FMSB
FICC MARKETS STANDARDS BOARD

Spotlight Review

Monitoring FICC markets and the impact of machine learning

August 2020

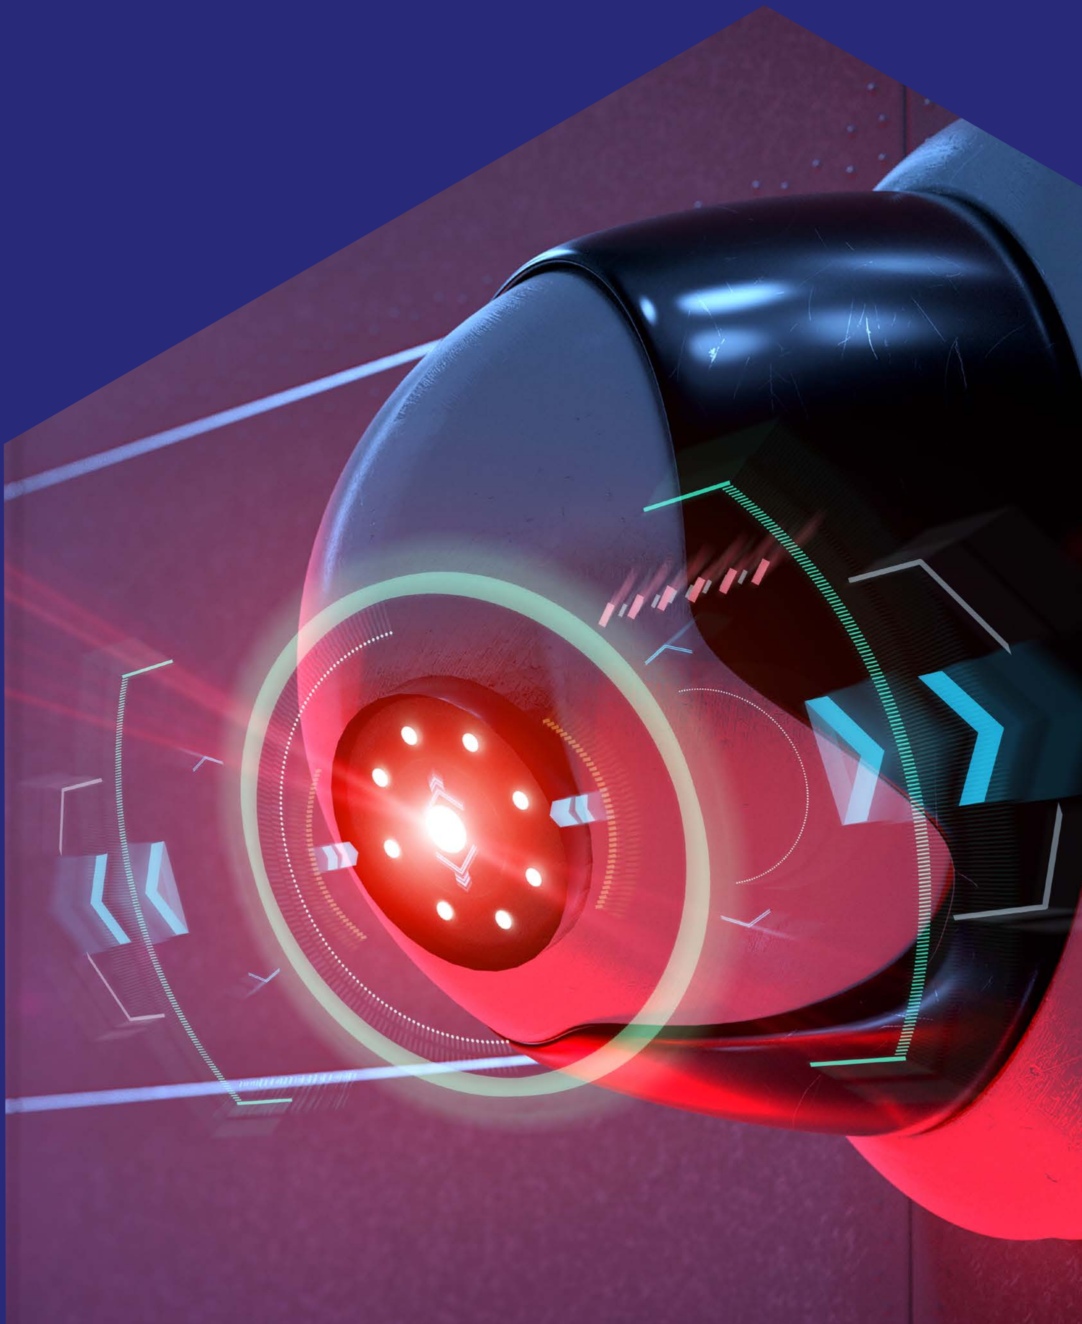

About FMSB

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 raising standards of behaviour and improving the quality, clarity and market-wide understanding of FICC trading practices. Producing guidelines, practical case studies and other materials that promote the delivery of transparent, fair and effective trading practices will help increase 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 committees, sub-committees 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.
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 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 are not intended to 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 structures 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.

Special acknowledgement

FMSB would like to thank The Alan Turing Institute for their time and input on the area of machine learning in this review.

Contents

1

Executive summary

Market surveillance in FICC has undergone, and continues to undergo, significant change as a result of regulation, the evolution of market structure, and technological developments. This Spotlight Review considers these structural and technological changes, in particular the emergence of machine learning trading strategies, and sets out some of the challenges associated with these developments for surveillance teams in FICC markets. The review then examines the role of technology as a potential solution to these challenges, creating as it does opportunities to improve market surveillance through the application of machine learning.

Market surveillance has been an area of regulatory focus over recent years. In particular, the UK's Financial Conduct Authority (FCA) has focused on the need for firms to continue to improve surveillance in FICC markets and to enhance both the quality and number of suspicious trade submissions relating to FICC markets activity. Furthermore, in the current remote working context, maintaining robust surveillance and suspicious transaction and order reporting has been flagged as a regulatory priority.[1]

However, further enhancing FICC surveillance in increasingly fast moving, complex and data-driven markets is not a simple task. This paper draws on two key structural challenges surveillance teams face in FICC markets: namely data quality and availability, and the increasing sophistication of trading strategies and technologies deployed to support them. The amount of data available to surveillance functions has increased significantly in recent years, driven by regulatory reporting requirements and the proliferation of electronic trading. However, extrapolating signals from these data sets remains challenging given variances in the accuracy, robustness, timeliness and consistency of such