FMSB
FICC MARKETS STANDARDS BOARD

Spotlight Review

Emerging themes and challenges in algorithmic trading and machine learning

April 2020

About FMSB

FICC Markets Standards Board Limited ('FMSB') is a private sector, market-led organisation created as a result of the recommendations in the Fair and Effective Markets Review ('FEMR') Final Report in 2015. One of the central recommendations of FEMR was that participants in the wholesale fixed income, currencies and commodities ('FICC') markets should take more responsibility for identifying and fixing poor market practice so that they operate in the best interest of their clients. Clear, practical guidance that delivers transparent, fair and effective practices will rebuild sustained trust in wholesale FICC markets.

FMSB brings together people at the most senior levels from a broad cross-section of global and domestic market participants and end-users.

In specialist, focused committees, subcommittees and working groups, industry experts debate issues and develop FMSB Standards and Statements of Good Practice, and undertake Spotlight Reviews that are made available to the global community of FICC market participants and regulatory authorities. FMSB has issued 18 publications since 2016. As part of its analysis on the root causes of market misconduct, FMSB is focusing on the challenges of new market structures.

Spotlight Reviews

Spotlight Reviews encompass a broad range of publications used by FMSB to illuminate important emerging issues in FICC markets. Drawing on the insight of members and industry experts, they provide a way for FMSB to surface nascent challenges market participants face and may inform topics for future work. Spotlight Reviews will often include references to existing law, regulation and business practices. However, they do not set or define any new precedents or standards of business practice applicable to market participants.

Find out more about the FICC Markets Standards Board on our website fmsb.com

The author

Rupak Ghose is leading the FMSB Spotlight Reviews on FICC market structure and the impact of regulatory and technological change on the fairness and effectiveness of wholesale FICC markets.

Rupak has more than 15 years' experience in FICC market structure. He was Head of Corporate Strategy for ICAP/NEX for six years providing advice on strategic direction, market structure change, customer/competitor landscape, new business ventures and M&A. He was a trusted counsellor to the PLC Board on the dramatic transformation of the business from interdealer broker to focused financial technology firm and the ultimate sale of the businesses.

Prior to this Rupak spent more than a decade at Credit Suisse as an equity research analyst where he was Top Three ranked on numerous occasions by buy-side clients in Institutional Investor and other surveys for his coverage of asset managers, exchanges and other capital markets focused companies.

Contents

1 **Executive summary** p.2

2 **Model risk**
Eight factors to consider in the importance of model risk management p.4

3 **New asset classes**
Expansion of algorithmic market making p.14

4 **Machine learning**
Growth and challenges p.18

5 **Execution algorithms**
Growing use in FICC markets p.24

6 **Best practices**
Benefits of practitioner-led solutions p.28

1

Executive summary

This Spotlight Review examines emerging themes in algorithmic trading in FICC markets including model risk management in market making, the adoption of new machine learning techniques and the increased use of execution algorithms. The latter refers to algorithms that are offered to clients on an agency basis and used for order execution. This Spotlight Review aims to generate further discussion on these topics and their relevance to future standards work by FMSB. Given the topical themes discussed, it will be of interest to a wide audience of participants in global wholesale FICC markets but it is specifically targeted at those senior managers with supervisory responsibility for algorithmic trading and those working on the application of machine learning in algorithmic trading.

The use of algorithmic trading is not new, and over the past two decades it has profoundly changed the nature of trading and market structure in many FICC markets in terms of the increased velocity of trading, levels of internalisation and cross asset/venue trading patterns. Algorithmic trading methods and electronic trading platforms have grown in a synergistic fashion.

As the adoption of algorithmic trading continues to grow the focus on governance of algorithmic trading has increased significantly. Central banks and other regulators have issued guidelines on the controls for algorithmic trading, focusing primarily on the documentation and controls expected for the development, testing and deployment of algorithms; and FMSB members are developing a Statement of Good Practice expanding on this area. However, the application of model risk management to algorithmic trading is an area that has received less attention. Nevertheless, the materiality of algorithmic model risks warrants a specialised practitioner-led approach.

Historically, algorithmic trading has been most prominent in highly liquid markets, which have significant amounts of high-quality data. As the application of algorithms has expanded into less liquid products and with increased utilisation of new machine learning techniques, the challenges of securing the quality and consistency of data needed are self-evident. Perhaps less obvious is the need to manage for increased model risk.

Progress towards increasing use of self-learning machines will be incremental and over an extended period. In the near term, machine learning in wholesale FICC markets looks likely to remain restricted to specific minor functions only and as a relatively small part of the overall trading and reporting process with tight controls in place. As in other businesses where machine learning is being adopted, there are nascent concerns about the conduct risks that might crystallise as a result of unintended design flaws, implementation and use. There is also increasing discussion within the industry about practices that can mitigate any market abuse or stability risks that may emerge.

