Securities Regulation Daily

Wolters Kluwer

July 31, 2020

# Behind the Curtain: The Role of Explainable AI in Securities Markets

*By Collin Starkweather* and Izzy Nelken***

The renowned statistician George Box famously observed that "all models are wrong, but some are useful." His statement implicitly acknowledges that scientific models are not just useful for their predictive abilities, but also for their ability to explain the world around us.[1]

Whether we care about the accuracy of a statistical model or want to understand why it gave a prediction is an unresolved tension whose consideration is timely and particularly germane to those adopting artificial intelligence and machine learning (AI and ML) techniques in financial markets.

To date, many AI and ML algorithms are considered "black-boxes" that may provide relatively accurate predictions, but at the cost of comprehensibility.[2] As concerns about the nature of model results (as opposed to merely the predictive accuracy) have gained traction, a new area of research has emerged to address these concerns.

More generally, the use of models to explain, rather than simply predict, is an ongoing and rapidly developing area of research in AI and ML,[3] where many models are characterized by a size and complexity (and corresponding opacity) that dwarfs their classical statistical brethren.[4]

In a recent overview of AI in the securities industry, FINRA observed that AI and ML technology is "transforming the financial services industry across the globe."[5] This article discusses how recent advances in explainability and interpretability inform the use AI and ML in securities markets. Though it has not been the focus of attention in the popular or financial industry press, it is becoming increasingly relevant in securities markets, where AI and ML techniques have seen rapid adoption across a wide range of market participants and operational roles.

Although the terms "explainability" and "interpretability" have been evolving in the vernacular and are often used relatively interchangeably, for ease of explication, we'll use the former term throughout.[6] We'll define "explainability" as the ability of an algorithm to explain or to present the reasons a given output was produced in terms understandable to a human.[7]

---

* Collin Starkweather, Ph.D., is the founder of Starkweather Economics LLC, and is an economics and technology consultant specializing in complex litigation and other matters related to artificial intelligence and machine learning, technology, finance, and competition economics.

** Izzy Nelken, Ph.D., is the founder of Super Computer Consulting, Inc. He is a member of the CBOE Product Development Committee and has served as a consultant or expert witness on several matters involving allegations of market manipulation and other securities matters.

## The Need for Explainability

AI and ML techniques have been making rapid inroads into a broad range of industries ranging from health care to logistics to banking and financial services. Their efficacy at tasks such as image recognition and manipulation, natural language processing, and autonomous vehicle navigation is by now well known,[8] and their broad application has also focused popular attention on some of the more opaque techniques in the AI and ML toolkit.

Many AI and ML techniques have the character of a "black box," in which the connection between model inputs and outputs may not be clear, even to the modeler.[9] For example, deep learning techniques, which rely on the application of deeply-layered neural networks, are commonly used for image recognition, among other things.[10] A deep learning application may enable a human face to be identified in a photo, but be unable to offer any insight as to how it recognized the face.[11]

There are a variety of use cases where there may be important considerations other than predictive accuracy. The ability to explain or interpret model results may be of considerable value to the model design process. It is also particularly important in settings where trust or fairness are integral to the process. It is one thing to be targeted for advertising or have friends identified in photos uploaded to social media by an opaque algorithm and quite another to be denied a loan, sentenced to prison,[12] or to receive sub-standard healthcare.[13]

The remainder of this article will review regulatory and market incentives driving development of explainable AI and ML techniques and provide a brief overview of recent developments in those techniques. Then we'll discuss increasing adoption of AI and ML techniques by securities market participants. Finally, we'll consider the ways that recent advancements will inform the nature of AI and ML adoption in securities markets, where it has implications ranging from trading model design processes to risk management to regulatory and compliance functions.

## Incentivizing Explainable AI

In recent years, regulators have become increasingly focused on compliance issues that arise in the context of AI and ML models and legislators have begun introducing explicit requirements for explainability into various jurisdictions' legal frameworks.

For example, banking regulators in the US have focused considerable attention on potential bias against protected classes in business practices such as consumer lending. The Equal Credit Opportunity Act (ECOA) requires creditors to provide consumers specific reasons for an adverse action such as a denial of credit.[14] The US Consumer Financial Protection Bureau (CFPB) recently commented that "AI may create or amplify risks, including risks of unlawful discrimination, lack of transparency, and privacy concerns" and that it is "particularly interested in exploring … [t]he accuracy of explainability methods, particularly as applied to deep learning and other complex ensemble models" and "[h]ow to convey the principal reasons [for an adverse action] in a manner that accurately reflects the factors used in the model and is understandable to consumers."[15]

 