data, in particular across different FICC asset classes. Furthermore, growth in algorithmic trading, systematic investment strategies and the nascent adoption of machine learning in trading is materially increasing the speed and complexity of FICC markets. This combination of increased data and trading complexity, and the possibility of new market abuse risks emerging as a result of these developments, may drive the adoption of new, or the improvement of existing, surveillance techniques.

Technological innovations are creating opportunities to change market surveillance. The rudimentary nature of traditional automated alert systems produced a remarkably high number of alerts, but with only a small number translating into suspicious transaction and order reporting (STORs) being captured, reported and investigated in FICC markets. Machine learning techniques, with their ability to process large, complex data sets efficiently from both structured and unstructured data sources, offer the opportunity to make surveillance significantly more effective.

Given the increasingly data driven nature of FICC markets and the potential for technological developments to significantly change the nature of market surveillance, a greater understanding of data science and technology is becoming central to the future of market surveillance professionals. However, it is likely that the full potential of the application of machine learning techniques to market surveillance will only be realised in the long term.

This Spotlight Review aims to create further discussion on this topic and its relevance to future standards work by FMSB. This review will be of interest to those senior managers with supervisory responsibility for managing conduct risk in FICC trading businesses, second lines of defence in compliance and surveillance, and those working more broadly on the application of machine learning across financial market participants.

This review outlines:

- factors driving the pace of change in market surveillance;
- the acute impact of data on surveillance effectiveness;
- surveillance of complex algorithms and machine learning;
- employing machine learning to empower surveillance; and
- the vital role agility plays in effective surveillance.

Contents

Factors driving the pace of change in market surveillance

The rate of evolution of wholesale FICC markets continues to quicken. Driven by a multitude of factors including innovation in traded products, new market structures and trading platforms, advances in technology and regulatory change, market participants are constantly revising how they trade in order to service clients, hedge risks and remain competitive. But with all this rapid change, there are enormous challenges for compliance functions in the second line of defence in staying one step ahead and maintaining the necessary levels of agility to monitor front office activity effectively.

