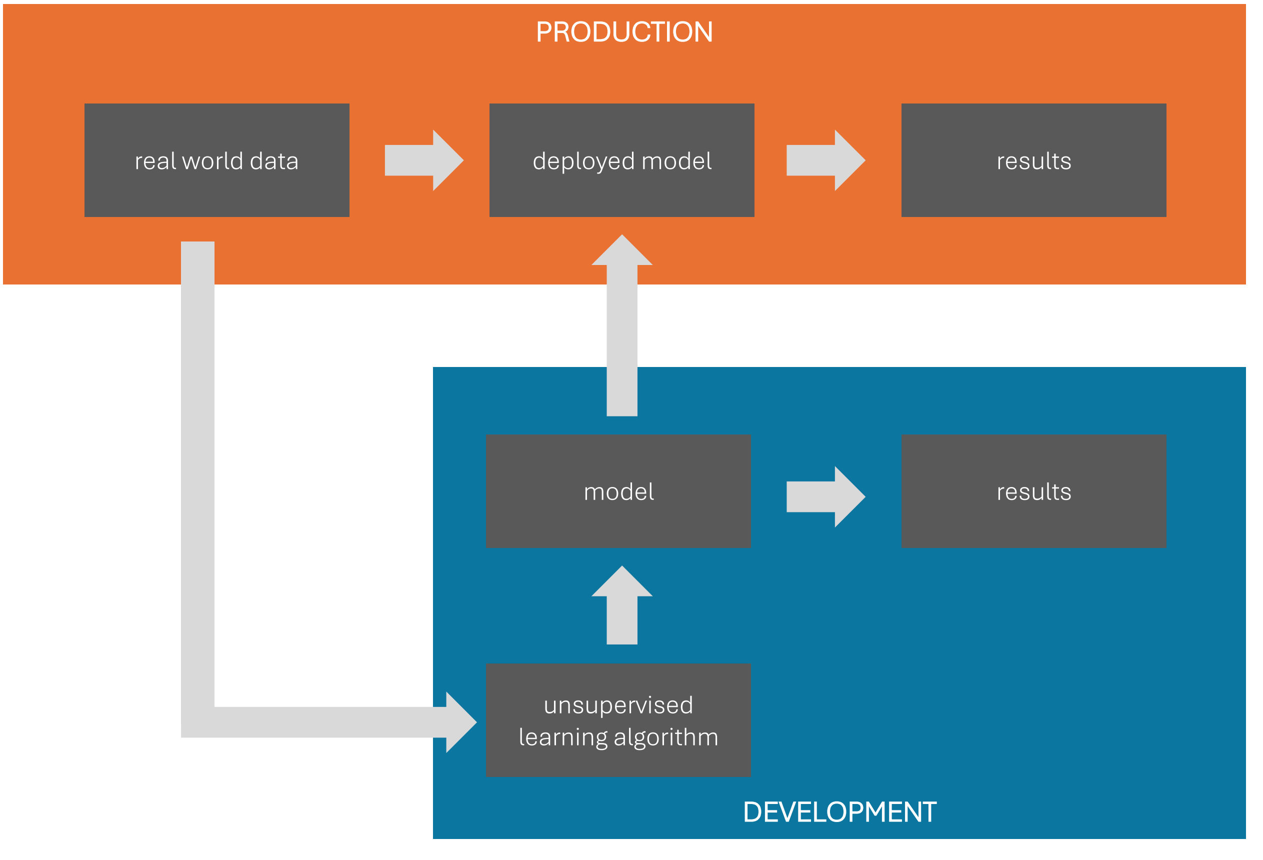

- May require a considerable amount of expert input to interpret results – a consequence of the lack of direction ‘up front’
- Best seen as a tool in the hands of data scientists
- Has a wide range of applications across different domains, but provides a very specific capability to do with giving insights into anomalies and variation in data

SELF- AND SEMI-SUPERVISED LEARNING

These, more than any other set of techniques, have powered the recent rise in interest, investment and usage of AI. Most distinctively these techniques underpin 'generative AI'. Unlike other learning techniques, the resulting model can take new data and generate data that is similar in some useful way, rather than a single 'result'. The original examples, appearing in the early 2020s, were to do with natural language. The development process involved exposed an algorithm to massive quantities of natural language texts, which would be encoded into numbers that capture the similarity of one word to another. The model is then able to take a new phrase, and generate the most likely next word. Combined with supporting techniques, and operating at scale, his leads to, for example, a better-than-human ability to summarize a piece of text into its essential points, or, applied to media other than language, generate a video based on a description of what that video should contain, and to do this iteratively in what seems like a conversation. The techniques can also be applied into other areas – generating future weather patterns or and the range of possible applications seems endless.

A huge industry around how to improve accuracy by for example using specific templates for data inputs ('prompt engineering') and additional data sources (for example RAG, or retrieval augmented generation) to reduce 'hallucinations', or incorrect outputs. This is a challenge - since the model has no understanding of the difference between 'fact' and fiction', but is just generating a plausible shape of results, then as far as it is concerned, it's all hallucination – a distorted version of the historical data that it's been exposed to.