California, whose legal framework on data privacy and related issues have informed legislative efforts at the federal level, recently put into effect the California Consumer Privacy Act (CCPA), which has been called "the most thorough privacy regulation in the US."[16] Mike Leone, the Senior Analyst at the Enterprise Strategy Group, commented to Forbes that "it will force those leveraging AI to prioritize explainability as a feature of their chosen AI platform, where insights derived from AI must be explained to a point where they can be understood by a human."[17]

In Europe, the EU's General Data Protection Regulation (GDPR) enacted a so-called "right to explanation" that provides a broadly encompassing prescription that organizations be able to explain their algorithmic decisions that "significantly affect" users.[18]

A more broadly relevant consideration may be the model design process itself. The ability to explain links between model inputs and outputs can provide model designers with important context by which to refine and improve their models, and when unusual or unexpected behaviors arise, may facilitate the analysis of those behaviors.

As one researcher observed, where gains from model design improvements exceed the gains in performance from the use of a "black box" model, "more interpretability leads to better overall accuracy."[19] This notion has been memorialized in a variety of widely-used data mining and knowledge discovery processes.[20]

## Recent Developments in Explainability

In response to the increasing focus on explainability by public and private interests, rapid improvements have been made in the last few years in the development of explainable AI and ML techniques. With many of these advancements, the oft-cited tradeoff between explainability and predictive accuracy[21] has been substantially diminished or effectively eliminated. As senior data scientists at H2O.ai, provider of a leading open source AI and ML platform, commented, "you can now have your accuracy and interpretability cake…and eat it too."[22]

There are three primary approaches to explainability, each with advantages and disadvantages. When considering relative advantages and disadvantages, it is important to bear in mind that explainability is not binary in nature, but instead exists on a continuum. And similar to statistical models, explainable models may provide you with important insights into various aspects of the underlying system, but they can't explain everything.

The first and most straight-forward approach to explainability involves simplifying models sufficiently to facilitate ready interpretation of results. While there is some appeal to this approach in its simplicity, an obvious drawback is the loss of performance that often accompanies simplification of the model. Although advances using other approaches have produced claims that you can "have your cake and eat it too," when modeling a system with a high degree of complexity, the tradeoff between performance and explainability is an important consideration with this approach.[23]

 

A second approach involves post hoc analysis of "black box" models, in which parameters may be adjusted or other means used to determine how changes to inputs might impact changes to outputs. This approach relies on a well-established theoretical foundation,[24] and because it retains the use of "black box" models, there is no loss of model efficacy.

However, there are limitations to this approach. The insights provided by this approach may be localized to the domain of the immediate analysis,[25] may lack fidelity relative to the model being explained, or may offer explanations that are insufficiently detailed.[26] As a result of these shortcomings, some prominent voices have recommended against the use of this approach, leading one scholar to go so far as to state that "explainable black boxes should be avoided in high-stakes decisions" in domains such as healthcare and criminal justice.[27]

A third approach which has acquired more mindshare recently involves designing ML techniques that are explainable from the ground up. This approach has gained traction in recent years as demand for explainability has led to active interest in the research community and a proliferation of new models including explainable neural networks (XNNs), explainable boosting machines (XBMs) , scalable Bayesian rule lists, and super-sparse linear integer models (SLIMs), among others.[28]

While the features baked into these models offer the modeler insights that their more opaque cousins do not,[29] there have been concerns among potential adopters that the explainability comes at a price in terms of model efficacy. Recent research indicates that for many of these models, and in a variety of settings, these concerns may be overblown or even unjustified.[30]

## Roles for Explainable AI in Securities Markets

In this section, we will focus on AI and ML adoption and usage in front office trading, portfolio management, and related operations and in mid office risk management and compliance operations, though AI- and ML-based applications are seeing deployment with securities markets participants in a wide range of operational roles in the front, mid, and back offices including customer engagement, brokerage account management, administrative functions, and cybersecurity.[31]

### AI and ML in the Front Office

AI and ML techniques have seen rapid adoption in front office operations in recent years. As far back as early 2016, the BBC was reporting that "[h]edge funds are increasingly turning to artificial intelligence in order to spot trends to try to make money for their customers"[32] and prominent asset managers such as BlackRock have in recent years been reported to be refocusing portfolio management operations on machine learning.[33] An AI powered ETF (AIEQ) was issued in October 2017 that "fully utilized" AI as a method for stock selection.[34] Incubators such as the CloudQuant Trading Strategy Incubator, Merantix, and others are actively supporting the development of innovative AI and ML trading models.[35,36]

 

As previously discussed, explainable techniques can provide model designers with important insights that can allow them to refine and improve their models, though they may come with costs in terms of model efficacy or development effort. Recent advances have enhanced the efficacy and decreased the costs associated with explainable techniques, in some cases to the point where those tradeoffs may no longer be a substantive factor, yielding greater potential for adoption of these techniques to enhance profitability of modeling efforts in trading, portfolio management, and related operations.

