

12-27-2019

Artificial Intelligence & Artificial Prices: Safeguarding Securities Markets from Manipulation by Non-Human Actors

Daniel W. Slemmer

Follow this and additional works at: <https://brooklynworks.brooklaw.edu/bjcfcl>



Part of the [Banking and Finance Law Commons](#), [Business Organizations Law Commons](#), [Internet Law Commons](#), [Organizations Law Commons](#), [Science and Technology Law Commons](#), and the [Securities Law Commons](#)

Recommended Citation

Daniel W. Slemmer, *Artificial Intelligence & Artificial Prices: Safeguarding Securities Markets from Manipulation by Non-Human Actors*, 14 *Brook. J. Corp. Fin. & Com. L.* ().

Available at: <https://brooklynworks.brooklaw.edu/bjcfcl/vol14/iss1/11>

This Note is brought to you for free and open access by the Law Journals at BrooklynWorks. It has been accepted for inclusion in Brooklyn Journal of Corporate, Financial & Commercial Law by an authorized editor of BrooklynWorks.

ARTIFICIAL INTELLIGENCE & ARTIFICIAL PRICES: SAFEGUARDING SECURITIES MARKETS FROM MANIPULATION BY NON-HUMAN ACTORS

ABSTRACT

Securities traders are currently competing to use Artificial Intelligence (A.I.) in order to make more profitable decisions in the marketplace. While A.I. provides superior abilities in recognizing market patterns, its complexity can obscure its decision-making process beyond human comprehension. Problematically, the current securities laws prohibiting manipulation of securities prices rest liability for violations on a trader's intent. In order to prepare for A.I. market participants, both courts and regulators need to accept that human concepts of decision-making will be inadequate in regulating A.I. behavior. However, the wealth of case law in the market manipulation doctrine need not be cast aside. Industry regulators should instead require A.I. users to harness the power of their machines to provide meaningful feedback in order to both detect potential manipulations and create evidentiary records in the event that allegations of A.I. manipulation arise.

INTRODUCTION

In an era where massive amounts of information are easily accessible via the internet and advances in computer hardware are continuously improving our ability to process more data, it is common to hear terms like “artificial intelligence” (A.I.) and “machine learning” in discussions of the latest cutting-edge technology. While these concepts are applied in a range of settings, they generally refer to the use of computers to think and act independently of humans, particularly when predictions need to be made based on large sets of data.¹ These methods are becoming increasingly relevant to financial institutions, especially those trading in securities markets.² Securities regulations, however, do not contemplate liability for A.I. actors and some violations require a finding of human intent.³ Often, the intent of an A.I. system's developer will not be reflected in the A.I.'s

1. See generally Calum McClelland, *The Difference Between Artificial Intelligence, Machine Learning, and Deep Learning*, MEDIUM (Dec. 4, 2017), <https://medium.com/iotforall/the-difference-between-artificial-intelligence-machine-learning-and-deep-learning-3aa67bff5991> (describing an example where “feeding huge amounts of data” to an algorithm can allow it to adjust itself and increase accuracy in determining whether an image contains a cat).

2. See Angelo Calvello, *With BlackRock's Artificial Intelligence Pivot, the Rubicon Has Been Crossed*, INSTITUTIONAL INV. (May 1, 2017), <https://www.institutionalinvestor.com/article/b1505p7qq0b511/with-blackrocks-artificial-intelligence-pivot-the-rubicon-has-been-crossed>.

3. The Securities and Exchange Act of 1934, which prohibits conduct like manipulation of securities prices, requires intent as an element for liability. See *ATSI Comme'ns, Inc. v. Shaar Fund, Ltd.*, 493 F.3d 87, 101 (2d Cir. 2007).

observable characteristics due to its complexity.⁴ Therefore, if A.I. actors exhibit harmful behavior in securities markets but do not reveal their intent, then there may be no one to hold accountable under the current securities regulations.

This Note suggests that regulatory and judicial action are necessary to prevent A.I. developers from escaping liability for harmful market behavior. Specifically, this Note addresses the actions that can be taken to prevent A.I. from manipulating securities prices and to support enforcement against developers⁵ of manipulative A.I. Part I of this Note explains the benefits and limitations of A.I. and the implications of using A.I. to trade securities. Part II discusses the legal doctrine of market manipulation and the problematic absence of regulations for A.I. traders. Part III explains how a clarified definition of market manipulation coupled with industry testing standards for A.I. could help prevent and detect A.I. manipulation. Part IV describes how courts might then revisit the existing doctrine of market manipulation to make use of new industry standards in light of the increasing involvement of A.I. in securities markets.