The use of execution algorithms is well established in cash equities markets and are increasingly being adopted in foreign exchange. As they move into rates, credit and emerging markets, a key challenge will be sourcing market data, given the less continuous nature of these markets. Moreover, banks providing execution algorithms to their clients need to be alert to any potential conflicts of interest that may arise in how they provide such agency products and how it relates to their core FICC market making businesses, which act on a principal basis.

The increasing usage of algorithmic trading and the growing complexity of models makes the topics and emerging themes discussed in this Spotlight Review extremely important. There are likely to be benefits from creating global best practices for model risks which are not fully covered by existing regulations. FMSB has a role to play in areas like this, where there may be knowledge gaps between the private sector and regulators and where there is scope for market participants to work together to address the issues rather than in isolation. We propose that market practitioners, given their deep domain expertise, are in a better position to provide solutions that are proactive on managing risks.

For global wholesale FICC markets this Spotlight Review outlines the:

> increased importance of model risk management in areas where algorithmic trading is being deployed;

> challenges faced as algorithmic market making expands into less liquid asset classes;

> adoption of new machine learning techniques in algorithmic market making;

> increased use of execution algorithms; and

> best practices, including the role of practitioner-led solutions.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

Eight factors to consider in the importance of model risk management

There are a number of existing regulatory requirements and associated guidance focusing on algorithmic trading both in Europe and beyond. Furthermore, guidance has been published in some jurisdictions on the topic of model risk. However, the application of model risk management to algorithmic trading is an area that has received less attention. In this section we outline eight factors to consider when looking at the importance of model risk management to areas where algorithmic trading is deployed.

1. Significant progress by regulators
2. The importance of model risk management
3. Unique nature of model risk in algorithmic trading
4. Crucial role of data inputs
5. The difficulty of benchmarking algorithmic models
6. The important role of testing and validation in model governance
7. Scenario analysis and capturing unintended consequences
8. The need for a robust second line of defence

1. Significant progress by regulators

Algorithmic trading is increasingly regulated in major global financial centres. In the UK, both the Financial Conduct Authority (FCA)[1] and the Prudential Regulation Authority (PRA)[2] have issued supervisory guidelines relating to governance, algorithm approval processes, testing and deployment, documentation of algorithms, and risk controls. Significant risks arise from the failure of systematic or operational controls that are intended to prevent or limit loss exposure for highly automated transactions. System runaway issues have the potential to cause material losses in a short period of time. The lack of a robust software development lifecycle process was cited as the main cause of high-profile incidents in recent years such as seen at Knight Capital.[3] The other regulatory focus has been on conduct, i.e. the risk of algorithmic strategies being coded, or learning, to disadvantage clients, abuse markets or cause disorderly markets.

The current regulatory guidelines, which are principally focused on operational and conduct risks, may mitigate some risks from models through the consolidated approach to documentation, testing, controls and performance analysis at a trading algorithm level. For instance, a lack of model robustness may lead to unexpected P&L losses but these would be bounded by a number of risk controls at an algorithm level. These include continuous validation in the form of P&L checks covering volatility/skew of returns and significant financial losses, position limits, price/spread limits. As a result, even though some models in algorithmic trading strategies may be highly complex, residual algorithmic model risk does not necessarily have to be high.

2. The importance of model risk management

At the same time, as algorithmic trading expands into new and more complex areas, there may be a benefit to best practices relating to how models are deployed here. Model validation in algorithms should consider factors such as model complexity, appropriateness of model methodologies, input data quality, controls around model assumptions and implementation. Execution controls, back testing, sensitivity analysis, erroneous data handling measures, and clear documentation are some of the key mitigants.

Risks can be greater in less liquid asset classes where pricing is less transparent, and the liquidity of the product should be considered when judgements about model risk are being made. At the same time, expectations around pricing precision should also be considered. For instance, in data-rich, heavily-traded instruments these expectations can be extremely high, while in data-light, infrequently-traded instruments pricing precision may have a larger allowable error term.