Over and above model efficacy, regulators have issued guidance and proposed "Risk Principles" that may inform the role of explainability in the model design process. FINRA has specifically noted that supervisory responsibilities under FINRA Rule 3110 applies to all of a firm's associated persons and its businesses, including supervising activities related to AI applications,[37] and promulgates supervisory best practices for firms engaging in algorithmic trading strategies. Its guidance includes, "at a minimum, a basic summary description of algorithmic strategies that enables supervisory, compliance and regulatory staff to understand the intended function of an algorithm without the need to resort, as an initial matter, to direct code review."[38]

In response to concerns about potential systemic risk associated with algorithmic trading, Regulation Automated Trading (Regulation AT) was proposed by the CFTC in 2015 that would impose substantial reporting requirements with respect to trading algorithms,[39] including requirements to establish and maintain a source code repository that documents the strategy and design of proprietary algorithmic trading software. That rule would potentially have made source code available for inspection by the CFTC or Department of Justice without a subpoena[40] and "attracted intense opposition."[41]

Regulation AT has recently been abandoned and replaced with a proposed rule addressing Electronic Trading Risk Principles. The proposed Electronic Trading Risk Principles regulation provides that exchanges "must adopt and implement rules governing market participants … to prevent, detect, and mitigate market disruptions or system anomalies associated with electronic trading."[42] There has been no clear indication yet of the form these rules will take in practice.

However, in the context of trading model design, beyond any specific regulatory requirements, whether and the degree to which explainable techniques are adopted ultimately remains a judgment call based on perceived costs and benefits.

### AI and ML in the Mid Office

In the context of risk management, algorithmic trading as a general matter can give rise to its own set of particular risk factors. Among other things, algorithmic trading leverages decision-making in trading operations, and with it potential risk.[43] The exceptional speed with which trading strategies can be executed can lead to exceptionally rapid accumulation of losses and limit the capability of management to perform timely interventions. Algorithmic trading may also add complexity to management and oversight, necessitating additional internal auditing, testing, and change management functions.[44]

With algorithmic trading based on traditional conditional logic that relies on a set of rules implemented in software, the software (if properly implemented) behaves in an entirely predictable manner, and given sufficient resources, internal auditing and review processes can determine the reason for any observed behaviors based on the specific rules implemented in the algorithm.

Many AI and ML techniques, however, may make use of hundreds, or even thousands, of inputs and result in dynamics that are characteristically challenging to explain, even to the model designers.[45] Explainable techniques may allow the modeler some insight into unexpected or aberrant behaviors.

Moreover, over and above model opacity, AI and ML techniques more generally can bring their own, unique set of attendant risks,[46] leading FINRA to recently comment that "[f]irms 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."[47]

With regards to explainability, FINRA characterized it as a "key consideration in the model risk management process for AI-based applications" and included review of the explainability of model output among "potential areas for firms to consider as they update their model risk management programs to reflect the use of AI models."[48]

In addition to risk management, AI and ML techniques have seen adoption in mid-office monitoring and compliance operations across a range of market participants.

In its recent FINRA survey of AI in the securities industry, industry participants indicated "they are spending significant time and resources in developing AI-based applications to enhance their compliance and risk management functions"[49] including surveillance and monitoring functions that "move beyond 'traditional rule-based systems ….'"[50]

FINRA commented specifically in the survey on explainability in compliance, audit, and risk management functions:

> 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.

In addition to internal monitoring and compliance functions adopted by securities market participants, exchanges have introduced AI and ML based trade surveillance[51] and a variety of third-party vendors have introduced AI- and ML-based surveillance products, with monitoring services targeting trading activities as well as communications.[52]

As an example of the ready application of AI and ML techniques to trade surveillance, as well as potential benefits and limitations of explainable techniques, consider "disruptive" trading activity that "is, or of the character of" spoofing, which the Dodd-Frank Act defines as "bidding or offering with intent to cancel the bid or offer before execution."[53]

Market participants have focused considerable attention and resources on implementing trade surveillance systems to detect spoofing and other trading activity of concern to regulators. In some cases, regulators have required targets of regulatory actions to maintain systems and controls "reasonably designed" to "detect patterns of activity that might constitute spoofing activity."[54]

Regulators have identified a variety of patterns of trading behavior they consider to be spoofing, but these patterns are merely examples and do not constitute a precise definition. Moreover, absent broader context, these trading patterns may be indistinguishable from legitimate trading activity and do not in and of themselves establish scienter. Commentators and defendants alike have argued that the definition of spoofing is impermissibly vague, though the courts have in various circumstances disagreed.[55]