I. A.I. AND ITS POTENTIAL . . . PROBLEMS

A. THE TECHNOLOGY, THE INDUSTRY VALUE, AND THE LEGAL GAP

To begin, A.I. is defined simply as the ability of a machine, such as a computer, to perform tasks commonly associated with intelligent beings.⁶ “Machine learning” is a method of achieving A.I. without explicitly instructing a machine how to perform its task.⁷ For example, one subset of machine learning called “deep learning” makes use of artificial neural networks (ANNs), which are computer architectures roughly modeled after the human brain.⁸ The novel structure of ANNs allows machines to reach unprecedented levels of complexity in analyzing abstract patterns, which

4. See Yavar Bathaee, *The Artificial Intelligence Black Box and the Failure of Intent and Causation*, 31 HARV. J.L. & TECH. 889, 901 (2018) (discussing how both complexity in an A.I. machine’s structure and its ability to process a large number of variables to reach a decision make it difficult, if not impossible, to understand exactly how a black box machine is arriving at a decision).

5. While “developers” in a colloquial sense might refer only to those who are directly involved in training the A.I., this Note uses “developers” more broadly, to include any parties or entities that fund the development of and profit from the operation of the A.I.

6. Stefan van Duin & Naser Bakhshi, *Part 1: Artificial Intelligence Defined*, DELOITTE (Mar. 28, 2017), <https://www2.deloitte.com/nl/nl/pages/data-analytics/articles/part-1-artificial-intelligence-defined.html>.

7. Rather, the machine is given training data to process; as the machine “trains,” it recognizes patterns in the data. From these patterns, the machine internally adjusts its own algorithm in order to make more accurate categorizations or predictions about future data. McClelland, *supra* note 1.

8. Charlie Crawford, *An Introduction to Deep Learning*, ALGORITHMIA (Nov. 4, 2016), <https://blog.algorithmia.com/introduction-to-deep-learning/>.

humans have channeled into impressive and beneficial applications.⁹ For instance, deep learning is notable for its implementation in systems that guide autonomous vehicles, translate language in real time, and that are superior to humans in other abstract tasks—like playing the ancient Chinese game Go.¹⁰ These definitions and examples are somewhat cursory considering the rate at which A.I. technology is evolving and being applied to a growing number of tasks. The important takeaway is that across all its applications, A.I.¹¹ is being used to more efficiently make predictions from large datasets.

Nowhere is the ability to make accurate predictions more valuable than in the world of finance. Financial institutions can use A.I. to anticipate future risks based on large sets of historical data, whether the task is determining the creditworthiness of potential borrowers or detecting fraudulent transactions by customers.¹² This predictive power is especially valuable in securities markets, where anticipating and responding to signals of market movement is already critical in gaining an advantage over competitors.¹³ Naturally, A.I. has attracted the interest of securities traders, and the world's leading asset managers are already experimenting with A.I. to make investment choices.¹⁴

While A.I.'s power might give traders a cause for celebration, the practical limitations of A.I. create a regulatory cause for concern. This is because A.I.'s ability to “learn” without explicit instruction correspondingly makes it difficult for humans to conclude how it makes decisions.¹⁵ This is referred to as the “black box” problem, and practitioners and academics have observed that it has the potential to create issues in various legal

9. ANNs process input data through “nodes,” as neurons do in the brain, and groups of node outputs are synthesized in “layers.” The “depth” in deep learning then occurs when multiple layers are interconnected, as additional layers allow for the aggregation of data from the previous layer to be reanalyzed with respect to other variables, achieving a higher level of abstraction in recognizing patterns. See Larry Hardesty, *Explained: Neural Networks*, MIT NEWS (Apr. 14, 2017), <http://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>.

10. Crawford, *supra* note 8.

11. While the term “A.I.” covers a broad range of machines from simple to complex, it has been used colloquially to describe complex, state of the art implementations of machine learning. For consistency, this is what will be referred to as A.I. for the purposes of this Note, even though only a small subset of artificially intelligent machines is likely relevant to the discussion herein.

12. See Larry D. Wall, *Some Financial Regulatory Implications of Artificial Intelligence*, 100 J. ECON. & BUS. 55, 58 (2018).

13. With computers trading at vastly greater speeds than humans, even 10 millisecond delays in responding to news with a market impact can lead to significant decreases in returns for traders. See Yesha Yadav, *The Failure of Liability in Modern Markets*, 102 VA. L. REV. 1031, 1067 (2016).

14. See, e.g., Wall, *supra* note 12 (discussing recent developments by Blackrock and J.P. Morgan to employ machine learning in choosing securities trades); Calvello, *supra* note 2.