Early supervisory guidance on model risk from the Board of Governors of the Federal Reserve System in the paper SR 11-7[4] was focused across all types of models, with reference to risk management and balance sheet/capital calculations, given the inadequacies exposed by the 2007-2008 global financial crisis. The paper defined a model as follows:

> *"...the term model refers to a quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates. A model consists of three components: an information input component, which delivers assumptions and data to the model; a processing component, which transforms inputs into estimates; and a reporting component, which translates the estimates into useful business information."*

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

2. The importance of model risk management continued

It goes on to state that:

"Model risk occurs primarily for two reasons:

> *The model may have fundamental errors and may produce inaccurate outputs when viewed against the design objective and intended business uses… the quality of model outputs depends on the quality of input data and assumptions…*

> *The model may be used incorrectly or inappropriately. Even a fundamentally sound model producing accurate outputs consistent with the design objective of the model may exhibit high model risk if it is misapplied or misused."*

3. Unique nature of model risk in algorithmic trading

There are fundamental differences in algorithmic model risk when compared to more traditional risk or capital calculation models. Consequently, any approach leveraging existing model risk validation processes may need adjusting. The risk associated with misspecification in any single model may be mitigated by bounds placed on how any model output data is used by the overall trading strategy. This combined with the dynamic feedback in a live electronic trading ecosystem means that residual model risk can be low in algorithmic trading. Consequently, less weighting can be placed on the accuracy of a model's estimates or predictions and more on the implementation testing, back testing and controls that minimise the conduct and operational risks.

The number of individual models deployed in an algorithmic trading system is much larger than traditional areas so documentation and model risk ratings, while still key, will need to be scalable to be effective. Moreover, the depth and frequency of model validation deployed should reflect the complexity and potential impact of individual models. There are often very simple model assumptions made within an algorithm, for instance the use of moving averages in price computation. For these 'de minimis' models, it is difficult or

3. Unique nature of model risk in algorithmic trading continued

impossible to perform an assessment of ongoing performance, especially determining the specific impact on the overall P&L generated. This should be considered when the governance framework is being applied.

Other important components might not meet the definition of a model and so could typically be out of scope for model risk review. Interpretations differ on an appropriate definition for a model within algorithmic trading. One approach is that an algorithmic trading model estimates or predicts an observable quantity, or that it involves some mathematical derivation of a non-observable quantity. Either of these approaches renders much algorithmic code as business logic and is therefore out-of-scope of the model definition.

4. Crucial role of data inputs

The amount, quality and consistency of data inputs represent crucial components of model risk management. The risks here include erroneous or stale input data and broader constraints such as sparse central limit order book transaction data, lack of depth and accuracy in other data sources, or single points of failure. Poor data quality and governance can create operational risks and conflicts of interest from inappropriate use of private client data and incorrect or inadequate interpretation of data sources.

In many liquid markets, there is a dependence on Central Limit Order Books (CLOBs) as reference prices and when a lack of depth or market structure issues drive price changes on these platforms that are not in line with fundamentals, there is a risk with following them 'blindly' as a key data input. A high-profile example in rates markets was the 15 October 2014 US Treasury flash crash[5] when despite an absence of material news flow there was a 37 basis points (bps) intraday trading range in 10-year US Treasury yields. Two examples in recent years in foreign exchange markets are the 7 October 2016 British sterling flash event[6] and the 3 January 2019 Japanese yen flash crash.[7] In the former despite limited news flow sterling depreciated 9% versus

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

4. Crucial role of data inputs continued

the US dollar in early Asian trading hours before retracing most of the move. The latter saw a 4% appreciation of the Japanese yen against the US dollar and much larger (circa 10%) moves against other currencies such as the Australian dollar. It had similarities to the sterling flash crash in terms of limited news flow and occurring during light trading in early Asian trading hours, but unlike other flash crashes in foreign exchange, it impacted a wider range of currency pairs than just the US dollar pairing.

In less liquid markets such as credit, post trade regulatory data may not give an accurate picture of liquidity given it tends to be focused on smaller size trades. Recent trade data may become irrelevant if market conditions change materially and a credit rating used as an input in pricing may become out of date relative to market conditions. A recent example has been the wide discount at which bond exchange-traded funds (ETFs) have traded relative to their net asset value reflecting the superior liquidity of the former and the lag of the latter, where third party pricing services may not have updated their valuation models to reflect changing conditions in credit markets.

The Federal Reserve stated in its model risk guidance that:

"The data and other information used to develop a model are of critical importance; there should be rigorous assessment of data quality and relevance, and appropriate documentation. Developers should be able to demonstrate that such data and information are suitable for the model and that they are consistent with the theory behind the approach and with the chosen methodology. If data proxies are used, they should be carefully identified, justified, and documented."

5. The difficulty of benchmarking algorithmic models

Another consistent focus of the 2011 Federal Reserve guidance on model validation is benchmarking.

Comparison with alternative theories and approaches is a fundamental component of a sound modelling process…Benchmarking is the comparison of a given model's inputs and outputs to estimates from alternative internal or external data or models."

