

Artificial Intelligence (AI) in the Securities Industry[1]

JUNE 2020

A REPORT FROM **THE FINANCIAL INDUSTRY REGULATORY AUTHORITY**

Contents

Introduction 1

SECTION I: Overview of Artificial Intelligence Technology 2

- Key Components of AI Applications 4

SECTION II: AI Applications in the Securities Industry 5

- Communications with Customers 5
- Investment Processes 7
- Operational Functions 8

SECTION III: Key Challenges and Regulatory Considerations 11

- Model Risk Management 11
- Data Governance 13
- Customer Privacy 15
- Supervisory Control Systems 16
- Additional Considerations 18

Request for Comments 20

Introduction

Artificial Intelligence (AI) technology is transforming the financial services industry across the globe. Financial institutions are allocating significant resources to exploring, developing, and deploying AI-based applications to offer innovative new products, increase revenues, cut costs, and improve customer service.[2] First developed in the early 1940s, AI technology has gained significant momentum over the past decade and become more mainstream due in part to the availability of inexpensive computing power, large datasets, cloud storage, and sophisticated open-source algorithms. In a recent survey-based report, executives at financial institutions noted that "AI is expected to turn into an essential business driver across the Financial Services industry in the short run, with 77% of all respondents anticipating AI to possess high or very high overall importance to their businesses within two years."[3]

Broker-dealers are exploring and deploying AI-based applications across different functions of their organizations, including customer facing, investment, and operational activities. In July 2018, FINRA solicited comments from the industry on the potential challenges associated with using and supervising AI applications at broker-dealer firms.[4] In response, commenters recommended that FINRA undertake a broad review of the use of AI in the securities industry to better understand the varied applications of the technology, their associated challenges, and the measures taken by broker-dealers to address those challenges. Based on this feedback, FINRA, through its Office of Financial Innovation (OFI), engaged in an active dialogue with the industry over the past year and held meetings with over two dozen market participants, including broker-dealer firms, academics, technology vendors, and service providers to learn more about the use of AI in the securities industry.

[1] This paper is not intended to express any legal position and does not create any new requirements or suggest any change in any existing regulatory obligations, nor does it provide relief from any regulatory obligations. While this paper summarizes key findings from FINRA's outreach and research on the use of AI applications in the securities industry, it does not endorse or validate the use or effectiveness of any of these applications. Further, while the paper highlights certain regulatory and implementation areas that broker-dealers may wish to consider as they adopt AI, the paper does not cover all applicable regulatory requirements or considerations. FINRA encourages firms to conduct a comprehensive review of all applicable securities laws, rules, and regulations to determine potential implications of implementing AI-based applications.

[2] PwC, Crossing the Lines: How Fintech is Propelling FS and TMT Firms Out of Their Lanes, 2019, PDF (reporting that financial services executives expect their AI efforts to result in increased revenue and profits (50%), better customer experiences (48%), and innovative new products (42%).).

[3] Ryll et al., Transforming Paradigms: A Global AI in Financial Services Survey, Jan. 2020, PDF.

[4] FINRA, Special Notice on Financial Technology Innovation in the Broker-Dealer Industry, July 30, 2018, Article.

This paper summarizes key findings from FINRA's review in three sections:

- Section I briefly defines AI and its scope as it pertains to the securities industry for the purposes of this paper.
- Section II provides an overview of broker-dealers' use of AI applications related to: (i) communications with customers, (ii) investment processes, and (iii) operational functions.
- Finally, Section III discusses key factors including potential regulatory considerations, securities market participants may want to consider as they develop and adopt AI-based tools.

The discussion below is intended to be an initial contribution to an ongoing dialogue with market participants about the use of AI in the securities industry. Accordingly, FINRA requests comments on all areas covered by this paper.[5] FINRA also requests comments on any matters for which it would be appropriate to consider guidance, consistent with the principles of investor protection and market integrity, related to AI applications and their implications for FINRA rules.

I. Overview of Artificial Intelligence Technology

Definition

The term artificial intelligence broadly refers to applications of technology to perform tasks that resemble human cognitive function and is generally defined as "[t]he capability of a machine to imitate intelligent human behavior."[6] AI typically involves "[t]he theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages."[7] John McCarthy, one of the founders of AI research, "once defined the field as getting a computer to do things which, when done by people, are said to involve intelligence."[8]

Scope[9]

While the definitions for AI discussed above provide a general outline of the meaning of the term, there is no single universally agreed upon definition of AI. In practice, AI is used as an umbrella term that encompasses a broad spectrum of different technologies and applications, some of which are described below.

- **Machine Learning (ML) –** Machine learning is a field of computer science that uses algorithms to process large amounts of data and learn from it. Unlike traditional rules-based programming, ML models[10] learn from input data to make predictions or identify meaningful patterns without being explicitly programmed to do so. There are different types of ML models, depending on their intended function and structure:

[5] See Request for Comments section on page 20 of this paper.

[6] Artificial Intelligence, Merriam Webster, Article.

[7] Artificial Intelligence, Oxford English Dictionary, Article.

[8] Shukla Shubhendu & Jaiswal Vijay, Applicability of Artificial Intelligence in Different Fields of Life, International Journal of Scientific Engineering and Research, 29–35 (2013), PDF.

[9] The definition and scope of AI presented here are intended purely to frame the discussion in this document and should not be interpreted as guidance. In our discussions with industry participants, there is a wide spectrum of viewpoints with no consensus on the definition or scope of the technology.

[10] An ML model generally refers to the combination of input data, key features identified from the data, algorithms, parameters, and outputs that are collectively used to build the AI application.

- **Supervised Machine Learning –** In supervised ML, the model is trained with labeled input data that correlates to a specified output. For example, a dataset of animal photos (input data) can be labeled as "cats" or "not cats" (output data). The model is continuously refined to provide more accurate output as additional training data becomes available. After the model has learned from the patterns in the training data, it can then analyze additional data to produce the desired output. Results of supervised ML models are typically reviewed by humans for accuracy and fed back into the model for further refinement. Supervised ML is successful when the model can consistently produce accurate predictions when provided with new datasets. For example, the ML model learns to recognize if a new picture is a cat or not.
- **Unsupervised Machine Learning –** In unsupervised ML, the input data is not labeled nor is the output specified. Instead, the models are fed large amounts of raw data and the algorithms are designed to identify any underlying meaningful patterns. The algorithms may cluster similar data but do so without any preconceived notion of the output. For example, a time series of trade events can be inputted into an unsupervised model, with the model identifying groups of similar trades as well as outliers. Results of unsupervised machine learning models are then interpreted by humans to determine if they are meaningful and relevant.
- **Reinforcement Learning –** In reinforcement learning, the model learns dynamically to achieve the desired output through trial and error. If the model algorithm performs correctly and achieves the intended output, it is rewarded. Conversely, if it does not produce the desired output, it is penalized. Accordingly, the model learns over time to perform in a way that maximizes the net reward. For example, in the securities industry, reinforcement learning models are being explored for options pricing and hedging.[11]
- **Deep Learning –** A deep learning model is built on an artificial neural network, in which algorithms process large amounts of unlabeled or unstructured data through multiple layers of learning in a manner inspired by how neural networks function in the brain. These models are typically used when the underlying data is significantly large in volume, obtained from disparate sources, and may have different formats (*e.g.*, text, voice, and video). For example, some firms in the securities industry are developing surveillance and conduct monitoring tools built on deep learning models. Deep learning applications can be supervised, unsupervised, or reinforcement based.