The pattern matching capabilities of AI and ML have made them a natural fit for trade surveillance for these kinds of trading patterns by exchanges, in-house compliance operations, and third-party service providers.[56] Among third-party monitoring and compliance service providers, incumbent providers, which had previously relied on rule-based approaches, have begun adopting AI and ML techniques while new service providers have emerged whose products have been built from the ground up based on AI and ML techniques.[57] And securities market regulators have been focusing on AI and ML techniques in their efforts to identify and prosecute malfeasance.[58]

While these techniques may identify patterns of trading activity consistent with activity regulators consider to be of concern, they may not provide compliance personnel with insights into why a particular cluster of activity was identified or the broader context of the activity. For that reason, products built using AI and ML techniques have begun including explainability features.[59] However, for trading patterns such as spoofing where interpretation may involve nuance, in making a determination on the nature of the trading activity, explainable AI and ML techniques in their present form may be able to enhance, but are by no means yet a substitute, for review by a human with relevant expertise.

Going forward, the adoption of AI and ML will likely continue to extend its reach beyond the regulatory contexts discussed in this article. Though it has not yet gained traction in areas such as regulatory reporting, in a recent joint white paper by Wolters Kluwer and PwC, based on a proof-of-concept exercise, the authors concluded that it was "likely that production reporting mechanisms will incorporate AI and ML in the near future."[60]

## Conclusion

The selection of explainable AI and ML techniques have improved substantially in recent years as has the efficacy of those techniques, and it remains an extremely active subject of research. As such, we anticipate that improvements in these techniques will continue going forward, as will incentives for broader adoption across a spectrum of securities market participants and operations. The choices various securities markets participants make to incorporate explainable AI and ML techniques into their modeling efforts and the nature of that adoption are multifaceted and will depend on the particular circumstances they face. We recommend that securities market participants who may be deploying AI and ML techniques in their organizations keep abreast of developments in the field and consider the adoption of explainable AI and ML techniques where they have the potential to enhance profitability, improve operations, and mitigate risk.

## ENDNOTES

1 The authors would like to thank Nicholas Schmidt, AI Practice Leader at BLDS, for helpful comments and suggestions.

2 For additional discussion of AI and ML model opacity in the context of securities trading, see Collin Starkweather and Izzy Nelken, "Artificial Intent: AI On The Trading Floor," Law360, January 23, 2019, available at https://www.law360.com/articles/1119871/artificial-intent-ai-on-the-trading-floor (subscription required).

3 To be clear on the distinction between the two, ML comprises many of the tools and techniques that underlie (and is considered to comprise "a large part of") AI, though the terms are often used interchangeably (Sara Castellanos, "What Exactly is Artificial Intelligence," The Wall Street Journal, December 6, 2018, available at https://www.wsj.com/articles/what-exactly-is-artificial-intelligence-1544120887 (subscription required)). To define the terms more precisely, AI has been said to "encompass[] the techniques used to teach computers to learn, reason, perceive, infer, communicate and make decisions similar to or better than humans," whereas ML broadly refers to the tools and techniques that "provide the technical basis of … the extraction of implicit, previously unknown, and potentially useful information from data." (*Ibid*; Ian H. Witten and Eibe Frank, Data Mining: Practical Machine Learning Tools and Techniques, Second Edition, 2005, pp. xxiii-xxiv and Table of Contents).

4 To be specific, model size and complexity in this context refers to the number of variables considered and their potential interactions rather than, for example, the number of observations in the dataset being analyzed.

5 "Artificial Intelligence (AI) in the Securities Industry," FINRA, June 2020, p. 1, available at https://www.finra.org/rules-guidance/key-topics/fintech/report/artificial-intelligence-in-the-securities-industry.

6 Broadly speaking, in the current vernacular, "explainability" is more often associated with post hoc summarization of AI techniques and the term explainable AI (XAI), whereas "interpretability" is more often associated with ML model architecture designed for the purpose of providing interpretable models and the term ML interpretability (MLI). However, the usage of these terms has evolved rapidly and may continue to evolve going forward.

7 Our definition is similar to that offered in Finale Doshi-Velez and Been Kim, "Towards a Rigorous Science of Interpretable Machine Learning," arXiv:1702.08608, March 2, 2017, available at https://arxiv.org/abs/1702.08608. While Hall and Gill support the Doshi-Valez and Kim definition, they also note that "more leading researchers use interpretable to refer to directly transparent modeling mechanisms …." Patrick Hall and Navdeep Gill, "An Introduction to Machine Learning Interpretability," 2nd Ed., O'Reilly Media, Inc., 2019.