In other segments of financial services, such as credit ratings data or securities valuation, there are third-party industry data providers that allow for independent benchmarking relative to peers. Algorithmic trading uses many publicly available data inputs and some comparisons here of inputs may be possible. However, peer group comparisons of the inner workings of algorithms and modelling assumptions are more difficult because of the proprietary nature of most algorithmic trading models and how they process and use these data inputs. Where peer group benchmarking is not appropriate, performance monitoring is critical.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

6. The important role of testing and validation in model governance

Given the limited ability to conduct detailed benchmarking against competing algorithmic trading models, it is important to have a rigorous model validation and performance monitoring process. With the drive for improved efficiency across the whole financial services sector it is natural for there to be a drive to re-use as many components of existing models as possible in new products and geographies. The question of whether a particular model is appropriate for use in a specific market, asset class or venue is not a new one, but likely to be more common than ever in future.

Core to model assessment is the testing of model robustness and reliability to ensure safe and sound implementation. However, SR 11-7 allows firms to take materiality of model risk into consideration when devising an approach to model risk management in order to meet supervisory expectations. Given the differences between pricing or risk and algorithmic trading models, different model validation approaches may need to be developed, where the control framework should be considered in deciding the model risk rating and any subsequent validation and testing requirements.

7. Scenario analysis and capturing unintended consequences

As model risk increases in complexity, scenario analyses that stress test data inputs and their impact on algorithmic models become increasingly important. This may include negative stress testing, which seeks to determine the conditions under which the model assumptions break down. Where model risks are found, controls should be put in place. Limitations to data inputs can add to the uncertainty of results, and the real world is generally more unpredictable and complex than models. Another unintended risk that is extremely hard to capture is that of similarities, and resulting interdependencies between, the algorithmic models of different firms.

Capturing the unintended consequences of algorithms and modelling components not performing in line with their intended aims is especially important. The behaviour of individual algorithms and modelling components may be as expected, but the combination of models up to the trading algorithm level may not be as expected. Unfortunately, it is very difficult to develop testing to demonstrate this, even with extremely clear guidelines on the aims of specific algorithmic components.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

8. The need for a robust second line of defence

Given the high degree of technical expertise needed, there is a fine balance to be drawn between having validation by deep subject matter experts in the first line of defence and an independent, unconflicted second line of defence with perhaps lower technical expertise. It may be difficult to have a second line of defence with the quantitative trading expertise able fully to challenge the first line, but it is crucial that the second line has enough product and technical knowledge to validate and test models properly. This will involve understanding the mitigating controls and being able to drive relevant scenario analysis covering how the model performs in different conditions to minimise any market abuse and market stability risks.

Many large banks have highly experienced and dedicated second line functions, but there remains a question about whether this is as embedded across all firms as, for instance, independent product valuation and balance sheet validation functions are. There is a different challenge for how smaller firms without such resources can perform these tasks.

Given the expected growth in machine learning and in automated trading in markets with less transparently priced products (which we discuss further in this Spotlight Review), there are likely to be benefits from creating global best practices for model risks which are not fully covered by existing regulations. FMSB has a role to play in areas like this, where there may be knowledge gaps between the private sector and regulators and where there is scope for market participants to work together to address the issues rather than in isolation. FMSB's work focuses on areas that impact transparent, fair and effective markets and supports open and competitive markets that deliver the right outcomes for end-users.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

Expansion of algorithmic market making

Traditionally, algorithmic trading has focused on near-continuous markets such as cash equities, spot foreign exchange, futures and on-the-run US Treasuries which are extremely liquid and can provide huge amounts of historical market data. These include both centralised marketplaces and more fragmented ones where bilateral trading dominates, but all these markets have publicly available reference prices generated by transactions on CLOBs. More recently, algorithmic market making has started to expand into other product categories such as over-the-counter (OTC) derivatives and bonds (beyond on-the-run US Treasuries).

These developments have been driven by a combination of factors: opportunities created by new technology, regulatory imperatives for more electronic trading and the need to reduce the costs of trading in a world where returns are under heavy pressure. The arrival of electronic and algorithmic trading in these new asset classes has brought significant benefits for market participants.

But algorithmic trading in these new product categories has also created new challenges and risks, given the more limited transaction data available, more limited transparency, the greater market concentration of counterparties, the lack of centralised marketplaces and the potentially longer holding period of positions. The use of algorithms in such markets can create different market fairness and effectiveness risks to those in faster markets and potentially result in higher tail risks. Three challenges are presented below.

1. Increased difficulty of sourcing market data

Historical market data is the fuel that powers algorithms, and most algorithmic trading models need detailed market data stretching back over a period that includes multiple types of market environment, including periods of stress. By definition, the amount and quality of historical market data is more limited in less liquid markets. For instance, there may not be detailed tick level data in less liquid markets or transaction data may be delayed in terms of reporting (e.g. most bond transaction data in Europe is reported with a one-month time lag, with only a limited number of individual issues currently reported in real-time or with a 15 minute time lag).