In recent years there has been significant regulatory focus on surveillance in FICC markets, including the requirements imposed by the Market Abuse Regulation (MAR)[2] which have dramatically increased spend on compliance and market surveillance. However, despite this investment, some firms remain dependent on legacy systems which saturate surveillance teams with false positives. Inflexible surveillance technology systems do not support the integration of trade and communication data sources. With the emphasis on innovation, the speed of deployment of new technologies and algorithmic tools on trading desks can quickly outstrip corresponding developments in surveillance.

Not seeing the wood for the trees

In the UK, the FCA has been particularly focused on the need for financial institutions to improve both the quality and number of STORs relating to FICC markets activity.

The FCA addressed its concerns around fixed income surveillance in September 2018:[3]

"STOR submission across asset classes remains inconsistent and we believe submissions are lower than they should be in some areas. In particular, our view is that submissions continue to be too low in fixed income products and we wish to provide some further observations from our recent visit programme where we have focused on fixed income markets.

In Market Watch 50, we explained that firms are often over-prescriptive with analysts and do not encourage them to look beyond the initial alert. Our observations indicate this continues to be the case with some firms. In some fixed income markets, for example, some analysts tended to take a narrow approach, reviewing only the activity in the product which triggered the alert and not considering other trading in correlated products. Because many fixed income products are inter-connected, consideration of trading activity *in correlated products – such as cash vs futures, or products with different durations – is an important element of effective surveillance."*

The concerns around fixed income surveillance were reiterated in a speech by Julia Hoggett, Director of Market Oversight at the FCA in February 2019:[4]

"…there is still more for the industry to do to improve its capacity to surveil for market manipulation – as opposed to insider dealing – and there remains a need to improve surveillance in the non-equity space…Insider dealing STORs in 2018 accounted for 86% of all STORs."

The FCA highlighted that it believed the key reasons behind the low level of STORs and high percentage of false positives is an over-reliance on the list of indicators of market abuse outlined in the guidelines on the MAR[5] being treated as exhaustive, firms being too dependent on current off-the-shelf calibration setting of software vendors for their alert parameters and not making sure that the techniques and methodologies used are appropriate to the products under surveillance. It would seem that a more tailored, discerning approach is called for, but this could bring the risk of firms adopting very different approaches to meeting the MAR guidelines.

Meeting surveillance challenges brought by remote working

The onset of the COVID-19 pandemic and the necessary rapid adoption of remote working arrangements brought into focus the challenges that market surveillance teams face in adapting to fast changing conduct risk landscapes. Following the introduction of government restrictions on movement, trading functions became geographically dispersed, resulting in trading taking place outside the usual control environment. These rapid developments have brought new challenges to control functions and surveillance teams, not least:

- the use of restricted or non-recorded voice or electronic communication channels with clients, creating gaps in audit trails; and
- new risks to the protection of material non-public information and client confidentiality, for instance where market participants from different parts of the same firm, different sell-side firms, or sell-side and buy-side firms share accommodation.

Furthermore, during this period of prolonged and widespread remote working, the FCA has reiterated the importance of maintaining robust market surveillance, controls around market abuse, conduct and managing conflicts of interest.[6]

Contents

The acute impact of data on surveillance effectiveness

This section highlights the challenges for surveillance teams in monitoring FICC markets with increasing amounts of data deriving from a multitude of structured and unstructured sources.

Data is a key factor shaping the ability of firms to monitor FICC markets effectively. However, firms face significant issues in sourcing relevant external market data, adjusting for inconsistencies in data collected from different business lines and internal systems and dealing with ineffective data screening and analysis.

This section outlines a checklist of the six critical factors which determine the quality of data and thus the overall effectiveness of a firm's market surveillance capability:

- i Amount and completeness of data
- ii Accuracy and robustness of data
- iii Relevance of data
- iv Appropriate use of unique and correlated data
- v Timeliness of data
- vi Consistency of data

Amount and completeness of data

The amount of data available to the compliance and market surveillance functions of large market participants has increased materially in recent years, in line with the emergence of electronic trading and new reporting rules.

However, the availability of transaction data varies significantly by asset class as a result of differing liquidity and activity levels. Foreign exchange, for example, is highly liquid, while certain corporate bonds may not trade for days (and in many cases weeks or months). For less liquid products transaction-related market data is therefore limited. Moreover, getting relevant data on block size liquidity remains challenging.