8 See, e.g. Agam Shah, "Consortium of Tech Firms Sets AI Benchmarks," Wall Street Journal, June 25, 2019, available at https://www.wsj.com/articles/consortium-of-tech-firms-sets-ai-benchmarks-11561490769 (subscription required); Olga Davydova, "7 types of Artificial Neural Networks for Natural Language Processing," Medium, September 26, 2017, available https://medium.com/@datamonsters/artificial-neural-networks-for-natural-language-processing-part-1-64ca9ebfa3b2; Isma-Ilou Sadou, "Convolutional neural networks and their application in self driving cars," Medium, June 3, 2018, available at https://medium.com/@ismailou.sa/convolutional-neural-networks-and-their-application-in-self-driving-cars-33fa0a1625c8.

9 These include widely-deployed classes of models including, for example, multilayer perceptron (MLP) neural networks (also often referred to as "deep learning" models) and gradient boosting machines (GBMs).

10 See, e.g., Anne Bonner, "The Complete Beginner's Guide to Deep Learning: Convolutional Neural Networks and Image Classification," Towards Data Science, February 2, 2019, available at https://towardsdatascience.com/wtf-is-image-classification-8e78a8235acb.

11 Deep learning can be defined as "a form of representation learning that uses multiple transformation steps to create very complex features," where representation learning can be defined as techniques which "transform features into some intermediate representation prior to mapping them to final predictions." Ian H. Witten, Eibe Frank, Mark A. Hall, and Christopher J. Pal, Data Mining: Practical Machine Learning Tools and Techniques, Fourth Edition, 2017, p. 418.

12 See, e.g., Nicholas Schmidt, Bernard Siskin, Syeed Mansur, "How Data Scientists Help Regulators and Banks Ensure Fairness when Implementing Machine Learning and

 

Artificial Intelligence Models," 2018 Conference on Fairness, Accountability, and Transparency, February 1, 2018; Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, "Machine Bias," ProPublica, May 23, 2016, available at https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing; Rebecca Wexler, "When a Computer Program Keeps You in Jail," The New York Times, June 13, 2017, available at https://www.nytimes.com/2017/06/13/opinion/how-computers-are-harming-criminal-justice.html (subscription required).

13 Ziad Obermeyer, Brian Powers, Christine Vogeli and Sendhil Mullainthan, "Dissecting racial bias in an algorithm used to manage the health of populations", Science, Vol. 366, Issue 6464, pp. 447-453.

14 Patrice Alexander Ficklin, Tom Pahl, and Paul Watkins, "Innovation spotlight: Providing adverse action notices when using AI/ML models," Consumer Financial Protection Bureau, July 7, 2020, available at https://www.consumerfinance.gov/about-us/blog/innovation-spotlight-providing-adverse-action-notices-when-using-ai-ml-models/. Similarly, the Fair Credit Reporting Act (FCRA) requires that creditors must disclose any key factors that adversely affect credit scores when an adverse action is based on a credit score.

15 *Ibid*.

16 Tom Taulli, "CCPA: What Does It Mean For AI (Artificial Intelligence)?" Forbes, December 27, 2019, available at https://www.forbes.com/sites/tomtaulli/2019/12/27/ccpa--what-does-it-mean-for-ai-artificial-intelligence/ (subscription required).

17 *Ibid*.

18 James Guszcza, Iyad Rahwan, Will Bible, Manuel Cebrian, and Vic Katyal, "Why We Need to Audit Algorithms," Harvard Business Review, November 28, 2018, available at https://hbr.org/2018/11/why-we-need-to-audit-algorithms. See also Bryce Goodman and Seth Flaxman, "European Union regulations on algorithmic decision-making and a 'right to explanation'," arXiv:1606.08813v3, August 31, 2016, available at https://arxiv.org/abs/1606.08813v3.

19 Cynthia Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," arXiv:1811.10154v3, September 22, 2019, p. 2, available at https://arxiv.org/pdf/1811.10154.pdf.

20 As Cynthia Rudin, Professor of Computer Science, Electrical and Computer Engineering, and Statistical Science and Principal Investigator at Duke University's Prediction Analysis Lab, observed, "any formal process for defining knowledge from data, would require an iterative process, where one refines the data processing after interpreting the results." Cynthia Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," arXiv:1811.10154v3, September 22, 2019, p. 2, available at https://arxiv.org/pdf/1811.10154.pdf.

21 See, e.g., Defense Advanced Research Projects Agency, "Broad Agency Announcement, Explainable Artificial Intelligence (XAI)," DARPA-BAA-16-53, August 10, 2016, p. 14, available at https://www.darpa.mil/attachments/DARPA-BAA-16-53.pdf.