For many OTC derivatives, corporate and emerging market bonds, the level of transaction data is too limited to drive algorithmic models. In these cases, pricing models can sometimes be built on data from related, more liquid markets, as a proxy for the less liquid instrument. There are opportunities to engineer 'artificial' data sets that have similar statistical properties to real market order and transaction data, in order to train algorithms. There are also opportunities, with machine learning algorithms, to use unstructured data from other sources in order to enrich historic price information (see below). Producing and maintaining such parallel, engineered or unstructured data itself carries serious and practical data governance challenges for firms attempting to use such strategies.

2. The role of public reference prices for hedging

The challenges discussed above in terms of sourcing market data are linked to a second phenomenon – liquidity sourcing during extreme market moves such as 'flash crashes' and the role of CLOBs.

There has been a significant growth of both single-bank and multi-dealer disclosed platforms, particularly in foreign exchange markets. Linked to this and the growth of algorithmic trading, has been a rapid growth in internalisation by large dealer banks where they avoid hedging into traditional interbank CLOBs or trading directly with other wholesale market participants. In normal markets liquidity providers try to avoid interbank platforms with public market data as much as possible as part of their efforts to minimise market impact and information leakage, which has often benefitted clients over the same period through spread compression. That said, the existence of CLOBs provide important places for hedging in more volatile markets. A recent and stark example of this was the unpegging of the Swiss franc by the Swiss National Bank in January 2015. A sudden, unprecedented move saw one-sided flow, with some banks unable to internalise to reduce risk, and reaching risk limits and liquidity rapidly disappearing on single-bank and multi-dealer disclosed platforms. This in turn led to a material increase in such liquidity providers' activity on the interbank CLOBs.

Most of the newer products where algorithmic market making is expanding, are less liquid and do not have the public liquidity on one or more CLOBs that is available in foreign exchange. This inevitably increases the tail risk associated with liquidity shocks or sudden gapping in prices in these markets.

The importance of public reference prices goes beyond the question of liquidity in times of stress. It also directly affects the question of fairness. Established manipulative techniques, e.g. inappropriate use of pre-trade information, spoofing and collusion (as discussed in 'FMSB Behavioural Cluster Analysis – Misconduct Patterns in Financial Markets'[8]) are all easier to perpetrate in conditions where public reference prices are harder to establish, as may be the case in these less liquid products. A key goal of algorithmic governance needs to be ensuring that algorithms that go to market are fair in terms of not creating market abuse and market stability risks.

3. More market and concentration risks in less liquid products

As algorithmic trading expands into markets that are less liquid, the associated risks will be greater, including the likelihood of 'gap' pricing driven by idiosyncratic events.

Hold times in liquid markets like foreign exchange are typically sub-seconds to minutes, but for other FICC markets these times may be days or even weeks. At the same time, it should be noted that as these other markets see more electronification, it is reasonable to assume that hold times will decrease. In recent years this has been seen in corporate bonds, albeit this has been mostly focused on smaller ticket sizes and bond issues that were already relatively liquid from larger issuers.

Whether it be longer hold times in less liquid markets or scope for greater losses from leverage in derivatives products, the market risk associated with algorithmic trading is likely to increase in coming years, as product coverage grows. In some instances, sporadic liquidity in one product may be compensated by hedging strategies in adjacent products, with associated basis risk.

Markets in less liquid products are also likely to be much more concentrated as there are unlikely to be more than one or two non-bank market makers who are willing to extend liquidity in all market conditions. Model validation in such cases is even more important and needs to take account of the tail risks of potentially disappearing liquidity.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

Growth and challenges

Significant increases in computing power and data storage in recent years have stimulated interest in using machine learning techniques to trade markets. Machine learning algorithms are fundamentally different from the traditional, deterministic rules-based algorithms that have been in use for several decades already, because they use neural networks and other 'deep learning' techniques.

Bank of England survey on machine learning

The October 2019 joint Bank of England and FCA report 'Machine learning in UK financial services'[9] started a public-private consultation process in this area. This report defined characteristics of machine learning (ML) models as follows:

"...while it will always depend on a multitude of factors whether a ML application poses a meaningful prudential or conduct risk, ML use can alter the nature, scale and complexity of IT applications and thus, a firm's IT risks. There are three dimensions to this...:

ML applications are more complex. *ML models are often very large, non-linear and non-parametric. This makes it harder to comprehensively understand their properties and to validate them. This means certain forms of risk-taking could go undetected. This type of complexity can constitute a significant change to existing systems.*

ML uses a broader range of data. *ML applications may often use entirely new types of complex, including unstructured, data. For instance, this could be data from news sources, satellite images or social media.*

ML systems are larger in scale. *ML systems increasingly consist of a multitude of interacting components. This can make it harder to validate if they always interact as intended. In many cases, this change is incremental."*

The report also incorporated a comprehensive survey, which highlighted the growing usage of machine learning from a low base. The highest penetration was in areas such as fraud detection and anti-money laundering, but the survey also showed a level of familiarity and adoption in sell-side trading.