FICC over-the-counter (OTC) products are predominantly traded bilaterally through disclosed channels, even in highly liquid instruments where electronic trading dominates. For instance, in foreign exchange, Central Limit Order Books (CLOBs) compose only a small proportion of market volumes. CLOBs produce significant amounts of publicly available market data, where prices are not specific to any counterparty and hence are a useful tool for surveillance. Disclosed liquidity forms of electronic trading typically have prices specific to the counterparty involved, although there has been a growth in pre-trade composite pricing. At the same time, there is limited access to quote or order data on most venues. There are also many platforms that do not provide publicly available market data.

With such rich diversity in trading activity, firms must make sure they have considered carefully the full range of alternative sources of data feeding their surveillance engines.

Contents

ii

Accuracy and robustness of data

The varying levels of liquidity and the bilateral nature of most FICC markets creates significant differences in the availability of accurate and robust transaction data across FICC asset classes. A systematic and rigorous approach taking into account the nature of the asset class in question is therefore crucial to understand the data and intentions that lie behind trade decisions. Examples of areas of focus for data accuracy include:

- **Communications on actual order/trade data** – the FCA has highlighted market abuse concerns around how a firm communicates to its clients or other market participants (via screen, instant message or voice), that it has bids or offers when they are not supported by, or derived from, an order, or that a trade has been executed at a specified price and/or size when no such trade has taken place.[7]
- **Use of non-transaction prices** – given the significant use of indicative, consensus or evaluated (model based) prices, it is important that surveillance is focused on ensuring that these data sources are robust. They should be free from any inherent bias or threat of manipulation. This may be in the form of rogue inputters into any consensus pricing products or spoofing type behaviour to manipulate reference prices. A fast-growing approach to execution in recent years has been interdealer broker mid-market auctions, particularly in less liquid markets. In many cases, these are led by indicative mid-prices from brokers and here it is important to ensure that the methodology is transparent.
- **Growth of alternative data sources** – the significant growth of new unstructured data poses new challenges. The creation of fake or distorted data, for example on social media with the intent to influence the price of bonds, currencies and commodities could constitute market manipulation. Furthermore, the growth of unstructured data online makes the ability to spread false information and scale these actions easier and greater than ever before.

iii Relevance of data

It is crucial when conducting market surveillance to be able to find the most appropriate data set and understand its relevance. In many financial instruments there may be sporadic transaction data, particularly in the relevant order size. For instance, transactions could be small size in a specific bond issue but to determine prices in larger sizes it may be necessary to see what the price of block size liquidity was in other, similar bonds. Moreover, in more volatile and stressed markets, it is important to determine what is a genuine market dislocation (e.g. a much quicker sell off in exchange-traded fund prices than in the underlying bond prices) as opposed to what could amount to manipulative activity.

Contents

iv Appropriate use of unique and correlated data

Unique data refers to market data on price-forming trades, while correlated data refers to the prices in financial instruments or on venues that are traded relative to each other. This increased complexity and interdependence makes it important for surveillance professionals to understand the sources of price-forming trades and to avoid conducting surveillance of one financial instrument or venue in isolation.

In recent years there has been an increase in the amount of trading activity across products and venues, fuelled by the growth of electronic and algorithmic trading. This has typically involved trading in correlated financial instruments (for instance, cash bonds and related markets such as futures and OTC derivatives) but increasingly it also involves trading across different asset classes that are correlated (for instance, commodity currencies such as the Australian dollar, South African rand and Brazilian real against cash or derivatives prices of related commodities such as gold and iron ore). Strategies deployed by market participants can vary from latency arbitrage to relative value approaches. However, where market participants use quote or transaction activity in one market or benchmark to influence the price of other correlated markets there is a risk of market manipulation. Surveillance professionals increasingly have to watch relevant trade data not just in one area but across all related instruments.