22 Patrick Hall and Navdeep Gill, "An Introduction to Machine Learning Interpretability," 2nd Ed., O'Reilly Media, Inc., 2019, p. 1. H2O.ai is the producer of H2O, which it states is "the leading open source data science and machine learning platform used by nearly half of the Fortune 500 and trusted by over 18,000 organizations and hundreds of thousands of data scientists around the world." (See https://www.h2o.ai/company/.)

23 With this approach, provided a model is not overfit, factors that may suggest such a tradeoff include underlying data-generating processes that are highly complex or are not well understood.

24 The most prominent examples of this approach use the Shapley value to determine the relative contribution of various inputs to the observed model estimates. The Shapley value is a foundational mathematical concept for cooperative game theory first proposed by economist Lloyd Shapley in 1951, for which he received a Nobel Prize in 2012.

25 See, e.g., Linwei Hu, Jie Chen, Vijayan N. Nair, and Agus Sudjianto, "Locally Interpretable Models and Effects based on Supervised Partitioning (LIME-SUP)," arXiv:1806.00663, June 1, 2018, available at https://arxiv.org/abs/1806.00663.

26 Cynthia Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," arXiv:1811.10154v3, September 22, 2019, pp. 3-5, available at https://arxiv.org/pdf/1811.10154.pdf.

27 Cynthia Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," arXiv:1811.10154v3, September 22, 2019, available at https://arxiv.org/pdf/1811.10154.pdf.

28 See, e.g., Navdeep Gill, Patrick Hall, Kim Montgomery, and Nicholas Schmidt, "A Responsible Machine Learning Workflow with Focus on Interpretable Models, Post-hoc Explanation, and Discrimination Testing," Information, February 2020. See also Patrick Hall and Navdeep Gill, "An Introduction to Machine Learning Interpretability," 2nd Ed., O'Reilly Media, Inc., 2019, pp. 3-4.

29 More opaque models include, for example, multilayer perceptron (MLP) neural networks and gradient boosting machines (GBMs).

30 See, e.g., Navdeep Gill, Patrick Hall, Kim Montgomery, and Nicholas Schmidt, "A Responsible Machine Learning Workflow with Focus on Interpretable Models, Post-hoc Explanation, and Discrimination Testing," Information, February 2020.

31 "Artificial Intelligence (AI) in the Securities Industry," FINRA, June 2020, pp. 5-6, available at https://www.finra.org/rules-guidance/key-topics/fintech/report/artificial-intelligence-in-the-securities-industry.

32 "Artificial intelligence takes on the stock market," BBC Click, February 10, 2016, available at https://www.bbc.com/news/av/technology-35405336/artificial-intelligence-takes-on-the-stock-market. Prominent hedge funds such as Bridgewater Associates, Renaissance Technologies, and Two Sigma Investments, with over $240 billion in assets under management (AUM) as of May 2020, have pioneered AI and ML in recent years. ("Machine Learning for Stock Trading

Strategies," Nanalyze, May 21, 2020, available at https://www.nanalyze.com/2020/05/machine-learning-for-stock-trading-strategies/.)

33 Ryan Vlastelica, "BlackRock's 'robot' stock-pickers are more of a tweak than a game-changer," Financial News, April 5, 2017, available at https://www.fnlondon.com/articles/blackrocks-robot-stock-pickers-are-more-of-a-tweak-than-a-game-changer-20170405.

34 AIEQ is the "[f]irst and only actively managed ETF to fully utilize artificial intelligence as a method for stock selection." Its issuer states "[t]he [AI] system mimics a team of 1,000 research analysts working around the clock analyzing millions of data points each day." See https://etfmg.com/funds/aieq/; Lizzy Gurdus, "This ETF run by a robot is beating the market—here's how it works ," CNBC, August 3, 2019, available at https://www.cnbc.com/2019/08/02/this-etf-run-by-a-robot-is-beating-the-marketheres-how-it-works.html.

35 See, e.g., CloudQuant's description of the Trading Strategy Incubator at https://info.cloudquant.com/trading-strategy-incubator/ and "Machine Learning for Stock Trading Strategies," Nanalyze, May 21, 2020, available at https://www.nanalyze.com/2020/05/machine-learning-for-stock-trading-strategies/.