Across all financial services, machine learning is typically being used to support current operations, rather than replace them. In sell-side trading this may involve it being used to assist manual trading desks in evaluating venue, timing and order size choices and determining the probability of an order being filled. The Bank of England/FCA report also found that machine learning has been used by some firms to determine order routing logic that is contained in 'algo wheels'.[10] However, in general, machine learning is used as part of a multi-layered execution process, which also involve algorithms based on simpler, rules-based models.

Despite the developments noted above, it is worth noting that virtually all the algorithmic trading that banks and large non-bank market makers conduct today is still built around relatively transparent rules-based deterministic models. There is very limited risk capital being deployed using machine learning algorithms alone as the basis for the whole market making process. Unsurprisingly, where it is deployed, machine learning is largely used in cash equities and foreign exchange, rather than rates and credit products, due to the longer history of rules-based algorithmic trading in these markets and the plentiful data sources and public venues.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

Challenges associated with machine learning[11]

1. Model drift

When trading engines are powered by machine learning, the relationship between data inputs and price outputs is much more obscure. The very low signal-to-noise ratio in the data, combined with the very large amounts of data that are mined and the complex often multi-layered decision trees that underlie the machine learning models themselves, may contribute to price formation being opaque. Unlike deterministic rules-based algorithms, where price formation is always performed in the same way with set inputs and steps, machine learning trading engines 'learn for themselves' how to create prices by repeated, constantly evolving, experimentation and it becomes very hard, or impossible, to trace how 'decisions' are made in the optimisation process.

The difficulty of tracing how decisions have been made by the machine make it very difficult to prevent in advance, or to correct afterwards, undesirable model outcomes. For example, the machine may discover complex, non-linear 'hidden' correlations that it is difficult or impossible for the programmer to anticipate or discover. Further, it is impossible to predict how a machine, trained on known historical data but 'making its own decisions' will react when it is live in the market with a much larger dataset and it encounters events that have not been seen before in the data that was used to train it.

Concerns about these transparency problems lie behind the increasing regulatory focus on 'explainability', model risk management and software validation, as well as how management and boards can satisfy themselves that they understand, at some level, what is going on inside the 'black box' of the model. In August 2019 the Bank of England published Staff Working Paper No. 816 'Machine learning explainability in finance: an application to default risk analysis'[12] which explored these topics in detail.

2. Bias

Ultimately, machine learning is all about discrimination, and unpredictable discrimination during the optimisation process, when an enormously wide range of factors are analysed. These biases could include unexpected or unfair changes in pricing or liquidity to certain types of market users, or even to individual customers, as a result of factors that are impossible to uncover because they lie effectively undiscoverable in the heart of the optimisation engine.

Another type of bias may also occur: the risk that a machine optimising on its own will 'discover' that unethical, manipulative trading practices are more profitable than ethical trading. Indeed, this is virtually a certain outcome, if the machine does not have an 'ethical governor' that tests the optimisation process against ethical benchmarks and rejects trading tactics that fall short of these standards. These ethical benchmarks are much more complex to describe than formal laws and regulations.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

3. Market concentration and correlation

The rules-based algorithmic trading developed in the past couple of decades has fostered competition, allowing non-bank market makers and traders to develop successful businesses and grow market share at the expense of banks. But network effects can create winner-takes-all market structures.

The way in which machine learning models improve by accessing more data is likely to create data network effects, which may in turn create barriers to entry for new firms. It remains to be seen whether these barriers will entrench the power of today's large financial services firms or, alternatively, allow technology-based competitors to create new oligopolies at the expense of today's financial sector.

Either way, unless they are carefully managed, concentrated market structures may disadvantage market users by unfair rationing of liquidity, skewed pricing, and other non-price based discriminatory barriers. As algorithms optimise big data from new sources, they may inadvertently increase, or create new correlations between macroeconomic or other input variables. Hungry algorithms will over time arbitrage the profit potential in these correlations – a machine learning version of the 'crowded trade' phenomenon – but in doing so they may make markets more fragile to unforeseen shocks and more interconnected, as multiple users depend on a limited number of underlying data relationships.

4. Resources

There is a big skills gap for the expert programmers, data scientists and risk managers who can safely develop, test and implement machine learning in financial markets. At an individual level, these skills are in short supply in the private financial services sector and among central banks and market regulators; and they contribute to a knowledge gap among senior management, in the boardrooms of financial services firms and at policy makers about the hazards of artificial intelligence and machine learning.