Where there is considerable venue fragmentation it is also important to look at how activity across multiple venues can impact each other. For instance, moving the quoted bid or ask on a primary venue and taking advantage of this move in the market price to make profitable and much larger disclosed liquidity trades. The risk of this may be higher in out-of-hours trading when CLOB market depth is thin. This risk highlights the need for data users not to be overly reliant on quotes on any single primary venue especially when it contributes a small percentage of market volumes in that asset class.

v Timeliness of data	Inconsistent time stamping, especially in emerging markets products and voice-brokered trades, can make it difficult to determine an accurate trail of illegitimate activity within trade data.
vi Consistency of data	Inconsistencies in the data reported by different functions/parts of a firm and different market venues also pose challenges for market surveillance teams. The granularity and type of data reported through a firm's front office trade capture systems tends to be very different from middle office surveillance systems. Data reported for surveillance purposes may not match up with trade reporting data in certain jurisdictions. Market fragmentation across multiple venues can also pose consistency challenges in terms of different venues having different rulebooks and levels of data quality. Although there is significant variability across jurisdictions and asset classes, these disparities are addressed to a degree in some areas by common standards. The Trade Reporting and Compliance Engine (TRACE) standard in certain US dollar securities is an example of this. Additionally, the movement towards new cloud-based data lakes as a centralised repository of data should help drive consistencies in data reporting across firms.

Contents

Surveillance of complex algorithms and machine learning

This section considers the impact of increasingly sophisticated trading strategies and, in particular, the adoption of machine learning on the ability of surveillance teams to monitor FICC markets effectively.

FMSB's Spotlight Review 'Emerging themes and challenges in algorithmic trading and machine learning'[8] noted that while the adoption of machine learning techniques for trading and investing in FICC markets is currently relatively low it is likely to increase considerably over time. The review highlighted that the complexity of machine learning can lead to increased model risks when it is deployed in algorithmic trading. This is also a topic of interest for other standard setting bodies such as the International Organization of Securities Commissions (IOSCO) which has published work looking at identifying and addressing potential conduct risks relating to the use of machine learning by market participants.[9]

This section focuses on the market abuse risks and challenges faced when conducting market surveillance of machine learning based investing and trading strategies.

Machine learning methods pose the following unique challenges for market surveillance:

- i Evidence of intent
- ii Complexity
- iii Learning to 'game' the system
- iv Collusion

Evidence of intent

Machine learning techniques inherently lack a mechanistic explanation of processes and how they lead to specific outcomes. In contrast, in rules-based algorithmic trading it is possible to trace how input data is directly impacting the output, the more linear cause and effect relationship making it easier to see if software is being developed that will result in illegitimate activity. This makes the approval process when using machine learning algorithms particularly important.

Moreover, machine learning and rules-based algorithmic trading, do not leave the 'breadcrumb' trails that manual traders do. Often the firm is alerted to human trader misconduct through email and instant messenger communications which are absent in the systems-driven decisions. Firms must ensure that the method by which decisions are reached by automated trading systems can be adequately and quickly interrogated and any inappropriate trading 'intent' weeded out.

Contents

4

ii

Complexity

To date, where there has been an adoption of machine learning in FICC market making, it has typically been for one component of the overall algorithmic trading strategy and the data sources used have been highly restrictive with a particular emphasis on quality control.

As more complex machine learning techniques are adopted, the risk of a small variation of input data leading to dramatically different outputs increases. If the algorithm is not robust with respect to data tampering it may therefore lead to unintended consequences.[10] Governance surrounding algorithmic market making already focuses on the sensitivity of outputs to certain key data input features and such governance is likely to be of increasing importance to act as a check on complex machine learning models.

The risk of machine learning adding complexity will vary depending on where and how it is being deployed. For instance, as systematic trading models deployed by the buy-side for investment decisions incorporate machine learning, they are likely to use a huge amount of data, often from disparate sources. When these self-learning machines are looking at correlations between unstructured data sources online and market prices, they may not always be able to differentiate between when the source is public or private and whether it should have been ringfenced or excluded from the decision process.

iii

Learning to 'game the system'

There is a risk that without the conscience of a manual trader or the clear roadmap of rules-based algorithms, machine learning techniques may not be able to understand the limits of permissibility. There may be a need for an 'ethical governor' that tests the optimisation process against ethical benchmarks and rejects trading tactics that fall short of these standards. This 'ethical governor' that restricts the freedom of action of self-learning machines could also be deployed across rules-based algorithms at the same time to seek uniformity in controls. One potential approach could be to include historical reference cases such as those in FMSB's Behavioural Cluster Analysis (BCA)[11] in the programming process, to train self-learning algorithms to have an idea of what illegal behaviour has looked like in the past. Nevertheless, operationalising ethical principles has been one of the greatest challenges so far faced by research in this area.[12]

At the same time some strategies, such as reinforcement-learning algorithms, seek the optimal strategy and/or action to maximise a predefined objective function. The ability of such self-learning machines to exploit fault lines is materially greater than manual traders or traditional algorithmic models that have a clearly defined remit. Such machine learning algorithms may systematically work out the parameters that compliance would naturally see as high conduct risk and look for gaps in this supervisory governance that it, or other self-learning machines, can find and exploit on a repeated basis – a kind of 'surveillance arbitrage'.