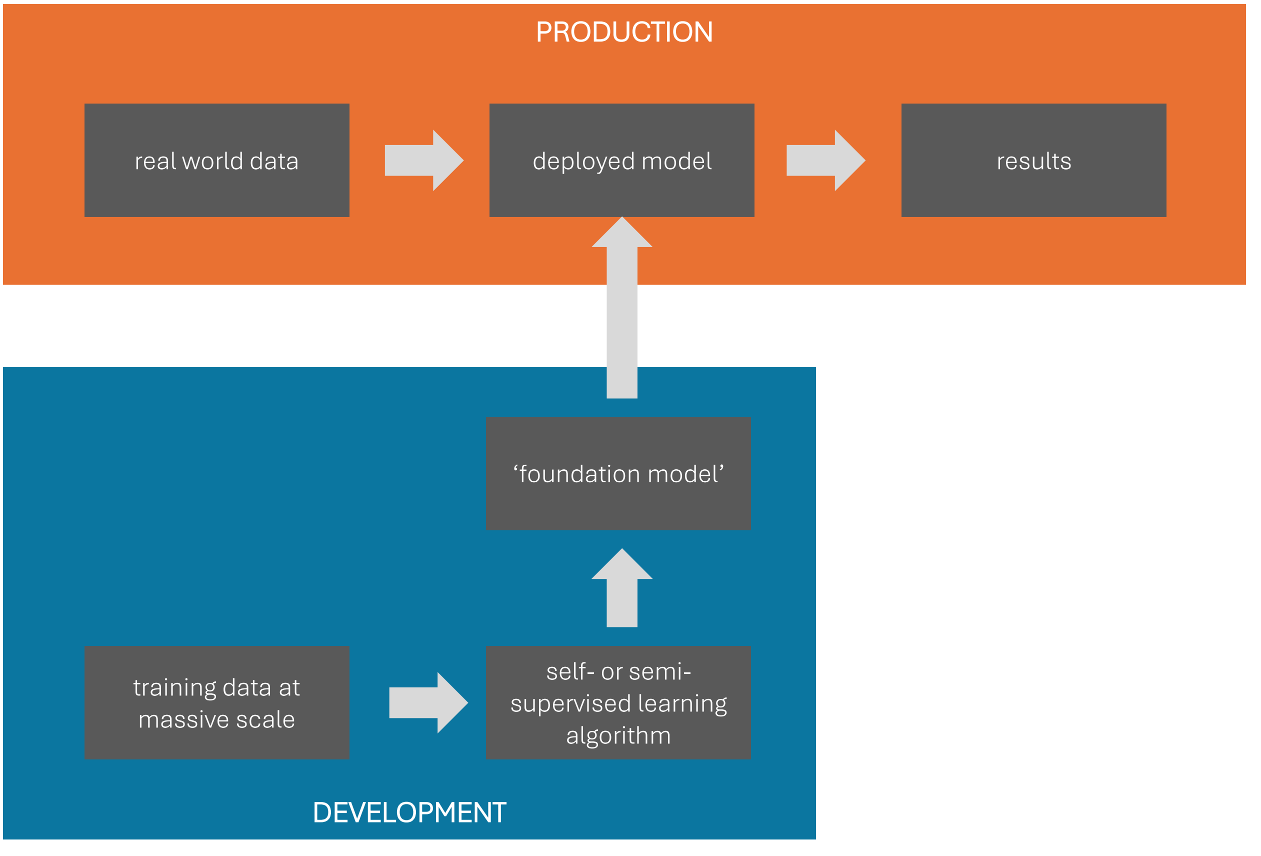

- The accuracy of results, and our limited ability to automatically test or assure these (without resorting to self-supervised learning!) - these models are not typically used in critical business situations
- The massive cost of the development process (borne by a small number of tech companies), and so ultimately the cost of using their models
- The limits of what can be achieved through further scaling – a matter of debate, but some would say that the technique has a performance ceiling that will require hybrid, or completely different, approaches, in order to make progress

REINFORCEMENT LEARNING

These learning algorithms are amongst the most interesting of AI techniques, in that rather than the 'statistics' or 'arithmetic' at massive scale, these seem closer to what we might understand as learning – given a specific goal, learning by trial and error what actions it takes best achieve that goal.

In our robot learning to walk example, the data that is required as an input to the learning algorithm is acquired from sensors monitoring the effect of actions being taken that have been suggested by the model. Working out what the right action to take, in terms of positioning a robotic limb to achieve forward motion, would be hard to define programmatically, and also hard to compile labelled data for. Defining what the goal is, however, is more straightforward, as is assessing the effect of the action against the goal. While running this process in a simulated world to develop a model is the norm, in principle the process can also be run in production, in so-called 'on-line learning'.

Other techniques are conceptually similar, for example evolutionary computing, which makes changes to the parameters of a population of models, observes the results that these generate, and then progresses to further rounds of 'evolution' with those that perform the best.

These techniques can work to multiple goals, across different time horizons (eg sacrifice short term 'losses' for longer term gains), and can be blended with varying degrees of guidance from human experts as well as other techniques such as search, planning and other AI techniques to solve complex, real-world problems.