For all these reasons it seems likely that market participants who do venture into machine learning for trading purposes will focus on building an extremely tight control 'sandbox' with significant P&L buffers.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

Growing use in FICC markets

Execution algorithms are growing in importance. These refer to algorithms that are offered to clients on an agency basis and used for order execution. Their order logic works in a systematic fashion, typically splitting a larger order into many smaller orders based on the available liquidity. As with the other topics addressed in this Spotlight Review, new technology and data are key to this development.

Rationale for using execution algorithms

The key reasons to use execution algorithms are to reduce execution costs and market impact. Additionally, execution algorithms are increasingly being used by the buy-side to help meet best execution obligations. Execution algorithms have also benefited from an increasing number and type of execution venues. We expect these factors to underpin rapid growth in the use of execution algorithms over the coming years, driven by the electronification of financial markets.

Expanding into FICC markets

Execution algorithms have been common in cash equities for a long time and penetration levels in those markets are extremely high.

As with electronic trading and algorithmic market making, penetration levels of execution algorithms are lower in FICC markets. Usage is greatest in foreign exchange markets, with recent growth underpinned by the increasingly fragmented FX market structure. Bond markets are earlier in the adoption cycle but there has been growing demand in recent years. Currently the most well adopted execution algorithms in bond markets are venue specific, in part due to the lower fragmentation in these markets.

Demonstrating best execution

The Markets in Financial Instruments Directive (MiFID) II has imposed more onerous best execution requirements on buy-side firms. Execution algorithms can play an important role in demonstrating compliance with these requirements, but it is critical that buy-side firms understand the algorithms they are using and that these allow them to deliver best execution. One way that many buy-side participants are trying to navigate this is the use of algo wheels that automatically select which algorithms to use, a little like smart order routing functionality with an execution management system. However, algo wheels are reliant on their inputs, which require codifying a limited number of normalised trading strategies and defining for each of them goals, constraints and flow characteristics.

Aligning technology with an asset manager's execution strategy

Asset managers outline their use of execution type (i.e. low touch algorithmic versus high touch) and rationale in their best execution summaries pursuant to MiFID II Regulatory Technical Standard (RTS) 28.[13] For smaller ticket sizes, algorithmic execution is common. The primary reason for using high touch in FICC markets was certainty of execution and limiting market impact.

With the proliferation in the number of algorithms being offered, having a clear view on execution strategy and which algorithm is most suited to delivering on these goals is important, as these vary depending on asset class and product liquidity. In liquid instruments with continuous markets, price and speed are generally the determining factors in execution decisions. In more episodic markets, price is important but can be secondary to certainty of execution, minimising risk and limiting market impact.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

Governance and model validation

Benchmarking relative to peers is easier when looking at execution algorithms in comparison with market-making algorithms. This is because information related to the type of model being used by execution algorithms (e.g. VWAP – volume weighted average price) is given to the users of such products, and easily available for peers to see.

Understanding model resilience, real-time controls, manual override functionality and other operational risk controls is particularly relevant in gauging how such execution algorithms may hold up in stressed markets. Article 6 of MiFID II RTS 6[14] imposes testing obligations for all investment firms running algorithms (and applies to both market making and execution algorithms), and has placed more of a burden on investment managers to test each new broker algorithm suite and parameter set as well as to demonstrate that they have a global 'kill switch' for all open/working orders.

Improving clarity of disclosure and managing conflict of interests

Many large buy-side firms already have sophisticated trading operations with deep electronic trading and data expertise, view execution as a core competency and have a deep understanding of the execution algorithms. In the coming years smaller asset managers and corporate clients will likely start to adopt such technology tools, making proper governance and disclosures very important.

Buy-side clients should be able to choose exactly what execution algorithms they want and for this they need to be provided with clear and easily understandable disclosures. If, for example, certain execution algorithms are naturally more likely to take liquidity from particular venues, or from the in-house principal trading desk, then that needs to be made clear. Disclosures need to be easy to understand for end-user clients of varying degrees of sophistication, so that they can match their individual execution requirements with the most appropriate execution algorithm.

There are two fundamental issues to consider as execution algorithms are rolled out in FICC markets:

i. The capacity in which the dealer or bank that provides the execution algorithm is acting. This differs between cash equities, where banks are largely trading in an agency role, and FICC markets, where banks are largely acting as principals and may be carrying significant inventory. It needs to be clear how the different conflicts of interest that may arise in relation to the potentially competing interests of the bank/dealer are to be managed, including what governance will be in place to protect information confidentiality.
ii. The availability of data inputs. Cash equities generally have dominant primary venues and plenty of continuous market data from different sources. In contrast, in many fixed income products there may not be a primary venue at all, and in liquid markets such as foreign exchange they may only represent a small portion of the overall market. These may pose considerable challenges in terms of sourcing accurate market data.