Contents

iv

Collusion

Self-learning machines can see patterns in data and are likely to identify other market participants that are investing or trading in a similar style to them. Manual traders have been doing this for hundreds of years, but the scale of data sets involved is now of a different magnitude and there is no 'evidence of intent' as highlighted earlier in this section. Surveillance of self-learning machines is entirely dependent on spotting trends in data.

There are some academic studies on potential game theory by self-learning machines and the scope to see collusive behaviour[13] but, given the limited use of machine learning in real-life trading, drawing any conclusions at this stage would be speculative. Moreover, it should be noted that adversarial networks (a technique employed in machine learning that attempts to fool models through malicious input) offer a good way to stress test models and determine how to control any collusive tendencies.

The observations outlined above should be taken within the context that computer-generated decisions can in principle be scrutinised, which might not always be the case for human decisions. Indeed, with the right governance, both the data used to train the algorithm, and the algorithm itself, can be investigated. There are many market participants that believe the existing governance and surveillance framework for algorithms should be extended to machine learning techniques but, as noted above, there are many new challenges as well. As the adoption of machine learning grows, more consensus needs to be built around the best practice in governance structures for managing these challenges.

Contents

Employing machine learning to empower surveillance

The rollout of new solutions in market surveillance is in its early days, and as with many promising new technologies, despite the marketing hype, it is not uncommon to see new vendor products fail to deliver on these promises when put to the test in real world market conditions. Implementation of these new technologies in a large multi-system, multi-jurisdiction environment is more cumbersome and complex than for smaller single firms with new technology infrastructure, which may have been the environment these tools were designed for and tested against. At the same time, there is scope for efficiency gains in the surveillance of FICC markets, and machine learning techniques have certain advantages over more traditional algorithmic surveillance systems. This section sets out an overview of the key benefits and risks of using machine learning in market surveillance.

Key benefits of using machine learning in market surveillance

1

Large and complex data sets

Machine learning could lead to fundamental changes in market surveillance given its ability to process complex, large and poorly structured datasets. FICC markets involve huge amounts of disparate structured and unstructured data including quotes, trades, email and voice communications data. In particular, machine learning techniques can proactively combine trade and communications data in a more systematic way through sophisticated natural language processing. The need for this will only increase with the growth of alternative data.

The October 2019 joint Bank of England and FCA report 'Machine learning in UK financial services'[14] incorporated a comprehensive survey, which highlighted the growing usage of machine learning from a low base. Moreover, it stated that the highest penetration of machine learning was in areas such as fraud detection and anti-money laundering (AML), which have similarities to market surveillance in the range and complexity of data sets.

Traditional screening of transactions for AML purposes typically leads to a high number of false positives, a pattern also seen in surveillance of financial markets. Perhaps an even more pressing concern is where market abuse goes undetected ('false negatives'). With machine learning and adversarial learning, firms can target illicit market behaviour through analysis of patterns of behaviour, sometimes in absence of hard and fast evidence of wrongdoing.

Another challenge, also present in AML, is the scarcity of labelled data sets, which are needed to train many machine learning algorithms. One solution developed for this in recent years has been the use of data sets with synthetically engineered features for testing purposes. This could be expanded to test and improve upon existing surveillance systems.

2

Learning new surveillance techniques

FICC markets are constantly changing with frequent emergence of new algorithmic and systematic trading strategies. Moreover, market structures are evolving, with fragmentation of trade venues, the proliferation of data sources and shifting correlations between financial instruments. In this context, rigid surveillance techniques struggle to keep up with the evolving landscape. Machine learning programmes are inherently less reliant than traditional algorithms on following a set of rules and can learn from experience and better leverage historical data including by finding relevant market conditions or structures.