15. See Bathae, *supra* note 4 (discussing how both complexity in a machine's structure and its ability to process a large number of variables to reach a decision make it difficult, if not impossible, to understand exactly how a black box machine is arriving at a decision).

contexts.¹⁶ The black box problem is particularly intimidating in instances where A.I. conduct is only governed by statutes that base liability in human-centric concepts like intent.¹⁷

This is currently the case for federal securities laws—specifically, the laws prohibiting market manipulation of securities prices. Market manipulation involves conduct “designed to deceive or defraud investors by controlling or artificially affecting the price of securities.”¹⁸ The Securities and Exchange Act of 1934 (the Act), which prohibits manipulation, requires scienter, or manipulative intent as an element for liability.¹⁹ Given the potential for A.I. use in securities trading and the attendant black box problem for A.I. decision-making,²⁰ there is the looming question as to how A.I. trader intent would be evidenced in a manipulation action. Bear in mind that the black box problem only pertains to *autonomous* trading systems, which are the focus of this Note. These are distinct from *automated* systems,²¹ which are already commonplace in securities markets.²²

B. A HYPOTHETICAL TO DEMONSTRATE THE RISKS OF A.I. IN SECURITIES MARKETS

There is no exhaustive list of manipulative conduct. Rather, the Act is “meant to prohibit the full range of ingenious devices that might be used to manipulate securities prices.”²³ But it is helpful to begin by considering how an A.I. machine might commit a classic example of market

16. See, e.g., *id.* at 906–28 (discussing the difficulties that may arise in constitutional, securities, and antitrust law due to the black box problem); Kevin Petrasic et al., *Three Big Questions About AI in Financial Services*, WHITE & CASE (July 18, 2017), <https://www.whitecase.com/publications/insight/ai-financial-services> (discussing the problems that artificial intelligence poses in the collusion, antitrust, and consumer discrimination contexts).

17. See, e.g., Bathae, *supra* note 4, at 892.

18. *Wilson v. Merrill Lynch & Co., Inc.*, 671 F.3d 120, 130 (2d Cir. 2011) (quoting *Ernst & Ernst v. Hochfelder*, 425 U.S. 185, 199 (1976)).

19. *ATSI Commc’ns, Inc. v. Shaar Fund, Ltd.*, 493 F.3d 87, 101 (2d Cir. 2007).

20. See Petrasic et al., *supra* note 16.

21. While automated and autonomous systems might seem largely the same in that they involve computers performing functions without much human intervention, automated systems do not necessarily create a black box problem. Even though they execute transactions independently, when their tasks are pre-defined, a developer’s intent may be imputed from the systems’ programming. See Stanton Jones, *Automation, Autonomy and the Messy In-Between*, INFO. SERVS. GROUP, <https://www.isg-one.com/research/articles/automation-autonomy-and-the-messy-in-between> (last visited Sept. 25, 2019). This does not mean automated systems have not created legal issues; they do not, however, create the evidentiary problems attendant to autonomous systems. See, e.g., *In re Facebook, Inc., IPO Sec. & Derivative Litig.*, 986 F. Supp. 2d 428, 441 (S.D.N.Y. 2013) (describing the automated system that failed and created liability for a stock exchange when the exchange ignored testing that suggested the automated system was not prepared for the anticipated order traffic).

22. See Yadav, *supra* note 13, at 1035 n.16 (“Reports suggest that algorithmic trading is responsible for around 70% of all trading in equities in the United States by volume.”).

23. *Santa Fe Indus., Inc. v. Green*, 430 U.S. 462, 477 (1977).

manipulation. First, imagine that an A.I. machine has broad access to information about securities markets and has the ability to transact in those markets.²⁴ The machine's developers use its wide capacity for data to feed it various categories of information associated with historical market behavior, thus training it to recognize market patterns.²⁵ At some point, the machine is allowed to actively trade on its predictions, guided by parameters that reflect the developing trader's own limitations.²⁶

If we assume that the developer trains the machine to make profitable trades, then as the machine learns it will assign more "weight"²⁷ to factors that correlate with profitability. Over time, the machine might recognize that certain variables like trading volume can correlate with a change in security price.²⁸ The machine might also recognize that its own transactions increase trading volume. Therefore, the machine might eventually calculate that by simultaneously buying and selling the same quantity of a security, it can itself create trading activity that drives up the security's price.²⁹ Since the transactions will offset, there will be a minimized risk of loss to the A.I., and it might then make other profitable transactions based on the anticipated rise in stock price.³⁰

Effectively, the system has just learned how to execute a "wash sale," which is an example of classically manipulative conduct under Section

24. See, e.g., Simon Kuttruf, *A Machine Learning Framework for Algorithmic Trading on Energy Markets*, MEDIUM (Jul. 16, 2018), <https://towardsdatascience.com/https-medium-com-skuttruf-machine-learning-in-finance-algorithmic-trading-on-energy-markets-cb68f7471475> (explaining how machine learning could be used to develop a real-time algorithmic trading system for European carbon emission certificates).