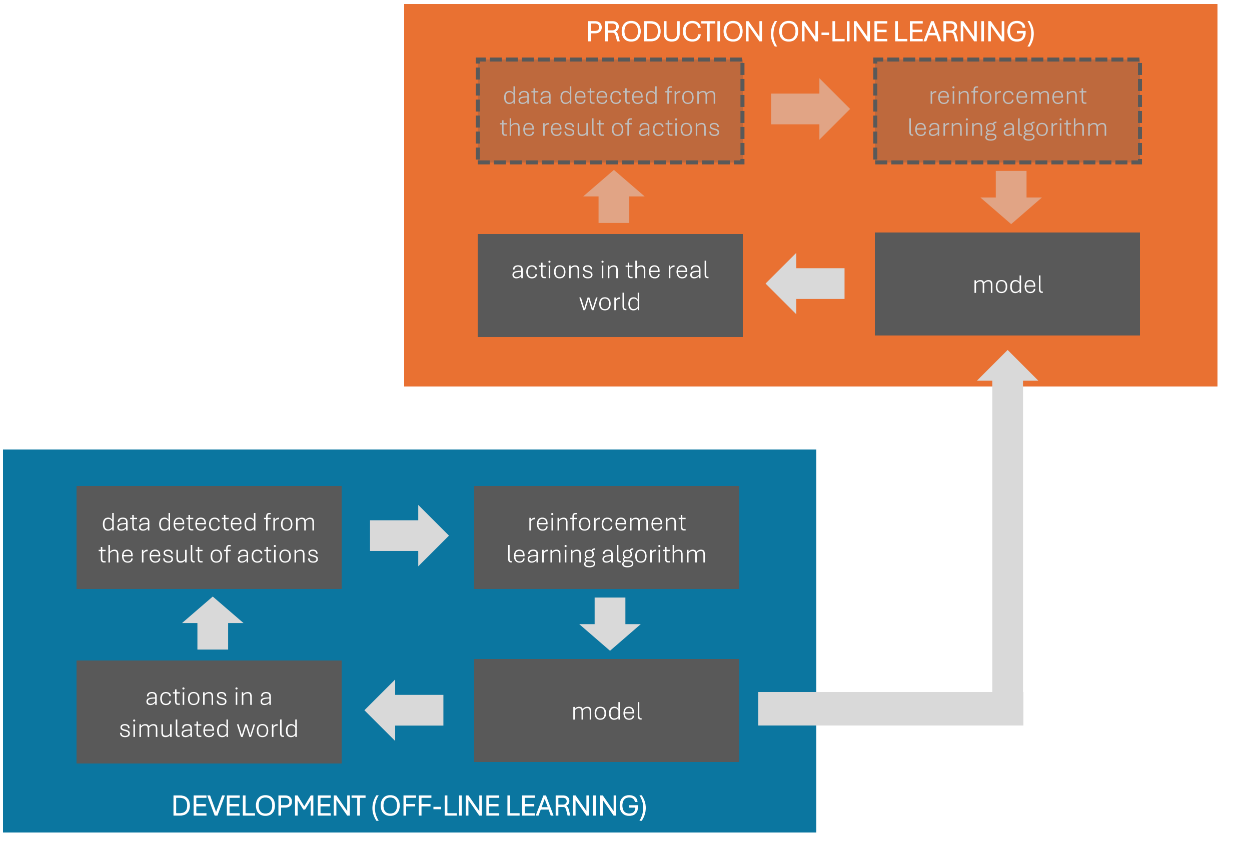

- A lot of assumptions are made in set rules for how the learning algorithm should behave (hyperparameters), and these can have a large effect on performance
- As with supervised learning, the resulting model can be hard to understand (an explainability issue)
- As with self- and semi-supervised learning, the results can be unpredictable

4. AI IN GLOBAL FINANCIAL MARKETS

In some ways, global financial markets is an industry that will be transformed by AI just like any other in the pursuit of productivity and growth, affecting many different types of organization, and 'core competence' processes as well as 'back office' processes:

- shifts in roles and ways of working as generative AI tools and agents answer questions; summarize and generate documents and other kinds of content; and increasingly take on end-to-end processes that are rich in information, such as customer service, IT management and software development, legal and finance, and regulatory reporting
- greater speed and efficiency in how a business stimulates demand for its services and goods, optimizes its supply chains and partnering processes to procure what it needs, and organizes its workloads and schedules through better prediction and optimization
- a renewed focus on curating the knowledge and data that differentiates a business, and that can be used to train business-specific models and agents that provide a competitive edge