Contents

3

Monitoring self-learning machines

Electronic and algorithmic trading have substantially increased the amount of data and diversity of data sources. This will only increase further with the adoption of machine learning strategies in investing and trading. Machine learning algorithms may be more likely than manual surveillance methods to identify trading patterns and the context in which such trades take place in an environment where there is less manual trader intervention and associated unstructured communications data to give the trading context.

Testing of machine learning systems using digital sandpits can help provide safe environments in which to identify potential risks. Digital sandpits are based around real data sets, together with synthetically generated scenarios, to enable testing under market conditions that are not present in historical data sets. Here adversarial learning provides an automated way of seeking such 'black swan'[15] scenarios.[16] The technique can also be used to stress test currently used rule-based models for detecting anomalies in the market. Digital sandpits are already in use in AML applications as per the FCA sandpit[17] and there is scope for this approach to be adapted to the surveillance of FICC markets.

Academic bodies such as The Alan Turing Institute are currently conducting research using several machine learning model types that focus on anomaly detection and identifying collusive behaviours.[18] These include multi-agent systems, seeking to simulate the behaviour of individual actors and how their interaction impacts the overall system. Added to this, the use of adversarial training creates dynamic data generators for testing, validation, and monitoring of live machine learning systems.

Key risks of using machine learning in market surveillance

1

The 'black box' dilemma

A key risk factor of relying on machine learning for surveillance is the 'black box' nature of such software. It creates a lack of transparency in areas where regulatory requirements, particularly those related to the ownership of supervisory duties by the front office, have increased significantly in recent years. A clear cause and effect relationship of how machine learning systems conduct surveillance may not be possible, but just as with its use in other fields, this can be compensated for by rigorous back testing and scenario analysis of the effectiveness of such self-learning machines. This would cover both the absolute and relative performance of these self-learning machines in terms of the number of STORs generated and the reduction of false positives. Such back testing and scenario analysis would also improve the ability of self-learning machines to detect different kinds of market abuses, including emerging threats and behaviours occurring in different market conditions.

Scenario analysis and testing should seek to leverage both historical training and synthetic data to reflect environments not seen in the past. The creation of synthetic market abuse scenarios is complex, requiring more data sources and more sophisticated models than is the case when simply fabricating market conditions.

Testing should also compare performance relative to rules-based algorithmic alert systems. In August 2019 the Bank of England published Staff Working Paper No. 816 'Machine learning explainability in finance: an application to default risk analysis'[19] which explored these topics in detail:

Contents

1

The 'black box' dilemma (continued)

"We propose a framework for addressing the 'black box' problem present in some Machine Learning (ML) applications. We implement our approach by using the Quantitative Input Influence (QII) method of Datta et al (2016) in a real-world example: a ML model to predict mortgage defaults. This method investigates the inputs and outputs of the model, but not its inner workings… We use clustering methods to arrive at groups of explanations for different areas of the input space. Finally, we conduct simulations on data that the model has not been trained or tested on. Our main contribution is to develop a systematic analytical framework that could be used for approaching explainability questions in real world financial applications. We conclude though that notable model uncertainties do remain which stakeholders ought to be aware of."

Ongoing research into new ways to aid explainability could be vital in supporting growth in machine learning innovation. New technologies designed to support tamper-proof transparency, such as the distributed ledger, could potentially also play a role in supporting improved explainability and transparency of machine learning algorithms.

2

Capturing 'black swans'

'Black swans' refer to an unpredictable event, with potentially severe consequences, that is outside the normally expected distribution of outcomes. The COVID-19 pandemic may be characterised as an example.

Machine learning programmes may be limited in their ability to adapt to unpredictable events due to their reliance on historical market data and emerging trends to underpin pattern recognition. Machine learning may therefore prove adept at spotting market abuse but less effective in assessing broader market stability risks, especially where data labelling is unclear.

However, machine learning programmes may permit more flexibility in the modelling of risks and interactions between different parts of the financial system, and if used judiciously, could prove better at proactively identifying potential weaknesses than rules-based algorithms.

Contents

This Spotlight Review highlights the significant challenges that firms face in ensuring their surveillance capability matches the fast pace of innovation in the front office. Not least in the face of increased use and sophistication of algorithmic trading and machine learning technologies.