- **Natural Language Processing (NLP) –** NLP is a form of AI that enables machines to read or recognize text and voice, extract value from it, and potentially convert information into a desired output format, such as text or voice. Examples of NLP applications in the securities industry range from keyword extraction from legal documents and language translation to more complex tasks, such as sentiment analysis and providing relevant information through chat-boxes and virtual assistants.
- **Computer Vision (CV) –** CV (also referred to as machine vision) is a "field of computer science that works on enabling computers to see, identify and process images in the same way that human vision does, and then provide appropriate output."[12] Frequently a CV application will use ML models to interpret what it "sees" and make predictions or determinations. Examples of CV-based applications include facial recognition, fingerprint recognition, optical character recognition, and other biometric tools to verify user identity.

[11] Halperin, Igor, QLBS: Q-Learner in the Black-Scholes (-Merton) Worlds, SSRN, Dec. 16, 2017, Article 1 or Article 2.

[12] What is Computer Vision?, Techopedia, Article.

- **Robotics Process Automation (RPA) –** RPA refers to the use of preprogrammed software tools that interact with other applications to automate labor-intensive tasks, resulting in increased accuracy, speed, and cost-savings. RPA tools are generally used for high-volume, repetitive processes involving structured data, such as account reconciliation, accounts payable processing, and depositing of checks. Some market participants do not consider RPA to be a form of AI because its focus is on automation of processes in a manner more akin to a rules-based system.[13] However, others consider it to be a rudimentary form of AI, particularly when it is combined with other technologies such as ML.

Key Components of AI Applications

AI applications generally involve the use of data, algorithms, and human feedback. Ensuring each of these components is appropriately structured and validated is important for the development and implementation of AI applications. The discussion that follows highlights how each of these components influences the development of AI applications.

- **Data –** Data generation in the financial services industry has grown exponentially over the past decade, in part due to the use of mobile technologies and the digitization of data. The importance of data has likewise rapidly increased, and some have even referred to data as a more valuable resource than oil.[14] Furthermore, cloud technology has enabled firms to collect, store, and analyze significantly large datasets at very low costs. Firms in the financial services industry now collect data from a variety of internal sources (*e.g.*, trading desks, customer account history, and communications) and external sources (*e.g.*, public filings, social media platforms, and satellite images) in both structured and unstructured formats, and analyze this data to identify opportunities for revenue generation as well as cost-savings. This explosion of data in the financial services industry is one of the key factors contributing to the increased exploration of AI in the industry.

  Data plays a critical role in the training and success of any AI application. AI applications are generally designed to analyze data by identifying patterns and to make determinations or predictions based on those patterns. The applications continuously and iteratively learn from any inaccurate determinations made by such applications, typically identified through human reviews as well as from new information, and refine the outputs accordingly. Therefore, AI applications are generally best positioned to yield meaningful results when the underlying datasets are substantially large, valid, and current.

- **Algorithms –** An algorithm is a set of well-defined, step-by-step instructions for a machine to solve a specific problem and generate an output using a set of input data. AI algorithms, particularly those used for ML, involve complex mathematical code designed to enable the machines to continuously learn from new input data and develop new or adjusted output based on the learnings. An AI algorithm is "not programmed to perform a task, but is programmed to learn to perform the task."[15] The availability of open-source AI algorithms, including those from some of the largest technology companies, has helped fueled AI innovation and made the technology more accessible to the financial industry.

[13] SIFMA, Re: Special Notice on Financial Technology Innovation in the Broker-Dealer Industry, Oct. 18, 2018, PDF (stating in its comment letter, published on July 30, 2018, that "…innovations such as the majority of robotic process automation ("RPA") do not use AI, but nevertheless equally deserve to be considered with regards to matters of supervision."). FINRA does not take an explicit view on the status of RPA in the context of AI, but for purposes of this report the use of the term AI does not encompass applications involving basic RPA.

[14] The World's Most Valuable Resource is no Longer Oil, but Data, The Economist, May 6, 2017, Article.

[15] Alexandre Gonfalonieri, What is an AI Algorithm?, Medium, Apr. 21, 2019, Article.

- **Human interaction –** Human involvement is imperative throughout the life cycle of any AI application, from preparing the data and the algorithms to testing the output, retraining the model, and verifying results. As data is collected and prepared, human reviews are essential to curate the data as appropriate for the application. As algorithms sift through data and generate output (*e.g.*, classifications, outliers, and predictions), the next critical component is human review of the output for relevancy, accuracy, and usefulness. Business and technology stakeholders typically work together to analyze AI-based output and give appropriate feedback to the AI systems for refinement of the model. Absence of such human review and feedback may lead to irrelevant, incorrect, or inappropriate results from the AI systems, potentially creating inefficiencies, foregone opportunities, or new risks if actions are taken based on faulty results.

II. AI Applications in the Securities Industry

The use of AI-based applications is proliferating in the securities industry and transforming various functions within broker-dealers. Some large firms have established centers of excellence to review, share, and build expertise and create synergies related to the use of AI across their organizations. In addition, firms are exploring and incorporating AI tools built by financial technology startups and vendors.

In this section, we highlight three broad areas where broker-dealers are evaluating or using AI in the securities industry: communications with customers, investment processes, and operational functions. Specifically, with respect to communications with customers, the discussion explores how firms are using AI to enhance customer experience and outreach targeting. With respect to investment processes, the discussion explores how firms are using AI to assist with brokerage account management, portfolio management, and trading. Finally, with respect to operational functions, the discussion explores how firms are using AI to assist with compliance, risk management, and administrative tasks.

While this section of the paper provides a high-level informational summary of key AI applications that securities industry participants shared with FINRA staff, it is neither an exhaustive list of possible applications nor intended to be an endorsement of any particular use case. Although the use cases noted below may offer several potential benefits, they also involve potential challenges, costs, and regulatory implications. Each firm should conduct its own due diligence and legal analysis when exploring any AI application to determine its utility, impact on regulatory obligations, and potential risks, and set up appropriate measures to mitigate those risks. Furthermore, use of AI applications does not relieve firms of their obligations to comply with all applicable securities laws, rules, and regulations.

Communications with Customers

Customer Experience

The use of AI in applications to enhance customer experience has gained significant traction, not just in the securities industry but broadly within the financial services industry. AI-based customer service applications largely involve NLP- and ML-based tools that automate and customize customer communications.

- **Virtual assistants –** A virtual assistant is an AI application that interacts with humans using voice recognition and synthesized speech and is programmed to perform certain tasks. Several firms have already deployed, or are in the process of deploying, virtual assistants to provide digital customer service. Firms have noted that their virtual assistants provide responses to basic customer inquiries, such as those related to account balances, portfolio holdings, market data, address changes, and password resets. At a few firms, virtual assistants also offer the ability to accept and process trade orders within certain thresholds.