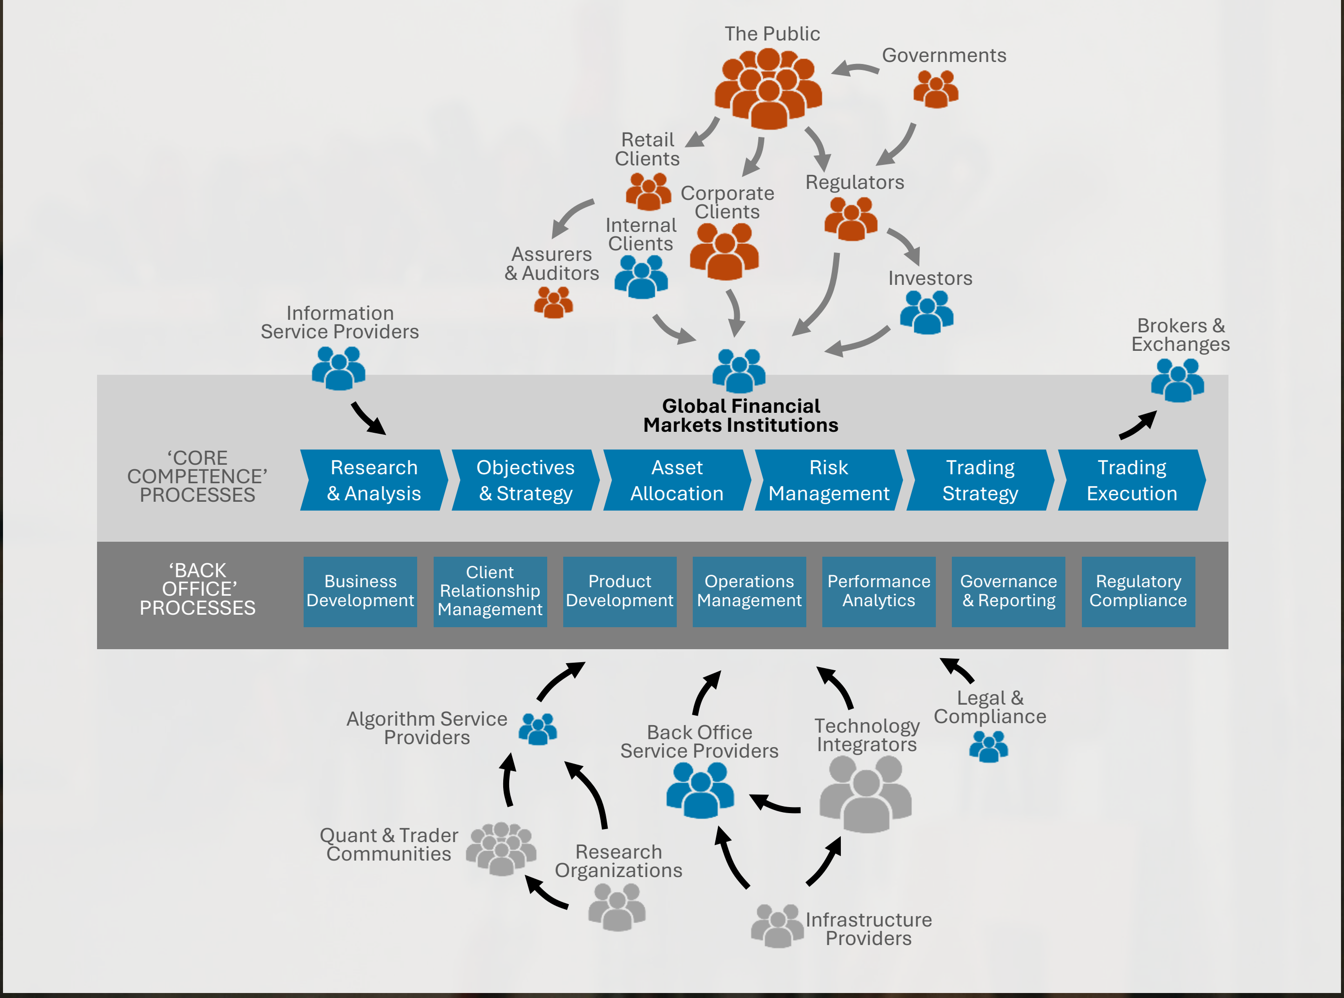

However, global financial markets is a particularly promising area for AI. Economists, even before the advent of computing, have used mathematical and statistical models to analyze data to explain and predict market activity and risk as they set strategies, allocate assets, consider how to measure and mitigate risk, and execute trades that realize the strategy.

There is competitive advantage to be gained in decision processes with speed, accuracy and consistency. Since financial markets are rich in data, it's not surprising that data science and AI techniques have long been trialed and applied.

In this section we take a step back in time to chart the progress of AI in global financial markets, helping to put current AI development in context.

1980s	THE BEGINNINGS OF AI IN GLOBAL FINANCIAL MARKETS
1990s	THE RISE OF ALGORITHMIC TRADING AND QUANTITATIVE MODELS
2000s	FIRST SUCCESSES FOR MACHINE LEARNING IN TRADING AND TEXTUAL DATA ANALYSIS
2010s	ROBO-ADVISORS, EXPERIMENTS WITH AUTONOMOUS TRADING, AND MAINSTREAM APPLICATIONS OF DEEP LEARNING
2020s	AI REACHES PEAK HYPE WITH SELF-SUPERVISED LEARNING, BUT MATURE SUCCESS STORIES EMERGE

1980s THE BEGINNINGS OF AI IN GLOBAL FINANCIAL MARKETS

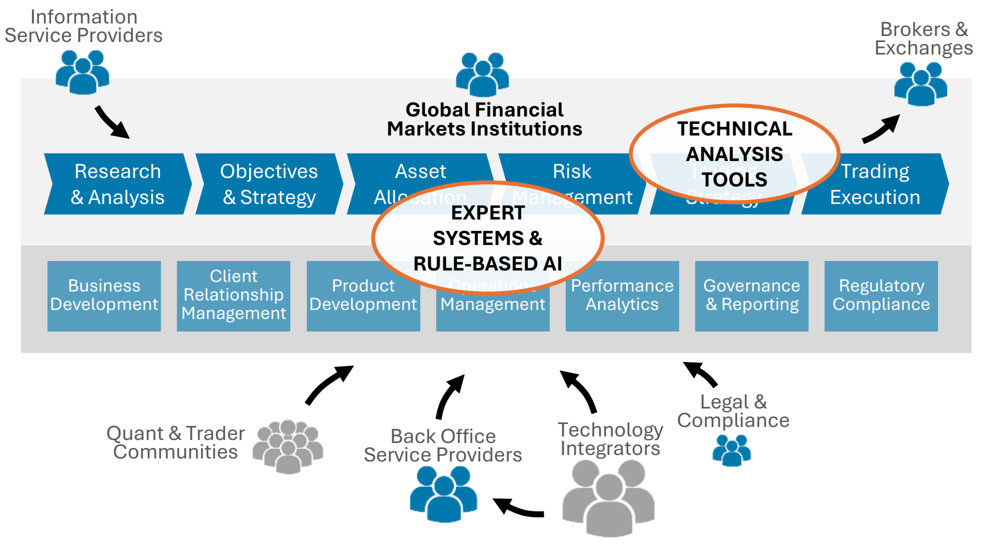

Prior to the 1980s, computing power in global financial markets, and in financial services more generally, was directed at the main transactional processes through mainframe applications for processes such as settlements and accounting. The more specialized processes such as stock selection and trading were supported by techniques with their roots in mathematics and statistics, and supported by simple tools such as calculators and charts that could be annotated with trends and thresholds following simple rules.

In the 1980s, we see the beginnings of AI in the development of rule-based systems for back office tasks, and better automated support for statistical techniques, in particular for time-series market data.

EXPERT SYSTEMS & RULE-BASED AI

Financial institutions of all types were experimenting with rule-based expert systems. For example, American Express developed the Authorizer's Assistant, an early AI-based system for credit approval and fraud detection.

TECHNICAL ANALYSIS TOOLS

AI-driven technical analysis tools were limited but began appearing in large banks like Goldman Sachs and Morgan Stanley, who used early statistical models for understanding market behaviours and setting trading strategy. These were based on concepts to explain and predict market, such as trend and resistance that had been developed throughout the 20th Century.