It has examined how poor management of the six key factors affecting the quality of data poses significant threats to effective surveillance. Although the amount of available data has increased significantly in recent years, there remain challenges in terms of accessing accurate, relevant data in a timely fashion. The huge amounts of structured and unstructured data also create noise, making it difficult to extract the data signals necessary to isolate and identify suspicious activity. Furthermore, the increasing complexity of trading strategies and the nascent deployment of machine learning techniques creates new challenges related to evidence of intent, complexity, and the risk of self-learning machines actively choosing to manipulate markets.

Given the view of regulators that there is a need for financial institutions to improve the suspicious trade submissions relating to FICC markets activity, and the ability of machine learning programmes to process large complex data sets efficiently, it is highly likely that machine learning will play a role in the future of market surveillance of market abuse risks. Working side by side with humans, over time, machine learning programmes may be better able to understand the semantics of data and the evolution of behavioural patterns and to adapt their machine learning algorithms. Consequently, a greater understanding of data science and technology is becoming central to the future of market surveillance professionals so they can effectively specify and test machine learning functionalities.

As the adoption of machine learning techniques in financial services matures in terms of usage, best practice on safe and effective deployment will emerge, and there may be an important role for practitioner-led industry standards to improve the consistency and effectiveness of market surveillance in fast-developing markets.

End notes

1 Market Watch No. 63 'Newsletter on market conduct and transaction reporting issues', May 2020, available at www.fca.org.uk/publication/newsletters/market-watch-63.pdf

2 Regulation (EU) No 596/2014 of the European Parliament and of the Council of 16 April 2014, available at https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32014R0596&from=EN

3 Market Watch No. 56 'Newsletter on market conduct and transaction reporting issues', September 2018, available at www.fca.org.uk/publication/newsletters/market-watch-56.pdf

4 'Market abuse requires a dynamic response to a changing risk profile', available at www.fca.org.uk/news/speeches/market-abuse-requires-dynamic-response-changing-risk-profile

5 European Securities and Markets Authority 'Final Report: Guidelines on the Market Abuse Regulation – market soundings and delay of disclosure of insider information', 13 July 2016, available at www.esma.europa.eu/sites/default/files/library/2016-1130_final_report_on_mar_guidelines.pdf

6 Supra note 1.

7 Market Watch No. 57 'Newsletter on market conduct and transaction reporting issues', November 2018, available at www.fca.org.uk/publication/newsletters/market-watch-57.pdf

8 April 2020, available at fmsb.com/our-publications/fmsb-spotlight-review-emerging-themes-and-challenges-in-algorithmic-trading-and-machine-learning/

9 IOSCO consultation paper CR02/2020 'The use of artificial intelligence and machine learning by market intermediaries and asset managers', 25 June 2020, available at www.iosco.org/library/pubdocs/pdf/IOSCOPD658.pdf

10 'Adversarial Attacks on Machine Learning Systems for High-Frequency Trading' by Micah Goldblum, Avi Schwarzchild, Ankit B. Patel and Tom Goldstein, 4 March 2020, available at https://arxiv.org/pdf/2002.09565.pdf

11 Supra note 1.

12 'Provably Beneficial Artificial Intelligence' by Stuart Russell, March 2017, available at https://people.eecs.berkeley.edu/-russell/papers/russell-bbvabook17-pbai.pdf

13 For example 'Artificial Intelligence, Algorithmic Pricing and Collusion' by Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pstorello, April 2019, available at www.tilburguniversity.edu/sites/default/files/download/Pastorello%20-%20Qlearning_3.pdf

14 Available at www.fca.org.uk/publication/research/research-note-on-machine-learning-in-uk-financial-services.pdf

15 A theory developed by Nassim Nicholas Taleb 'black swans' refer to an unpredictable event, with potentially severe consequences, that is outside the normally expected distribution of outcomes.

16 'Adversarial Machine Learning at Scale' by Samy Bengio, Ian J Goodfellow, and Alexey Kurakin, February 2017, available at https://arxiv.org/pdf/1611.01236.pdf

17 See '2019 Global AML and Financial Crime TechSprint' webpage referring to the event held from 29 July to 2 August 2019, available at www.fca.org.uk/events/techsprints/2019-global-aml-and-financial-crime-techsprint

18 'Anomaly detection in networks with application to financial transaction networks' by Andrew Elliott, , Mihai Cucuringu, Milton Martinez Luaces, Paul Reidy, and Gesine Reinert, May 2019, available on The Alan Turing Institute's website at www.turing.ac.uk/research/publications/anomaly-detection-networks-application-financial-transaction-networks

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

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