  Firms are deploying virtual assistants through different platforms. Several online firms have deployed conversational chatbots through their websites and mobile apps.[16] Some firms noted that they are experimenting with the development of virtual assistant applications available to customers through third-party platforms, such as Amazon's Alexa, Google's Home Assistant, and Apple's Siri. In addition, some firms are integrating AI-based interactive voice response (IVR) systems at their call centers to respond to basic inquiries from callers or collect sufficient information to facilitate triaging of calls to the appropriate human customer service agents.

  These online, device-based, and telephonic virtual assistants employ NLP (including speech-to-text/text-to-speech conversion, tone recognition, and text generation), ML, and sophisticated customer authentication tools, including the use of facial recognition, fingerprints, and voice biometrics. The applications are trained with large sets of historical and current data, including customer account information, trading history, and past inquiries, as well as market data, to formulate appropriate responses to incoming customer inquiries.

- **Email inquiries –** Some firms indicated that they are using AI-based applications to screen and classify incoming client emails based on key features, such as the sender's identity, the email's subject line, and an automated review of the email message itself. Such applications may also automatically respond to emails containing common or routine inquiries, while routing emails with complex inquiries to the appropriate personnel. Firms also noted using similar applications to process and triage internal inquiries (*e.g.*, those received by internal help desks), to provide automated responses where possible, and to route more complicated inquiries to relevant subject matter experts.

Outreach Targeting

A number of broker-dealers are exploring the use of AI to target outreach to customers or potential customers. Some firms are using AI tools to analyze their customers' investing behaviors, website and mobile app footprints, and past inquiries, and in turn, to proactively provide customized content to them. This could include curated educational information, news, and research reports on specific investment products or asset classes. This content could be delivered to customers by email or directly through the firm's website or mobile app. In addition, firms have also indicated AI tools are being explored to determine whether individuals would be interested in certain services based on their customer profile and browsing history within the firms' websites.

Industry participants noted that the exploration and deployment of AI-based applications for customer communications, particularly automated virtual assistants, have gained significant traction because they may help to enhance customer experience, create significant internal cost efficiencies for firms, and provide opportunities for greater engagement between the firm and customer. However, these applications may also pose certain challenges and potential risks, such as those associated with customer authentication, data privacy, cybersecurity, and recordkeeping.

[16] A chatbot is a computer program or a software that simulates conversations with humans in the form of text or voice messages.

Investment Processes

Brokerage Account Management

Brokerage account management functions at some large firms are starting to incorporate AI-based tools to provide individual brokers with customized, real-time customer information and better insights into customer preferences and trading behaviors. Registered representatives use such information to enhance customer relationships, to provide better service and recommendations, and to potentially minimize customer attrition. The following are some examples shared by firms for the use of AI in understanding customer needs and managing their brokerage accounts:

- **Holistic customer profiles –** Firms are starting to develop AI-based applications that create real-time, holistic customer profiles, which incorporate information from a broad range of sources, such as customer assets (held both at the firm and outside the firm), spending patterns, and debt balances obtained through data aggregation tools; updates on social media and other public websites; browsing history on the firm's website and mobile apps; and past communications (*e.g.*, from emails, chat messages, and meeting notes). All this information is analyzed using AI tools to provide the registered representative with a broader picture of customer needs, along with tailored suggestions of what investment products the customer may be interested in. Industry participants indicated that registered representatives use this information to augment their existing knowledge and expertise when making suggestions to their customers. Industry participants noted taking a cautious approach to employing AI tools that may offer investment advice and recommendations directly to retail customers, citing several legal, regulatory, and reputational concerns.[17]
- **Customized research –** Firms have also indicated growing use of AI tools to provide curated market research directly to customers to share relevant information on various investment opportunities. For example, as noted in the earlier section, AI-based tools may offer customers social media data and related sentiment analysis on investment products and asset classes.

While these AI tools offer the potential to customize investment suggestions for customers, firms should be cognizant of potential concerns and challenges related to data privacy, use of corrupt or misleading data, and adapting to each customer's unique circumstances.[18]

Portfolio Management and Trading

Broker-dealers are also exploring and using AI applications within their portfolio management and trading functions.

- **Portfolio management –** Within portfolio management, firms noted the use of AI applications to identify new patterns and predict potential price movements of specific products or asset classes. These applications tap into vast amounts of data available from internal and external sources, including from non-traditional sources like social media and satellite imagery, which is used as proxy data for economic activity to identify insights that may signal price movement. Some broker-dealers that are also investment advisors aim to incorporate these predictions into their investment strategies to generate alpha for the portfolio.

---

[17] In the U.S., digital investment platforms (commonly referred to as "robo-advisors") that offer investment advice directly to clients via online platforms, currently largely use rules-based models to develop those recommendations. See FINRA, Report on Digital Investment Advice, Mar. 2016, PDF.

Firms may wish to review their AI-based investment tools to determine whether related activity may be deemed as offering discretionary investment advice and therefore implicate the Investment Advisors Act of 1940.

[18] FINRA, Social Sentiment Investing Tools – Thing Twice Before Trading Based on Social Media, Apr. 2019, Article.

- **Trading –** Securities industry participants are also exploring AI tools to make their trading functions more efficient by maximizing speed and price performance. Examples include using ML for smart order routing, price optimization, best execution, and optimal allocations of block trades.

Firms should bear in mind that use of AI in portfolio management and trading functions may also pose some unique challenges, particularly where the trading and execution applications are designed to act autonomously. Circumstances not captured in model training – such as unusual market volatility, natural disasters, pandemics, or geopolitical changes – may create a situation where the AI model no longer produces reliable predictions, and this could trigger undesired trading behavior resulting in negative consequences. In addition, some industry participants have expressed concern that AI trading models across the industry may start to learn from each other, potentially leading to collusive activity, herd behavior, or unpredictable results.[19]

Operational Functions

Compliance and Risk Management

In conversations with FINRA staff, industry participants noted that they are spending significant time and resources in developing AI-based applications to enhance their compliance and risk management functions. This is consistent with FINRA's 2018 research on the use of regulatory technology (RegTech), where we observed that "market participants are increasingly looking to use RegTech tools to help them develop more effective, efficient, and risk-based compliance programs."[20] According to an April 2018 research study conducted by Chartis Research and IBM, which surveyed more than 100 relevant risk and technology professionals, 70% of respondents noted using AI in risk and compliance functions.[21]

Broker-dealers have to keep pace with complex and evolving domestic and international regulations, as well as a rapidly changing risk landscape (*e.g.*, cybersecurity, internal threats, and financial risks). At the same time, they now have access to vast amounts of data, inexpensive computing power, and innovative technologies that present opportunities for them to develop automated compliance and risk-management tools. Below are some examples that firms shared of how they are incorporating AI in their compliance and risk management tools.[22]

- **Surveillance and monitoring –** AI technology offers firms the ability to capture and surveil large amounts of structured and unstructured data in various forms (*e.g.*, text, speech, voice, image, and video) from both internal and external sources in order to identify patterns and anomalies. This enables firms to holistically surveil and monitor various functions across the enterprise, as well as monitor conduct across various individuals (*e.g.*, traders, registered representatives, employees, and customers), in a more efficient, effective, and risk-based manner. Market participants noted that these tools could significantly reduce the number of false positives, which in turn, free up compliance and supervisory staff time to conduct more thorough reviews of the remaining alerts, resulting in higher escalation rates. Firms indicate that these tools offer the ability to move beyond "traditional rule-based systems to a predictive, risk-based surveillance model that identifies and exploits patterns in data to inform decision-making."[23] For example, some firms noted the use of AI-based surveillance tools to monitor

[19] R. Jesse McWaters & Matthew Blake, Navigating Uncharted Waters: A Roadmap to Responsible Innovation with AI in Financial Services, World Economic Forum, Oct. 23, 2019, Article.