36 Of note, while AI and ML techniques have seen rapid adoption in trading operations, it remains to be seen the degree to which these trading models are the kind of "game changer" that many advocates anticipate. AllianceBernstein analysts have expressed skepticism that AI will be able to generate significantly different results by "the mere fact that analyzing more and more data results in increasingly similar strategies," and at hedge funds and incubators alike, a new generation of modelers are becoming more intimately acquainted with Eugene Fama's Efficient Markets Hypothesis. AIEQ, advertised by the issuer as the "[f]irst and only actively managed ETF to fully utilize artificial intelligence as a method for stock selection," was reported to have "failed miserably" in its first couple of years of operation, substantially underperforming the S&P 500 despite "over a million market signals, news articles, and 6,000 U.S. companies analyzed daily." (Evelyn Cheng, "Just 10% of trading is regular stock picking, JPMorgan estimates," CNBC, June 13, 2017, available at https://www.cnbc.com/2017/06/13/death-of-the-human-investor-just-10-percent-of-trading-is-regular-stock-picking-jpmorgan-estimates.html; Bradley Hope and Juliet Chung, "The Future Is Bumpy: High-Tech Hedge Fund Hits Limits of Robot Stock Picking," Wall Street Journal, December 11, 2017, available at https://www.wsj.com/articles/the-future-is-bumpy-high-tech-hedge-fund-hits-limits-of-robot-stock-picking-1513007557 (subscription required); Jason Bowling, "What Happened When I Tried Market Prediction With Machine Learning," Towards Data Science, November 22, 2019, available at https://towardsdatascience.com/what-happened-when-i-tried-market-prediction-with-machine-learning-4108610b3422; Keith Speights, "An AI-Powered ETF Failed Miserably at Beating the Market in 2018 -- Here's What You Can Learn From Its Mistakes," Motley Fool, January 6, 2019, available at https://www.fool.com/investing/2019/01/06/an-ai-powered-etf-failed-miserably-at-beating-the.aspx; https://etfmg.com/funds/aieq/.)

37 "Artificial Intelligence (AI) in the Securities Industry," FINRA, June 2020, p. 12, available at https://www.finra.org/rules-guidance/key-topics/fintech/report/artificial-intelligence-in-the-securities-industry.

38 "Guidance on Effective Supervision and Control Practices for Firms Engaging in Algorithmic Trading Strategies", FINRA Regulatory Notice 15-09, March 26, 2015, available at https://www.finra.org/rules-guidance/notices/15-09.

39 The CFTC proposed Regulation Automated Trading (RegAT) in November 2015, which would require entities engaging in automated trading to "put in place pre-trade risk controls, establish standards for how they develop, test, and monitor their ATS, make annual reports to their designated contract markets attesting to their pre-trade compliance controls, and register with the CFTC if they haven't already. Furthermore, the CFTC also plans to require automated trading firms to keep their trading software's source code available for inspection." "Regulators, quant up! New rules from FINRA, SEC and CFTC target automated algorithmic trading," Bloomberg, February 29, 2016, available at https://www.bloomberg.com/professional/blog/regulators-quant-up-new-rules-from-finra-sec-and-cftc-target-automated-algorithmic-trading/; Peter Y. Malyshev, Kari S. Larsen and Michael S. Selig, "Regulation Automated Trading: The CFTC's Supplemental Proposal and Beyond," Thomson Reuters Futures & Derivative Law Report, January 2017, available at https://www.reedsmith.com/-/media/files/perspectives/2017/01/regulation-automated-trading/regulationautomatedtrading.pdf.

40 "Algorithmic Trading Regulation With Enforcement Edge," The National Law Review, February 23, 2016, available at https://www.natlawreview.com/article/algorithmic-trading-regulation-enforcement-edge.

41 "Statement of Commissioner Dan M. Berkovitz on Proposed Rules for Electronic Trading Risk Principles and Withdrawal of Regulation AT," CFTC Statement, June 25, 2020, available at https://www.cftc.gov/PressRoom/SpeechesTestimony/berkovitzstatement062520.

42 "CFTC Approves Two Final Rules and Two Proposed Rules at June 25 Open Meeting - Agency Also Withdraws Proposed Rule and Supplemental Proposal for Regulation AT," CFTC Press Release, June 25, 2020, available at https://www.cftc.gov/PressRoom/PressReleases/8188-20; "Electronic Trading Risk Principles," CFTC Notice of Proposed Rulemaking, June 25, 2020, p. 20, available at https://www.cftc.gov/media/4056/votingdraft062520b/download; Michelle Price, "U.S. derivatives regulator to scrap Obama-era algo-trading rule," Reuters, October 17, 2018, available at https://uk.reuters.com/article/uk-usa-cftc-trading/u-s-derivatives-regulator-to-scrap-obama-era-algo-trading-rule-idUKKCN1MR2E7.