1990s THE RISE OF ALGORITHMIC TRADING AND QUANTITATIVE MODELS

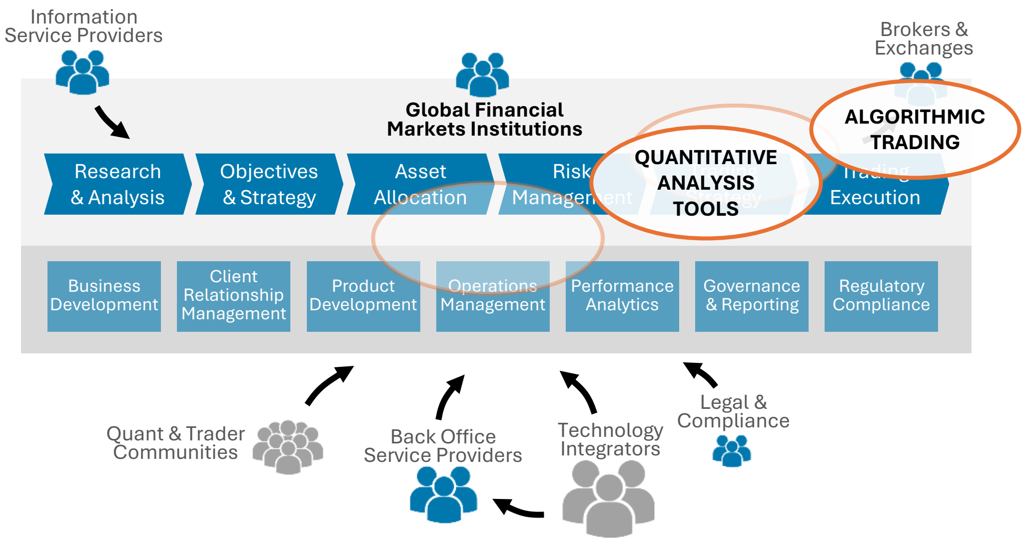

In the 1990s, the greater availability of data, connectivity and computing power means that the automation of statistical techniques with connectivity through to execution enables some major advances.

QUANTITATIVE ANALYSIS TOOLS

Firms like Renaissance Technologies popularized quantitative trading with the use of mathematical and statistical models that can automated the technical analysis generation of signals for trading. Unsupervised learning techniques such as classification are being applied both to stock selection and trading strategy as well as risk management through the identification of correlations across portfolios, as at Citadel.

ALGORITHMIC TRADING

Goldman Sachs was among the early pioneers of algorithmic trading. By the late 1990s, Deutsche Bank, Merrill Lynch, and Barclays developed proprietary algorithmic trading platforms to execute orders, with others such as QuantLab starting to take advantage of timing-based arbitrage opportunities.

2000s FIRST SUCCESSES FOR MACHINE LEARNING IN TRADING AND TEXTUAL DATA ANALYSIS

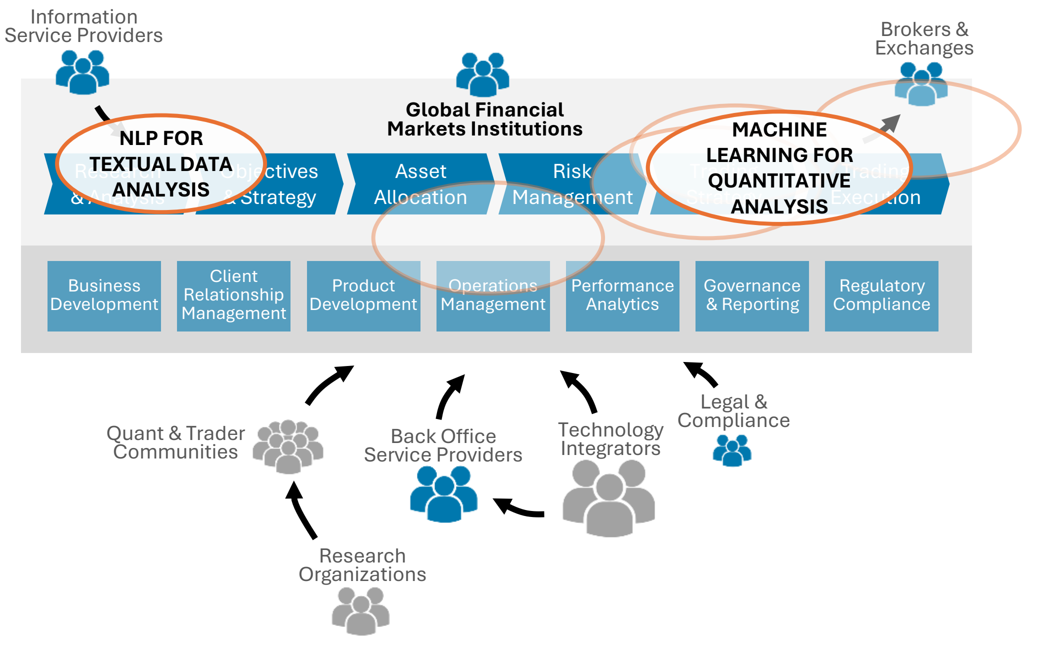

In the 2000s, the first experiments in unsupervised techniques for natural language processing, and supervised techniques to build better models, start to make it into production usage, supporting by a growing amount of research into AI techniques now enabled by the reducing cost of computing power.

NLP FOR TEXTUAL DATA ANALYSIS

Hedge funds like Point72 and Bridgewater Associates started incorporating natural language processing (NLP) tools to analyze news reports, earnings call transcripts, and other textual data to forecast market sentiment and adjust trading strategies accordingly. These are based on unsupervised 'brute force' statistical techniques that pre-date the era of generative AI.

MACHINE LEARNING FOR QUANTITATIVE ANALYSIS

Two Sigma, a quantitative hedge fund founded in 2001, applied machine learning and AI to financial data. By this point, AQR Capital Management, for example, had also embraced AI-driven quantitative strategies for portfolio management. These see the early application of supervised learning in financial markets, operating at limited scale with regards to data, and effectively implementing simpler statistical techniques such as linear regression.