[20] FINRA, Technology Based Innovations for Regulatory Compliance ("RegTech") in the Securities Industry, Sept. 2018, PDF [hereinafter FINRA RegTech White Paper].

[21] IBM and Chartis Research, AI in RegTech: A Quiet Upheaval, 2018, Article.

[22] Please note that FINRA does not endorse or validate the use or effectiveness of any specific tools in fulfilling compliance obligations. FINRA encourages broker-dealers to conduct a comprehensive assessment of any compliance tools they wish to adopt to determine their benefits, implications and ability to meet their compliance needs.

communications with customers across various channels, such as emails, social media, and text messaging. Firms noted that these tools gave them the ability to move beyond a traditional lexicon-based review to a more risk-based review, such that they could decipher tone, slang, or code words, which may be indicative of potentially risky or non-compliant behavior.

- **Customer identification and financial crime monitoring –** AI-based tools are also being developed for customer identification (also referred to as "know-your-customer" (KYC)) and financial crime monitoring programs, for example, to detect potential money laundering, terrorist financing, bribery, tax evasion, insider trading, market manipulation, and other fraudulent or illegal activities.[24] Market participants noted that many traditional KYC and financial crime monitoring methods are cumbersome and not as effective as desired, often resulting in high rates of false positives. Consequently, firms have started incorporating AI technologies, such as ML, NLP, and biometrics, to make their programs more effective and risk based. Firms indicated that these tools enable them to identify and track customer activity with greater accuracy and efficiency, and to conduct more holistic and detailed analysis of customer transactions.[25]
- **Regulatory intelligence management –** Broker-dealers use a variety of regulatory intelligence management programs and processes to identify, interpret, and comply with new and changing rules and regulations across jurisdictions. While this has traditionally been a manual process, firms are now exploring the use of AI tools to digitize, review, and interpret new and existing regulatory intelligence (including rules, regulations, enforcement actions, and no-action letters) and to incorporate appropriate changes into their compliance programs. Some industry participants noted that automated regulatory intelligence management programs have the potential to increase overall compliance, while reducing both costs and time spent implementing regulatory change. According to a research report that explores the use of AI by financial institutions for risk and compliance functions, "[a]utomating the process of regulatory change management is something of a 'holy grail' in the use of AI."[26] Some regulators are also exploring and adopting the concept of "machine-readable" rulebooks, which could potentially enable firms to automate the process of identifying, categorizing, and mapping the rules to relevant regulatory obligations within their internal workflows.[27]
- **Liquidity and cash management –** In our discussions, some firms noted that they are employing ML applications to optimize their financial liquidity and cash management. Such applications analyze substantial historical data along with current market data to identify trends, note anomalies, and make predictions, for example, related to intra-day liquidity needs, peak liquidity demands, working capital requirements, and securities lending demand.

[23] FINRA RegTech White Paper.

[24] The Bank Secrecy Act of 1970 (BSA) requires all broker-dealers to, among other things, implement compliance programs to detect and prevent money laundering. In addition, FINRA Rule 3310 (Anti-Money Laundering Compliance Program) requires all broker-dealers to develop and maintain a written AML program to comply with the requirements of the BSA. FINRA Rule 2090 (Know Your Customer (KYC)) requires broker-dealers to "use reasonable diligence, in regard to the opening and maintenance of every account, to know (and retain) the essential facts concerning every customer and concerning the authority of each person acting on behalf of such customer."

[25] In December 2018, Treasury's FinCEN and federal banking agencies issued a joint statement, encouraging "banks and credit unions to take innovative approaches to combating money laundering, terrorist financing, and other illicit financial threats." Financial Crimes Enforcement Network (FinCEN), Joint Statement Encouraging Innovative Industry Approaches to AML Compliance, Dec. 3, 2018, Article.

[26] IBM and Chartis Research, Demystifying Artificial Intelligence in Risk and Compliance, Apr. 2018, PDF.

[27] FINRA, Special Notice on Financial Technology Innovation, July 30, 2019, PDF (seeking comments from the industry on, amongst other things, the development of a taxonomy-based machine readable rulebook.).
U.K. Financial Conduct Authority, Digital Regulatory Reporting, Jan. 11, 2017, PDF (stating that, in 2017, the U.K. Financial Conduct Authority would launch an initiative to make their rulebooks machine-readable).

- **Credit risk management –** Broker-dealer firms are also employing AI-based models to assess creditworthiness of their counterparties, which both speeds up the credit review process and allows the incorporation of non-traditional criteria (*e.g.*, information available through social media). However, some AI-based credit-scoring systems have faced criticism for being opaque and potentially biased and discriminatory. These models not only analyze traditional credit-evaluation criteria, such as current financial standing and historical credit history, but may also identify other demographic factors as deterministic criteria, which could lead to unfair and discriminatory credit scoring based on biases present in the underlying historical data. (Refer to Section III for additional discussion on the topic.)
- **Cybersecurity –** Cybersecurity continues to be a top challenge for the financial services industry. Perpetrators are continuously evolving and using sophisticated technology, including AI, to conduct their attacks. In addition, regulators are requiring financial institutions to develop comprehensive cybersecurity controls. In response, broker-dealers are starting to incorporate AI as an essential component of their cybersecurity programs. A recent research report noted that "sixty-nine percent of organizations believe AI will be necessary to respond to cyberattacks."[28] Incorporating AI into cybersecurity programs may allow firms to assist overwhelmed cybersecurity staff to predict potential attacks, detect threats in real-time, and respond to them faster and at lower costs. Use of AI in cybersecurity programs often begins within insider risk programs where normal behavior can be learned and then deviations or anomalies can be flagged as a risk and reviewed.

Administrative Functions

Another area where AI is being widely used by the securities industry is the automation of administrative functions. Firms are incorporating software augmented by AI technologies (*e.g.*, ML, NLP, and CV) to automate high-volume, less complex, repetitive, and manual tasks that traditionally involve significant human staff time. Firms indicated that automating such tasks with AI tools have the potential of high returns in terms of cost savings and efficiency increases. The following are select examples shared by some broker-dealers of AI use cases for administrative functions:

- **Automation of manual, paper-based processing –** Some firms have deployed AI tools to automate functions that involve manual review of documents, such as processing faxed trade orders, depositing physical checks, and searching, ranking, and retrieving documents. These applications incorporate CV and NLP to identify customers, review signatures, read orders, and digitize documents. Firms indicated that such applications not only increase productivity, but also accelerate important processes, such as trade and payment processing.
- **Document review and information extraction –** NLP enables firms to review significant volumes of documents (*e.g.*, legal contracts, custody documents, and loan agreements) at a fraction of the time it takes for human analysis. In addition to time and cost savings, some firms have indicated that these automated reviews have yielded results with higher accuracy. For example, some firms noted that they use NLP and unsupervised ML to review and classify legal contracts, or to look for certain clauses or key data points within a category of contracts. Another firm noted the use of AI to automate the review of prospectuses for key information, such as expense ratios.
- **Other applications –** Industry participants also indicated they are exploring other areas related to operations where AI tools help to improve the accuracy and efficiency of existing processes, such as reporting and invoice reconciliation.