[43] "Algorithmic Trading Briefing Note," Senior Supervisors Group, April 2015, available at https://www.newyorkfed.org/medialibrary/media/newsevents/news/banking/2015/SSG-algorithmic-trading-2015.pdf. In addition to considerations at the level of the trading desk, regulators and market commentators have also debated the potential impact of such leveraging at the systemic level. See also, e.g., Robin Wigglesworth, "Volatility: how 'algos' changed the rhythm of the market," Financial Times, January 8, 2019, available at https://www.ft.com/content/fdc1c064-1142-11e9-a581-4ff78404524e (subscription required).

[44] See, e.g., "Guidance on Effective Supervision and Control Practices for Firms Engaging in Algorithmic Trading Strategies", FINRA Regulatory Notice 15-09, March 26, 2015, available at https://www.finra.org/rules-guidance/notices/15-09.

[45] For example, Demis Hassabis, the CEO of Google's DeepMind, which created AlphaZero and AlphaGo, the AI that beat the reigning world chess and go champions, said of AlphaZero that "It doesn't play like a human, and it doesn't play like a program, it plays in a third, almost alien, way .... It's like chess from another dimension." Will Knight, "Alpha Zero's 'Alien' Chess Shows the Power, and the Peculiarity, of AI," MIT Technology Review, December 8, 2017, available at https://www.technologyreview.com/s/609736/alpha-zeros-alien-chess-shows-the-power-and-the-peculiarity-of-ai/.

[46] See, e.g., Andrew Burt and Patrick Hall, "What to Do When AI Fails," May 18, 2020, available at https://www.oreilly.com/radar/what-to-do-when-ai-fails/; Bluford Putnam, "Using AI to analyze financial markets," CME Group, November 18, 2019, available at https://www.cmegroup.com/education/featured-reports/videos/using-ai-to-analyze-financial-markets.html.

[47] "Artificial Intelligence (AI) in the Securities Industry," FINRA, June 2020, p. 11, available at https://www.finra.org/rules-guidance/key-topics/fintech/report/artificial-intelligence-in-the-securities-industry.

[48] *Ibid*.

[49] *Id*., p. 8.

[50] *Ibid*.

[51] See, e.g., James Rundle, "Nasdaq Deploys AI to Detect Stock-Market Abuse," Wall Street Journal, July 31, 2019, available at https://www.wsj.com/articles/nasdaq-deploys-ai-to-detect-stock-market-abuse-11564571534 (subscription required).

[52] See, e.g., "TT® Score: Trade Surveillance with Machine Learning," Trading Technologies, available at https://www.tradingtechnologies.com/data/surveillance/tt-score/; "eCommunications Surveillance & Investigations," Catelas, available at https://catelas.com/wp-content/uploads/2018/11/Catelas-eCommunications-Surveillance-Investigations.pdf.

[53] Dodd-Frank Act § 767.

[54] Jay Biondo, "Spoofing, Surveillance, and Supervision," National Society of Compliance Professionals, June 13, 2018, available at https://medium.com/@Trading_Tech/spoofing-surveillance-and-supervision-7aef16a45fee.

[55] See, e.g., Catriona Coppler, "The Anti-Spoofing Statute: Vague as Applied to the 'Hypothetically Legitimate Trader'," American University Business Law Review, Volume 5:2, 2016, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3157409.

[56] See, e.g., "For the First Time, Nasdaq Is Using Artificial Intelligence to Surveil U.S. Stock Market," Nasdaq, November 7, 2019, available at https://www.nasdaq.com/articles/for-the-first-time-nasdaq-is-using-artificial-intelligence-to-surveil-u.s.-stock-market; "TT® Score: Trade Surveillance with Machine Learning," Trading Technologies, available at https://www.tradingtechnologies.com/data/surveillance/tt-score/.

[57] For example, while the products of incumbent trade surveillance providers such as Nasdaq SMARTS were initially built prior to the widespread adoption of AI and ML techniques, products of more recent market entrants such as Trading Technologies' TT Score was built from the ground up to take advantage of AI and ML techniques.

[58] For example, the CFTC recently announced the inaugural competition under the Science Prize Act of 2015 was a challenge to participants to "leverage artificial intelligence (AI) and other technologies to identify unregistered foreign entities potentially engaging in illegal activity subject to the CFTC's jurisdiction." CTFC, "LabCFTC Launches 'Project Streetlamp' Science Prize Competition," Press Release No. 8154-20, available at https://www.cftc.gov/PressRoom/PressReleases/8154-20.

[59] For example, Trading Technologies' TT Score provides a setting that allows users to adjust the degree to which the model produces explainable results.

[60] Wouter Delbaere and Dr. Andreas Deppeler, "Artificial Intelligence in Regulatory Reporting," White Paper, Wolters Kluwer and PwC, May 29, 2019.