2010s THE DEEP LEARNING REVOLUTION

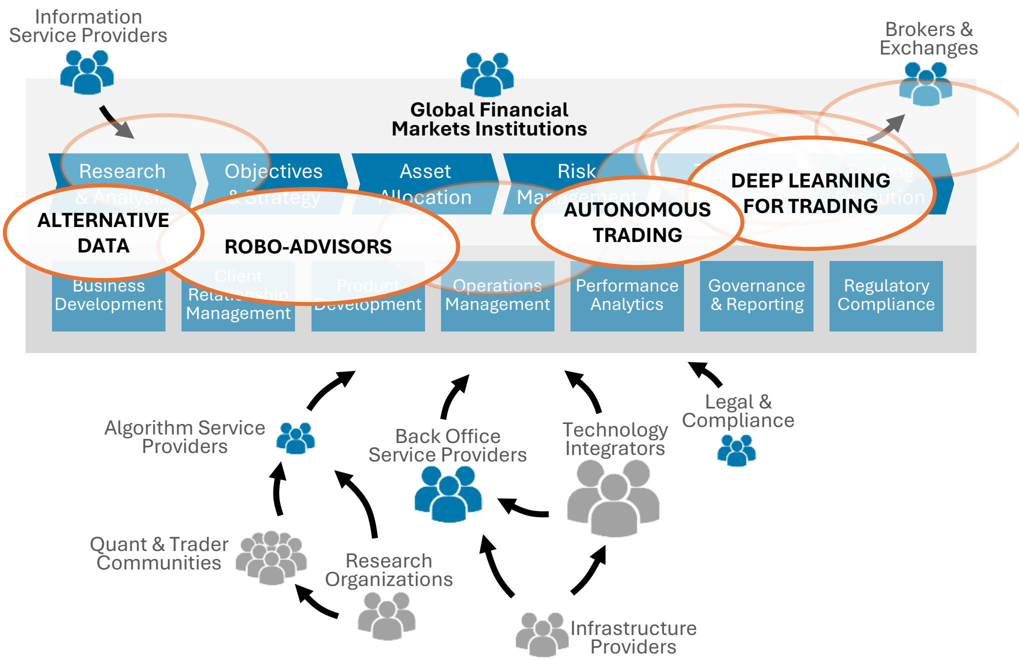

In the 2010s, supervised machine learning techniques become mainstream for a specific set of applications in hedge funds and for research, supported by specialist AI and algorithm providers and start-ups and the emerging cloud infrastructure providers, and we experience a short-lived and premature application to a more ambitious use in autonomous trading based on supervised learning techniques.

DEEP LEARNING FOR TRADING

Deep learning algorithms found applications in hedge funds such as Man Group (AHL Division), which started applying advanced neural networks for market prediction. Citadel, Two Sigma, and Point72 also adopted deep learning models for time-series analysis, risk management, and portfolio optimization with Numerai adding the twist of crowdsourcing the development of its models.

AUTONOMOUS TRADING

Several organizations attempted to use AI capabilities against asset price and/or contextual data to decide when to buy, hold and sell assets. An example is the AI-enabled hedge fund Aidyia, trading autonomously with great success for a short period of time, then incurring high losses before ultimately closing.

ROBO-ADVISORS

The rise of robo-advisors in the 2010s was led by companies like Betterment and Wealthfront, both of which used AI to provide automated, low-cost investment solutions to retail investors. The technology becomes more mainstream with for example Vanguard and providing what they promote as ‘AI-driven wealth management services’ to a broader audience.

ALTERNATIVE DATA

Supervised and unsupervised learning is also applied to contextual data such as news and financial reports, as well as the beginnings of social media data and other public data sources such as satellite imagery, by BlackRock with its Aladdin platform, and Thomson Reuters (now Refinitiv) and Bloomberg.

2020s AI REACHES PEAK HYPE, BUT MATURE SUCCESS STORIES EMERGE

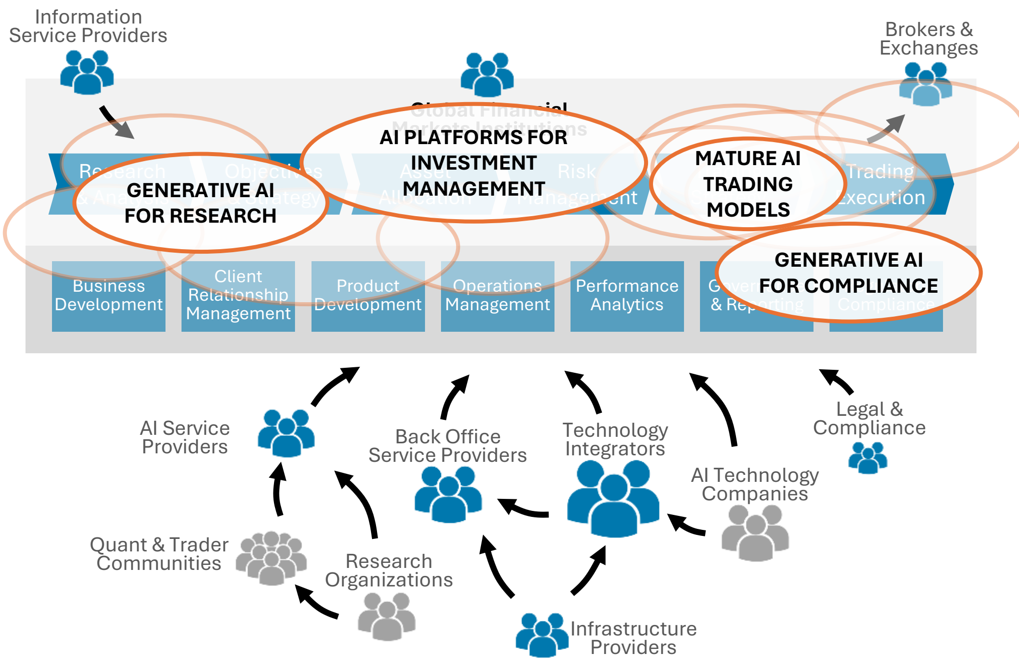

Despite the inflated hype around AI, and in particular generative AI, the 2020s have seen the emergence of a number of stable and mature applications of AI. Service Providers now provide not just models but AI platforms, while technology companies blur the lines between service provider and financial institutions. Although this is most notable in the area of crypto currencies and challenger banks, the same trend is happening in global financial markets as technology becomes intrinsic to how business processes run. Generative AI, as in many industries, is being used in many areas of business process such as handling customer interactions, but is also finding more specific uses in investment management, and the role of generative AI in forecasting and prediction is starting to be explored, even if only in generating synthetic time series data for traditional ML training.

GENERATIVE AI FOR RESEARCH AND FOR COMPLIANCE

As in many industries, generative AI is being applied to industry-specific problems, such as the need to summarize and report on the performance of portfolios and funds for customers and regulators. It is also a logical progression from previous NLP techniques in automating the review and analysis of large amounts of macroeconomic information to support research and strategy, and combining this with ‘chat’ interfaces for simpler adoption by non-technical analysts.