[28] Capgemini Research Institute, Reinventing Cybersecurity with Artificial Intelligence: The New Frontier in Digital Security, July 11, 2019, PDF.

III. Key Challenges and Regulatory Considerations

AI-based applications offer several potential benefits to both investors and firms, many of which are highlighted in Section II. Potential benefits for investors include enhanced access to customized products and services, lower costs, access to a broader range of products, better customer service, and improved compliance efforts leading to safer markets. Potential benefits for firms include increased efficiency, increased productivity, improved risk management, enhanced customer relationships, and increased revenue opportunities.

However, use of AI also raises several concerns that may be wide-ranging across various industries as well as some specific to the securities industry. Over the past few years, there have been numerous incidents reported about AI applications that may have been fraudulent, nefarious, discriminatory, or unfair, highlighting the issue of ethics in AI applications. As such, several organizations have established initiatives or developed principles to promote the ethical use of AI.[29]

AI-based applications present some particular challenges that securities market participants may wish to consider as they explore and adopt related technology tools. Specifically, where applicable, factors for market participants to consider when seeking to adopt AI-based applications include model risk management, data governance, customer privacy, and supervisory control systems. Other factors for potential consideration include cybersecurity, outsourcing/vendor management, books and records, and workforce structure. This section provides a brief discussion of each of these factors and highlights certain related regulatory considerations.[30]

While this section highlights certain key thematic areas, it is not meant to be an exhaustive list of all factors or regulatory considerations associated with adopting AI-based applications. Broker-dealers should conduct their own assessments of the implications of AI tools, based on their business models and related use cases.

Model Risk Management

Firms that employ AI-based applications may benefit from reviewing and updating their model risk management frameworks to address the new and unique challenges AI models may pose. These challenges may include those related to model explainability, data integrity, and customer privacy. Model risk management becomes even more critical for ML models due to their dynamic, self-learning nature.

A comprehensive model risk management program typically includes areas such as model development, validation, deployment, ongoing testing, and monitoring. Where applicable, the following are potential areas for firms to consider as they update their model risk management programs to reflect the use of AI models:

- Update model validation processes to account for complexities of an ML model.[31] This includes reviewing the input data (*e.g.*, review for potential bias), the algorithms (*e.g.*, review for errors), any parameters (*e.g.*, verify risk thresholds), and the output (*e.g.*, determine explainability of the output).

[29] For instance, in Apr. 2019, the European Commission published a set of non-binding "Ethics guidelines for trustworthy AI", prepared by the Commission's High-Level Expert Group on AI. European Commission, Ethics Guidelines for Trustworthy AI, April 8, 2019, Article.

[30] Supra note 1. While the paper highlights certain regulatory and implementation areas that broker-dealers may wish to consider as they adopt AI, the paper does not cover all applicable regulatory requirements or considerations. FINRA encourages firms to conduct a comprehensive review of all applicable securities laws, rules, and regulations to determine potential implications of implementing AI-based applications.

[31] Model validation refers to "the set of processes and activities intended to verify that models are performing as expected, in line with their design objectives and business uses. Effective validation helps to ensure that models are sound, identifying potential limitations and assumptions and assessing their possible impact." Board of Governors of the Federal Reserve System, Supervisory Letter (SR 11-7) on Guidance on Model Risk Management, Apr. 4, 2011, Article.

- Conduct upfront as well as ongoing testing, including tests that experiment with different and stressed scenarios (*e.g.*, unprecedented market conditions) and new datasets.
- Employ current and new models in parallel and retire current models only after the new ones are thoroughly validated.
- Maintain a detailed inventory of all AI models, along with any assigned risk ratings such that the models can be appropriately monitored and managed based on their risk levels.
- Develop model performance benchmarks (*e.g.*, number of false negatives) and an ongoing monitoring and reporting process to ensure that the models perform as intended, particularly when the models involved are self-training and evolve over time.

Model Explainability

Many ML models allow for some level of explainability with respect to the underlying assumptions and factors considered in making a prediction. Some ML models, however, are described as "black boxes"[32] because it may be difficult or impossible to explain how the model works (*i.e.*, how its predictions or outcome are generated).[33]

While firms indicated that operational deployment of black box models in the near term within the securities industry was unlikely, they also noted that some cutting-edge applications of AI had presented explainability challenges.

An appropriate level of explainability may be particularly important in AI applications that have autonomous decision-making features (*e.g.*, deep learning-based AI applications that trigger automated investment decision approvals). Against this backdrop, firms noted that their compliance, audit, and risk personnel would generally seek to understand the AI-models to ensure that they conform to regulatory and legal requirements, as well as the firms' policies, procedures, and risk appetites before deployment.

FINRA Rule 3110 (Supervision) requires firms to establish and maintain a system to supervise the activities of its associated persons that is reasonably designed to achieve compliance with the applicable securities laws and regulations and FINRA rules. This rule applies to all activities of a firm's associated persons and its businesses, regardless of the use of technology. As such, in supervising activities related to AI applications, firms have indicated that they seek to understand how those applications function, how their outputs are derived, and whether actions taken pursuant to those outputs are in line with the firm's legal and compliance requirements.

The following are some potential areas for firms to consider, as applicable, when establishing policies and procedures that address concerns related to explainability.

- Incorporating explainability as a key consideration in the model risk management process for AI-based applications. This may involve requiring application developers and users to provide a written summary of the key input factors and the rationale attributed to the outputs. The

[32] In its February white paper, the European Commission also noted the 'black box effect' as one of its top concerns: "opacity, complexity, unpredictability, and partially autonomous behavior may make it hard to verify compliance with, and may hamper the effective enforcement of, rules of existing EU law." European Commission, White Paper on Artificial Intelligence – A European Approach to Excellence and Trust, Feb. 19, 2020, PDF.