Execution algorithms should become very useful tools for driving efficiency in FICC market structure and to alleviate some of the cost pressures that the FICC industry faces, but they are unlikely to be a 'magic bullet' in the delivery and measurement of best execution.

Contents

Executive summary

Model risk

New asset classes

Machine learning

Execution algorithms

Best practices

Benefits of practitioner-led solutions

Introducing guidelines that make the traditional model validation process more suitable for algorithmic trading could have significant benefits, in terms of efficiency and appropriateness, as well as reducing the risks of market abuse and potential threats to market stability. Such standards could ensure firms use appropriate data inputs and have controls over the appropriate use of model type and assumptions. They could also create a common understanding of how best to test whether models and model components are robust in all market conditions, through appropriate stress testing. Where there are existing model risk teams, ensuring there is a suitable level of integration with algorithmic trading oversight committees, so that there is a consolidated approach to governance frameworks, would avoid duplication from the second line of defence.

The benefits of such guidelines are likely to increase in the coming years, as the level of complexity in algorithmic market making continues to rise. Model risks, data quality issues and the need for transparency are likely to be greater as algorithmic market making expands to a wider set of less liquid products, asset classes and geographies. There is considerable debate about how more complex machine learning techniques should be governed. Many market practitioners believe that existing governance arrangements with a tighter 'sandbox' in terms of controls and limits are appropriate. Others believe that machine learning can create new market fairness and stability threats that require a new distinct governance framework. It is too early in the evolution and usage of these new techniques to be definitive either way but there are likely to be new model risks especially related to the more limited transparency.

The use of execution algorithms must be properly aligned with asset managers' specific execution policies and strategy. It is important to ensure clarity about when it is appropriate for execution algorithms to direct flow to an in-house principal desk, and controls over the sharing of potentially inappropriate pre-trade information are also issues.

In summary, the increasing usage of algorithmic trading and the growing complexity of models makes the topics and emerging themes discussed in this Spotlight Review extremely important. Areas of such rapid technology change are also often best addressed by market practitioners with deep domain expertise who can develop solutions that are clear, practical and proactive in managing risks.

End notes

1 'Algorithmic Trading Compliance in Wholesale Markets', February 2018, available at www.fca.org.uk/publications/multi-firm-reviews/algorithmic-trading-compliance-wholesale-markets

2 Supervisory Statement 5/18 'Algorithmic trading', June 2018, available at www.bankofengland.co.uk/prudential-regulation/publication/2018/algorithmic-trading-ss

3 'Knight Capital Group provides update regarding August 1st disruption to routing in NYSE-listed securities', 2 August 2012, available at www.sec.gov/Archives/edgar/data/1060749/000119312512332176/d391111dex991.htm

4 Board of Governors of the Federal Reserve and Office of Comptroller of the Currency SR Letter 11-7 Attachment 'Supervisory Guidance on Model Risk Management', 4 April 2011, available at www.federalreserve.gov/supervisionreg/srletters/sr1107a1.pdf

5 'Joint Staff Report: The U.S. Treasury Market on October 15, 2014', July 2015, available at www.treasury.gov/press-center/press-releases/Documents/Joint_Staff_Report_Treasury_10-15-2015.pdf

6 Bank for International Settlements 'The sterling 'flash event' of 7 October 2016', January 2017, available at www.bis.org/publ/mktc09.pdf

7 International Monetary Fund Working Paper 19/136 'Anatomy of Sudden Yen Appreciations', July 2019, available at www.imf.org/en/Publications/WP/Issues/2019/07/01/Anatomy-of-Sudden-Yen-Appreciations-46944

8 July 2018, available at fmsb.com/wp-content/uploads/2018/11/FMSB_BCA_Book_Final_v2.pdf

9 Available at www.bankofengland.co.uk/report/2019/machine-learning-in-uk-financialservices

10 Algo wheels are used to automatically select which algorithms to use and are explained in more detail in the section on execution algorithms.

11 See also the summary of speeches by Mark Yallop, FMSB Chair, titled 'Behaving Fairly: Artificial Intelligence and Conduct in Wholesale Markets' on page 21 of the FMSB Annual Report 2019: fmsb.com/wp-content/uploads/2020/02/FMSB-Annual-Report-2019_Final.pdf

12 Available at www.bankofengland.co.uk/working-paper/2019/machine-learning-explainability-in-finance-an-application-to-default-risk-analysis

13 https://ec.europa.eu/finance/securities/docs/isd/mifid/rts/160608-rts-28_en.pdf

14 https://ec.europa.eu/finance/securities/docs/isd/mifid/rts/160719-rts-6_en.pdf

luminous

Design and production
www.luminous.co.uk

125 Old Broad Street, London EC2N 1AR
secretariat@fmsb.com
+44 (0)20 3961 6150
fmsb.com