AI PLATFORMS FOR INVESTMENT MANAGEMENT

Tools to recognize types of market conditions, match client requirements against strategies and monitor on an on-going basis, with branded technologies such as MDOTM’s ‘Sphere’ and Solactiv’s ‘ARTIS’ combining supervised and unsupervised learning with other more traditional techniques.

MATURE AI TRADING MODELS

Supervised and unsupervised learning techniques and toolkits, enabled by advances made by AI Technology Companies, are now a standard part of methodology for algorithmic model construction, just as statistical models had been decades earlier. Experimentation shifts to reinforcement learning and generative AI to improve predictive power, albeit with the ever-present issue of how to make the leap from performance in the lab to performance in the real world.

5. CAN AI DELIVER ON ITS PROMISE?

There has been a huge amount of investment by global financial markets institutions, over decades, in the techniques and technologies of data and modelling, and the integration of these into the core activities from research through to execution. There have been two intentions:

1. to increase productivity through faster, more efficient business processes
2. and to enhance performance and results through better research, strategies, risk management and higher returns

As a steady evolution of the industry's long heritage of embracing data and analysis, the adoption of the latest techniques as they become mainstream is arguably 'table stakes' to remain competitive and up to date with the emerging skills and expectations both of the workforce and of customers.

But questions are being asked in every industry about whether AI will achieve the benefits being claimed. Many industry commentators have noted that businesses are yet to see the operational returns being promised by AI adoption, and the most visible indicator of the AI revolution is still simply the heightened share price of the AI companies providing the hardware and software infrastructures required to support these initiatives.

In this section we look at some of the evidence of what is being achieved in the real world rather than in the lab, consider why AI might not be delivering the benefits we expect, and share our opinion on what needs done to enable AI to deliver on its promise.

EVIDENCE FROM REAL WORLD RESULTS

Hedge funds, one particular type of global financial markets institution, have been the most enthusiastic adopters of AI technology over the last five decades. These should be the place to look to evidence the benefits that AI can achieve in real-world organizations, as opposed to the theoretical returns from proof of concepts or research. The largest hedge funds have built dedicated teams of mathematicians, data scientists and technologists using the techniques described in the last section to support their investment activities. Taking MAN Group as an example, these teams currently make up 675 people out of a total workforce of 1700, with access to the latest technologies from the AI technology companies, and access to dedicated AI research from their Oxford-Man Institute – a comprehensive commitment to technology[1].

HEDGE FUND RETURNS 1990s – 2020s

Although they follow diverse strategies and practices, the stated aims of most hedge funds are firstly to deliver 'alpha' returns that are above average, and to reduce risk through being uncorrelated with other asset classes, in particular the traditional long-only strategies of most funds.

The performance of hedge funds over the last twenty years, however, has been well documented by industry commentators and academics as consistently falling below other less technology-focused funds, and below benchmarks such as the S&P500.

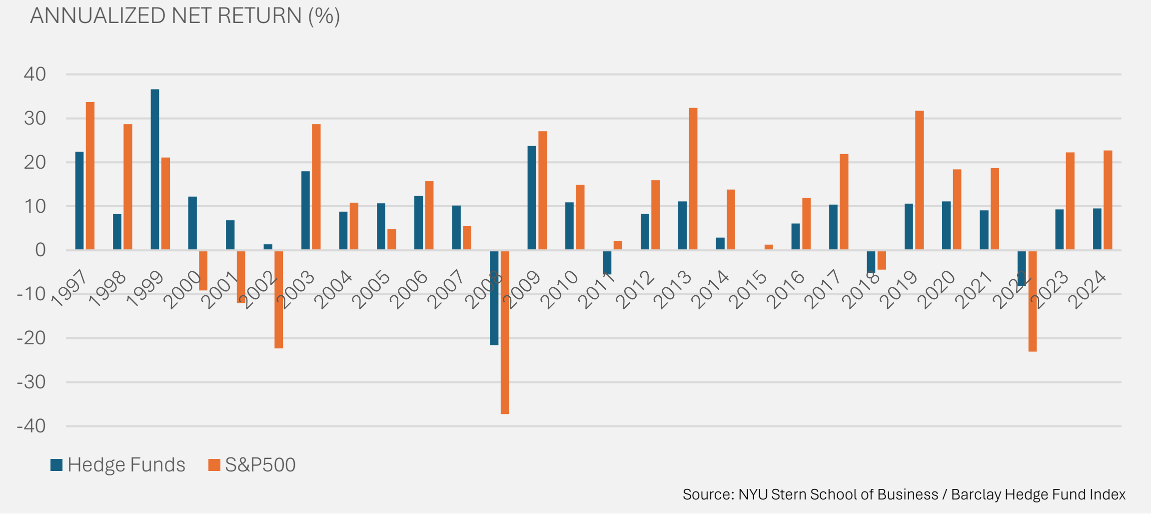

More detailed analyses show that hedge funds do achieve lower volatility overall, and perform better when markets are themselves volatile, but also that they typically achieve much of their returns from long-only strategies, or even cash[2], and by falling back on arbitrage as a way of generating results, rather than discovering alpha opportunities in the market.

So if the area of global financial markets that is best placed to leverage advances in AI is underperforming much simpler and cheaper approaches, what is it that AI is not achieving? What's the problem?

[1] Man Group Annual Report 2023
[2] eg Darden Business School, Twenty-Five Years of Hedge Fund Returns

WHY ISN'T AI ACHIEVING MORE?

AI can and has produced success stories across various industries, where a new machine learning approach to an old problem has achieved a step-change in results, but it seems that global financial markets is not yet one of them.

A PHARMACEUTICALS AI SUCCESS STORY

One emerging success story is in drug discovery. Proteins are small molecules made up of amino acids that play many important roles in living organisms. Combining amino acids in different ways results in proteins with different properties, but previous methods to predict this, by modelling how amino acids interact, could only work at small scale, and slowly. Researchers can now feed machine learning with known patterns of amino acids, and the proteins that these result in, and then much more quickly focus on promising proteins to progress to the next stage in the process. This approach unblocks a fifty-year-old industry problem[1], and translates into tangible, real world results in the form of the faster delivery of new and more effective drugs[2].