[33] In simple machine learning models (e.g., models that use traditional statistical methods, such as logistic regression or decision trees), one can follow the logic used by the models and the factors that contribute to the final outcome. However, more complex models (e.g., deep learning models) involve multiple layers and a dynamic, iterative learning process, where the internal learnings are opaque, making it difficult to identify the specific factors and their interrelationships that lead to the final outcome. Despite the challenges, these more complex AI models continue to garner interest from the industry because they offer the potential to be more powerful in identifying patterns and making more precise predictions relative to simpler models.

models can then be tested independently by the model validation teams or by external parties. Some firms noted that they test certain ML models using techniques that involve isolating specific data variables or features in the model to determine their impact on the output. For instance, if eliminating an important feature does not significantly change the output, it may indicate that the model is not appropriately incorporating that feature in the decision-making. Some firms also noted introducing new datasets during the model validation process to ensure that the model operates consistently. However, in considering factors around explainability, it is important to guard against ex-post rationalization of ML models based on correlations that may not link to any underlined causality.[34]

- Building a layer of human review of the model outputs, where applicable, to ensure that the results are in line with business goals, as well as firms' internal policies, procedures, and risk appetite. In our discussions with industry participants, the vast majority noted that their ML-based applications do not involve autonomous action but are instead used to aid human decision-making.
- Establishing appropriate thresholds and guardrails, where ML models trigger actions autonomously. For example, some firms exploring ML for trading applications indicated that they have risk-based limitations built into those applications, such as amount or threshold limits for trade orders.

There are several efforts underway in the financial services industry, as well as more broadly, to develop tools that can provide transparency and explainability for AI models. One such notable effort is the Explainable AI (XAI) program undertaken by the Defense Advanced Research Project Agency (DARPA) to "[p]roduce more explainable models, while maintaining high levels of learning performance (prediction accuracy); and [e]nable human users to understand, appropriately trust, and effectively manage the emerging generation of artificially intelligent partners."[35]

Data Governance

Data is the lifeblood of any AI application. AI applications benefit from large amounts of data to train and retrain models, conduct comprehensive analyses, identify patterns, and make predictions. Accordingly, the quality of the underlying dataset is of paramount importance in any AI application.

Industry participants noted that one of the most critical steps in building an AI application is to obtain and build the underlying database, such that it is sufficiently large, valid, and current. Depending on the use case, data scarcity may limit the model's analysis and outcomes, and could produce results that may be narrow and irrelevant. On the other hand, incorporating data from many different sources may introduce newer risks if the data is not tested and validated, particularly if new data points fall outside of the dataset used to train the model. In addition, continuous provision of new data, both in terms of raw and feedback data, may aid in the ongoing training of the model.

[34] Several research studies have noted that data mining can lead to incorrect or misleading results because of the identification of spurious correlations. See, for instance, Hou, Kewei and Xue, Chen and Zhang, Lu, Replicating Anomalies, Working Paper No. 2017-03-010 presented at the 28th Annual Conference on Financial Economics and Accounting, Fisher College of Business, June 12, 2017, Article.

[35] Matt Turek, Explainable Artificial Intelligence (XAI), Defense Advanced Research Projects Agency, Article.

Data Bias

Data integrity is a key area of consideration for ML models. An important factor for maintaining data integrity centers on considering issues of data bias.[36] One component of data bias involves data sets that may not include relevant information or include a skewed subset of information, resulting in potential distortions in the output produced by ML models. To limit the potential for this type of data bias it is helpful to have data that is complete, current, and accurate. Another component of data bias involves issues around demographic biases.[37] Such biases may relate to individuals or institutions and may exist in data due to historical practices and could foster discriminatory outcomes if not appropriately addressed.[38]

Data bias, in general, may also be inadvertently introduced during the data preparation process, as data scientists determine which data fields and related features to incorporate in the ML model. Any biases in the underlying data may propagate through the ML model and may lead to inappropriate decision-making.

FINRA Rule 2010 requires firms, in the conduct of their business, to observe high standards of commercial honor and just and equitable principles of trade. These general requirements apply to activity engaged by the firm including, where applicable, those resulting from AI applications.

When reviewing and modifying data governance policies and procedures to address potential data-related risks that may emerge in AI applications, some areas for consideration, where applicable, include:

- **Data review for potential bias –** When building an AI application, it is important to review the underlying dataset for any potential built-in biases. Some firms undertake steps during the testing process to review for potential biases. For example, some firms adjust or eliminate certain data features of the AI model to see how the changes impact the model output. Using such data filters may provide indications of potential biases in the model output based on those features. Testing the models using proxies instead of demographic data is another technique that may reveal potential biases. Firms may also involve multiple participants representing different functions to review the dataset as well as to test the outputs of the models. Recent reports have noted that introducing diversity in the staffing teams that build and test AI systems may provide wider perspectives and enhanced reviews for potential bias in the data. "A more diverse AI community would be better equipped to anticipate, review and spot bias and engage communities affected." Furthermore, providing training on this issue to all individuals that are involved in the development and testing process will likely make them more cognizant of the issue. Firms may also consider using open-source tools created by large technology companies to assist in identifying unwanted bias in data.[41]
- **Data source verification –** AI models often incorporate data from both internal and external sources. As firms compile data, many indicated that they regularly review the legitimacy and authoritativeness of data sources. This is particularly important where the data is obtained from external sources (*e.g.*, open-source data libraries, market data providers, data aggregators, and social media channels). When sourcing data from open source platforms or social media, firms benefit from incorporating appropriate verification steps, particularly given the proliferation of deep fakes[42] and social media manipulation.[43]

[36] Digital Challenges: Overcoming Barriers to AI Adoption, May 28, 2019, Article (in which EY and Massachusetts Institute of Technology conducted a survey at the 2019 EmTech Digital Conference, and found that "42% cite a lack of quality, unbiased data as the greatest barrier to AI adoption.").

[37] For example, Federal Study Confirms Racial Bias of Man Facial Recognition Systems, Casts Doubt on Their Expanding Use, The Washington Post, Dec. 19, 2019, Article.

- **Data integration –** As firms tap into various data sources to power their AI applications, they seek to obtain and integrate the data effectively into their systems so that it can be leveraged across their organizations. While traditionally, data may have resided in silos in different parts of their organizations, firms are now creating central data lakes to ensure consistency in data usage, to maintain appropriate entitlement and access levels, and to create synergies in data usage.
- **Data security –** Another key consideration is the security of the data that is made available to various stakeholders, both internal and external, in order to develop, test, and use AI applications. It is critical that firms develop, maintain, and test appropriate entitlement, authentication, and access control procedures, as well as use encryption techniques for sensitive data. As discussed in the following section, ensuring customer data privacy is a key objective in establishing data security measures.
- **Data quality benchmarks and metrics –** As part of a comprehensive data governance strategy, firms may also wish to consider developing and monitoring benchmarks and metrics to measure and assess the effectiveness of their data governance programs.

Customer Privacy

AI applications used in the securities industry may involve the collection, analysis, and sharing of sensitive customer data, as well as ongoing monitoring of customer behavior. For example, AI-based customer service tools may involve collection and use of personally identifiable information (PII) and biometrics. Similarly, certain customer focused AI applications monitor information, such as customer website or app usage, geospatial location, and social media activity. Some tools also involve recording written, voice, or video communications with customers. While AI tools based on these types of information may offer firms insights into customer behavior and preferences, they also may pose concerns related to customer privacy if the information is not appropriately safeguarded. Broker-dealers benefit from considering the applicability of relevant customer privacy rules when developing and using such applications, both with respect to the data that is used in AI models and the information that is made available by their outputs.

Protection of financial and personal customer information is a key responsibility and obligation of FINRA member firms. As required by SEC Regulation S-P (Privacy of Consumer Financial Information and Safeguarding of Personal Information), broker-dealers must have written policies and procedures in place to address the protection of customer information and records.[44] In addition, as detailed in *NASD Notice to Members 05-49* (Safeguarding Confidential Customer Information), firms are required to maintain policies and procedures that address the protection of customer

[38] Over the last few years, there have been several reports and notable incidents across various industries of ML models producing allegedly biased and discriminatory results. For example, the European Commission, in its February 2020 white paper, noted its concern about bias and discrimination: "The use of AI can affect the values on which the EU is founded and lead to breaches of fundamental rights, including the rights to freedom of expression, freedom of assembly, human dignity, nondiscrimination based on sex, racial or ethnic origin, religion or belief, disability, age or sexual orientation, as applicable in certain domains, protection of personal data and private life, or the right to an effective judicial remedy and a fair trial, as well as consumer protection. These risks might result from flaws in the overall design of AI systems… or from the use of data without correcting possible bias…" European Commission, White Paper on Artificial Intelligence – A European Approach to Excellence and Trust, Feb. 19, 2020, PDF.

[39] Rubin Nunn, Workforce Diversity Can Help Banks Mitigate AI Bias, American Banker, May 30, 2018, Article. See also, Jake Silberg & James Manyika, Tackling Bias in Artificial Intelligence (and in Humans), June 2019, Article.

[40] James Manyika et al, What Do We Do About the Biases in AI?, Harvard Business Review, Oct. 25, 2019, Article.

[41] Some examples of tools that reportedly assist in identifying potential bias in data include IBM's AI Fairness 360 and Google's Responsible AI Practices.

[42] Deep fakes refer to a form of synthetic or manipulated communication in which an existing video, image, or audio clip is replaced or superimposed with someone else's likeness in order to create false impressions or communications.

[43] FINRA, Social Sentiment Investing Tools – Thing Twice Before Trading Based on Social Media, Apr. 2019, Article.

[44] U.S. Securities and Exchange Commission, Final Rule: Privacy of Consumer Financial Information (Regulation S-P), Article.

information and records, and ensure that their policies and procedures adequately reflect changes in technology. Firms also should provide initial and annual privacy notices to customers describing information sharing policies and informing customers of their rights. Additionally, SEC Regulation S-ID (the Red Flags Rule) requires broker-dealer firms that offer or maintain covered accounts to develop and implement written "Identity Theft Prevention Programs." Moreover, numerous international, federal, and state regulations and statutes set forth specific rules and requirements related to customer data privacy. Firms should assess the applicability of these laws as they build their AI applications and any underlying infrastructures.

Firms also should update their written policies and procedures with respect to customer data privacy, to reflect any changes in what customer data and information is being collected in association with AI applications, and how that data is collected, used, and shared. In this regard, below are some questions for firms to consider.

- Have appropriate consents been obtained from customers to track, collect, and monitor their information, including information obtained directly from customers (*e.g.*, PII and biometrics) as well as other sources (*e.g.*, website usage, social media platforms, or third-party vendors)?
- Has the applicable data been authorized for each relevant use case related to an AI application?
- Have user entitlements and access procedures been updated as new, shared databases or centralized data lakes are created?
- Has sensitive data been obfuscated[45] (as needed) and does the data continue to remain protected as it is applied across different AI applications?
- Has the data governance framework been appropriately updated to reflect any changes related to customer data privacy policies and procedures?

Supervisory Control Systems

FINRA rules require firms to establish and maintain reasonable supervisory policies and procedures related to supervisory control systems in accordance with applicable rules (*e.g.*, FINRA Rules 3110 and 3120). This includes having reasonable procedures and control systems in place for supervision and governance of AI-based tools and systems across applicable functions of a broker-dealer.

As discussed earlier in this report, use of AI-based applications may pose unique and complex challenges, such as those related to model explainability and bias. As noted in FINRA's 2020 Risk Monitoring and Exam Priorities Letter, "Firms' increasing reliance on technology for many aspects of their customer-facing activities, trading, operations, back-office, and compliance programs creates a variety of potential benefits, but also exposes firms to technology-related compliance and other risks."

As broker-dealer firms employ AI-based tools and services across the firm, they should update and test related supervisory procedures and reflect those updates in their written supervisory procedures (WSPs). In addition to the topics discussed earlier in this report, some areas for consideration, where applicable, when adopting AI applications are noted below.[46]

[45] Obfuscation may be accomplished with encryption, tokenization, or anonymization techniques.

[46] These are some of many possible areas that broker-dealers may wish to consider as they explore adjusting their supervisory processes. This does not express any legal position, does not create any new requirements or suggest any change in any existing regulatory obligations, nor does it provide relief from any regulatory obligations. It is not intended to cover all applicable regulatory requirements or considerations. FINRA encourages firms to conduct a comprehensive review of all applicable securities laws, rules, and regulations to determine potential implications of implementing AI-based tools and systems.

- **Establish a cross-functional technology governance structure –** As previously stated by FINRA, firms may find it beneficial to establish a cross-disciplinary technology governance group to oversee the development, testing, and implementation of AI-based applications.[47] Such a group could include representation from different functions across the organization, including business, technology, information security, compliance, legal, and risk management. FINRA has previously stated in the trading context that: "[A]s the use of algorithmic strategies has increased, the potential of such strategies to adversely impact market and firm stability has likewise grown. When assessing the risk that the use of algorithmic strategies creates, firms should undertake a holistic review of their trading activity and consider implementing a cross-disciplinary committee to assess and react to the evolving risks associated with algorithmic strategies."[48]
- **Conduct extensive testing of applications –** Testing new tools and applications across various stages of their lifecycle can help identify potential concerns in a timely manner and limit potential issues. This could involve extensive testing of the applications by various user groups, and by using new data sets and new scenarios in the testing process. In addition, this could also include maintaining existing parallel systems as firms test new ones.
- **Establish fallback plans –** Establishing back-up plans in the event an AI-based application fails (*e.g.*, due to a technical failure or an unexpected disruption) can help ensure that the relevant function is carried on through an alternative process. FINRA Rule 4370 (Business Continuity Plans and Emergency Contact Information) requires firms to create and maintain a written business continuity plan with procedures that are reasonably designed to enable firms to meet their obligations to customers, counterparties, and other broker-dealers during an emergency or significant business disruption.
- **Verify personnel registrations –** The skillsets of securities industry personnel are evolving rapidly to keep pace with the adoption of emerging technologies. Technical and operational roles are starting to blend, as information technologists and data scientists are playing key roles in operational functions like trading and portfolio management. Firms may need to evaluate the roles of these personnel to ensure that they have the appropriate FINRA licenses and registrations. For instance, as stated in *Regulatory Notice 16-21*, FINRA requires registration of associated persons involved in the design, development, or significant modification of algorithmic trading strategies. Furthermore, FINRA Rule 1220(b)(3) and *FINRA Regulatory Notice 11-33* state that certain firm personnel engaged in "back office" covered functions must qualify and register as Operations Professionals.

AI technology has the potential to disrupt and transform supervisory functions within a broker-dealer. Firms may benefit from conducting an overall assessment of the functions and activities that are employing AI-based tools, and updating their supervisory procedures accordingly. The following are examples of some areas that firms may wish to review.