Although this is just one example, and still with its own challenges, it has two interesting high-level characteristics:

- stable, high-quality data is used to enable machine learning algorithms do the 'heavy lifting' for a specific problem, making a previously human-bounded scaling issue newly tractable
- a specific scientific problem is solved, but the approach can also be embedded in industry workflows to implement other processes, such as quickly finding ways to repurpose drugs to treat previously neglected diseases, including those affecting the poorest of the world's population[3]

Together these deliver results that are unmistakably better than what had been achieved before, and open up new opportunities for further innovation, for example by applying generative AI to improve the new process with additional data.

THE TROUBLE WITH GLOBAL FINANCIAL MARKETS

AI applied in Global Financial Markets, particularly in the core processes of asset management from asset allocation through to trading strategy, lacks many of the characteristics of the success story above. There are two main reasons:

1. The challenging data environment. Economists and researchers consistently describe financial markets as having lots of noise and very little signal, generated from a 'complex adaptive system'[4] in which 'patterns disappear as soon as they are discovered'[5]. Overfitting is always a concern, but in addition the actors in the system, such as exchanges, may not release data for training that accurately reflects operational experience, while orders might not necessarily be fulfilled as expected. This means that R&D results seldom translate into real world impact, and it can seem easier to focus on more durable, if less effective, strategies such as rule-based arbitrage.
2. The fragmented nature of how AI is being applied. Regulators, and quite possibly the institutions themselves, struggle to explain the real-world applications of techniques that are or claim to be based in AI[6]. While the asset management industry is notoriously secretive, it is possible that many AI successes are not really there to be shared at all – many institutions still require large staffs and 'star traders' to achieve even the results they do.

[1] Science 2021 Breakthrough of the Year
[2] Nature, Inside the nascent industry of AI-designed drugs
[3] Drugs for Neglected Diseases Initiative
[4] Eugene Fama, Efficient Capital Markets: A Review of Theory and Empirical Work (1970)
[5] Marcos López de Prado, Advances in Financial Machine Learning (2018)
[6] US HSGAC Report, AI in the Real World (2024)

CAN AI EVER DELIVER ON ITS PROMISE? MODELOMNI'S OPINION

AI, and in particular generative AI, is gradually lifting efficiency and productivity across many industries, and global financial markets is no exception. What is more debatable is whether applications of AI in our industry will ever be held up as one of those 'breakthrough of the year' examples. Many are pinning hopes on the use of generative AI to improve the speed and scale of macroeconomic information analysis and the decision making based on this[1]. Whilst we are not as scathing as some in our analysis of the value that generative AI can bring[2], we do believe that this kind of application of AI will simply be 'table stakes' for institutions in the future, and pinning hopes on this to deliver competitive breakthrough advantage is misplaced.

Instead, we think there are two things financial institutions should be doing to ensure that AI delivers the kinds of transformative benefit that are promised.

1. GET SPECIFIC

The challenges of applying AI in the core processes of global financial markets are real, and its 'complex adaptive systems', like amino acids, are almost certainly too hard to model convincingly by conventional methods. Instead, we need to identify the particular problems that are human-bound, fully understand the data that is appropriate for the problem, and match with the AI techniques that make most sense for that problem rather than simply what is in vogue[3].

For example, there are many simple models that attempt to capture the dynamic aspects of financial markets through concepts such as momentum and trend. The configuration and application of these models can be adjusted by 'instinct' by the most skilled traders to be successful, but generating these models from first principles, even with machine learning, is simply too hard. So how do we scope down on the different parts of these models, and data related to how these operate, and make them operate significantly better?

2. USE THE ADVANCES TO TRANSFORM BUSINESS MODELS

Better models can only deliver small-scale, localized improvements unless they are applied to re-shape the way a business operates, and so transform aspects of the industry itself.

Despite the huge amount of interest in machine learning for forecasting and prediction, the shape of the activities, who carries them out, and the business model stays much the same for asset managers, traders, and those who manage different types of risk.

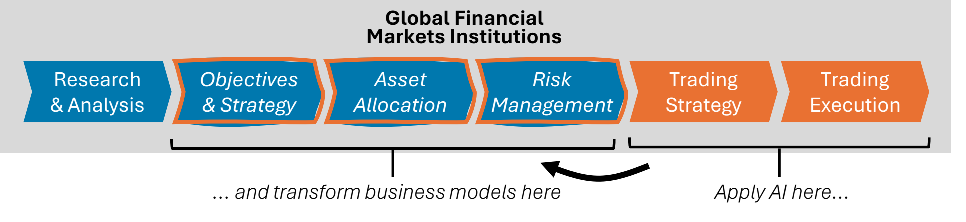

For example, if we were able to make any trading strategy generate significant results, and at the same time reduce volatility to a negligible level, what would that let us do? Manage currency risk without recourse to traditional and costly strategies such as forward contracts and options? Enable greater foreign direct investment into developing countries?

MODELOMNI OPINION

Only if we use our understanding of the real problems of global financial markets, combined with a deep understanding of where and how specific AI improvements can be made, will we achieve the transformative breakthroughs that really make a difference.

[1] see for example [Bridgewater Associates starts $2 billion fund that uses machine learning ...]
[2] Goldman Sachs, [Gen AI: too much spend, too little benefit?]
[3] see for example this research paper: [Do We Really Need Deep Learning Models for Time Series Forecasting?]

ABOUT THE AUTHORS

Kamran Usmani
CEO and founder
kamran@modelomni.com

Kamran has worked in global financial markets for over 30 years as a hands-on trader, asset manager and brokerage CEO, experiencing the issues at first hand that he is now working to resolve with Modelomni.

Dr Simon Smith
Head of Innovation & Strategy
simon.smith@modelomni.com

Simon has over 30 years experience in academic research and software start-ups, working on tooling and automation for modelling, decision management, and the development and assurance of AI systems.