- **Trading applications:** FINRA has previously stated in *Regulatory Notice 15-09*, "In addition to specific requirements imposed on trading activity, firms have a fundamental obligation generally to supervise their trading activity to ensure that the activity does not violate any applicable FINRA rule, provision of the federal securities laws or any rule thereunder." As firms adopt AI algorithms and strategies in their trading functions, they benefit from reviewing and testing their supervisory controls to ensure that there is continued compliance with applicable rules and regulations, including but not limited to FINRA Rules 5210 (Publication of Transactions and Quotations), 6140 (Other Trading Practices) and 2010 (Standards of Commercial Honor and Principles of Trade), SEC Market Access Rule, and SEC Regulation NMS, Regulation SHO and Regulation ATS.

[47] FINRA RegTech White Paper.

[48] FINRA, Regulatory Notice 15-09 on Effective Supervision and Control Practices for Firms Engaging in Algorithmic Trading Strategies, Mar. 2015, Article.

- **Funding and liquidity risk management:** As firms employ AI applications across functions like liquidity and cash management, portfolio management, and trading, they may wish to consider reviewing their supervisory procedures to ensure that the applications and the underlying models incorporate appropriate risk thresholds and relevant regulatory requirements, and do not create an environment of excessive risk-taking. This is particularly relevant where AI tools are used for liquidity and cash management, cases in which the models may generate aggressive recommendations for liquidity and leverage or may lead to unsound recommendations in unforeseen or stressed situations. As firms review their supervisory procedures, some factors to review include adequacy of existing controls, monitoring tools, and reporting tools to manage such risks. As previously stated in *Regulatory Notice 15-33*, "As part of a firm's obligation to supervise the businesses in which it engages, FINRA expects each firm to regularly assess its funding and liquidity risk management practices so that it can continue to operate under adverse circumstances, whether these result from an idiosyncratic or a systemic event." Further, as stated in *Regulatory Notice 10-57*, "FINRA expects broker-dealers affiliated with holding companies to undertake these efforts at the broker-dealer level, in addition to their planning at the holding-company level."[49]
- **Investment advice tools:** Market participants are exploring the use of AI tools that generate client risk profiles and potential investment recommendations. These tools may aid in developing a new investment strategy, rebalancing portfolios, suggesting specific products or asset classes, or offering tax-minimization strategies. Market participants may benefit from considering how SEC Regulation Best Interest (BI) and FINRA Rule 2111 (Suitability) would apply in these contexts. In addition, as noted in FINRA's Report on Digital Investment Advice, firms should ensure "sound governance and supervision, including effective means of overseeing suitability of recommendations, conflicts of interest, customer risk profiles and portfolio rebalancing."

Additional Considerations

- **Cybersecurity –** As noted in Section II of this report, cybersecurity continues to be a key threat for the financial services industry. While AI technology empowers the industry to identify potential security threats and attacks in real-time, use of related applications may also pose new vulnerabilities and threats. For instance, AI-based applications that pull in data from multiple sources may expose the firm to new security risks. In addition, customer-facing tools offered by firms on third-party platforms (*e.g.*, virtual assistants offered on third-party consumer devices) could also pose security risks, such as those introduced through vulnerabilities of those third-party platforms or through inadequate customer authentication procedures. Firms would benefit from incorporating cybersecurity as a critical component of the evaluation, development, and testing process of any AI-based application. For additional resources on this topic, including applicable rules, guidance, and FINRA's report on Cybersecurity Practices, refer to FINRA's webpage on cybersecurity.

[49] FINRA, Regulatory Notice 10-57 on Funding and Liquidity Risk Management Practices, Notice; FINRA, Regulatory Notice 15-33 on Guidance on Liquidity Risk Management Practices, Notice.

- **Outsourcing and vendor management –** As firms look to take advantage of the benefits offered by AI-based tools, many are choosing to outsource specific functions or purchase turnkey applications from vendors. Some vendors are developing niche products that leverage AI for specific activities (*e.g.*, financial crime monitoring and trade surveillance). Use of such vendor tools can be appealing to both small and large firms that seek to implement AI-based technology with low upfront capital investment and faster implementation time.[50] Firms are reminded that outsourcing an activity or function to a third-party does not relieve them of their ultimate responsibility for compliance with all applicable securities laws and regulations and FINRA rules. As such, firms should review and update their WSPs to ensure that they appropriately address outsourcing arrangements (see, *e.g.*, *Notice to Members 05-48 (Outsourcing)*) and to ensure that the security of the third-party meets or exceeds that expected by the firms. Firms may also wish to consider introducing language in contracts with third-party vendors that includes, but is not limited to, requiring vendors to notify firms in the event of a security breach and giving firms the right to audit, including the ability to review third-party System and Organization Controls (SOC) reports.
- **Books and records –** The use of AI applications may lead to the creation of new records. Firms should review the use of their AI tools and systems to ensure compliance with recordkeeping obligations, such as those associated with Exchange Act Rules 17a-3 and 17a-4 and FINRA Rule 4510 (Books and Records Requirements). For example, the use of AI tools with respect to chatbots and virtual assistants may create novel issues in the context of compliance with applicable recordkeeping requirements.
- **Workforce structure –** Adoption of AI tools and services could potentially impact a firm's workforce in multiple ways. AI-based applications may conduct certain tasks previously performed manually in a more effective manner and in a fraction of the time. This could result in a reduced number of jobs in certain skillsets and an increased number in others. This transition may be challenging for many reasons, including that there is a significant shortage of individuals with AI-related skills in the industry.[51] Firms may also face challenges with adjusting the culture of their technology functions as they shift away from a traditional waterfall methodology to a more agile process. Accordingly, as firms seek to adopt AI-based technology tools, they may want to consider reviewing the potential impact on their workforce, and take the appropriate steps to, for example, review any new staffing needs or training required for the use of new applications.

[50] However, some market participants have raised concerns that use of a limited number of vendors to develop AI tools for an industry, if not managed appropriately, may create overreliance and concentrate risks related to errors or malfunctioning of AI systems.

[51] Scott Likens, How Artificial Intelligence Is Already Disrupting Financial Services, Barron's, May 16, 2019, Article (stating that "[A]lmost a third of the financial executives in the AI survey are worried that they won't be able to meet the demand for AI skills over the next five years.").

IV. REQUEST FOR COMMENTS

FINRA encourages comments on this paper, including areas where guidance or modifications to FINRA rules may be desired to support adoption of AI applications while maintaining investor protection and market integrity. Comments are requested by August 31, 2020.

Member firms and other interested parties can submit their comments using the following methods:

- Emailing comments to pubcom@finra.org; or
- Mailing comments in hard copy to: Marcia E. Asquith, Office of the Corporate Secretary, FINRA, 1735 K Street NW, Washington, DC 20006-1506.

To help FINRA process comments more efficiently, persons should use only one method to comment on the proposal.

Important Notes: All comments received in response to this paper will be made available to the public on the FINRA website. In general, FINRA will post comments as they are received.

Direct inquiries regarding this paper to Haimera Workie, Senior Director and Head of Office of Financial Innovation, at (202) 728-8097; or Kavita Jain, Director, Office of Financial Innovation, at (202) 728-8128.