25. During training, the machine "guesses" what will happen, and when its predictions are fed back to it, the machine refines its algorithm by giving more weight to variables that are more predictive. This example suggests a "supervised" learning approach. Some machines can learn without supervision, but for the sake of a simplified example, they are not differentiated here. See Justin Gage, *Introduction to Machine Learning*, ALGORITHMIA (Mar. 19, 2018), <https://blog.algorithmia.com/introduction-to-machine-learning/>.

26. Even though a machine would make its predictions using historical market data, a developer will still have to manually set parameters related to a trader's own personal limitations, like "a threshold for the model confidence of a given prediction, what position do you place on the market, what position size, for how long do you hold a position in the given state of the market etc." See Kuttruf, *supra* note 24.

27. See, e.g., John McGonagle et al., *Backpropagation*, BRILLIANT, <https://brilliant.org/wiki/backpropagation/> (last visited Dec. 29, 2018) (explaining how ANNs adjust their weights via a method called "backpropagation").

28. See Walter Sun, *Relationship Between Trading Volume and Security Prices and Returns* (Feb. 6, 2003), <http://sbg.mit.edu/~waltsun/docs/AreaExamTR2638.pdf> (unpublished Ph.D. Area Exam Report, Massachusetts Institute of Technology Laboratory for Information and Decision Systems).

29. See Comment, *Market Manipulation and the Securities Exchange Act*, 46 YALE L.J. 624, 626–28 (1937) (discussing that when "the operator is both buyer and seller of the same stock . . . the price slowly rises, [and] a complex publicity apparatus is set into motion to aid the stimulation of demand").

30. For example, when a party holds an option to purchase a security at some price lower than the market value, a higher market value will allow the party to profit when they exercise the option and immediately sell at the heightened market price. See *id.* at 628.

10(b) of the Act.³¹ This is manipulative because there is no actual change in ownership, but the transactions give other investors a false impression that the security is being legitimately traded.³² Realistically, such classically manipulative behavior would be an indicator of improper training, and an improperly trained A.I. system is unlikely to be deployed by an enterprise whose mission is profitability.³³ However, this example demonstrates a couple of reasons why it may be more difficult to detect manipulation when A.I. is involved.

First, the wash sale example shows that market manipulation can involve capitalizing on simple relationships between interrelated market variables. A.I. decision-making already creates a “dimensionality problem,”³⁴ in that humans may be unable to conceptualize the patterns a machine is acting on. Considering A.I. machines’ superior ability to recognize patterns, it is not only conceivable that they could learn classic forms of manipulation, but that they might also manipulate in ways that would not be readily apparent to humans.³⁵ This might be especially problematic if multiple A.I. actors are on each side of a transaction or are working together.³⁶ Second, the hypothetical demonstrates that manipulative activity can be the result of multiple individual transactions, which might appear legitimate standing alone.³⁷ Again, while a conspicuously manipulative A.I. system is unlikely to be deployed by a sophisticated developer, bad actors might avoid detection by training A.I. to conceal any potential manipulations beneath blankets of legitimate trades. It

31. See *Santa Fe Indus., Inc. v. Green*, 430 U.S. 462, 476 (1977) (discussing that manipulation is essentially a “term of art” in reference to practices like wash sales).

32. See *Ernst & Ernst v. Hochfelder*, 425 U.S. 185, 205 n.25 (1976) (defining wash sales as “transactions involving no change in beneficial ownership”).

33. FINANCIAL STABILITY BOARD, ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING IN FINANCIAL SERVICES 18–19 (2017), <http://www.fsb.org/wp-content/uploads/P011117.pdf> (discussing the common sentiment that many “quant” funds are not comfortable using their trading models if they cannot understand how a particular prediction is made).

34. See Bathaee, *supra* note 4, at 892.

35. This could potentially compound an existing problem, as there is already uncertainty as to how much market manipulation occurs as a result of human actors, who are presumably easier to monitor. See Merritt B. Fox, Lawrence R. Glosten & Gabriel V. Rauterberg, *Stock Market Manipulation and Its Regulation*, 35 YALE J. ON REG. 67, 77–80 (2018) (discussing studies showing the existence of certain types of manipulative behavior but commenting that overall there is little data showing how often market manipulation occurs on the whole).

36. Ariel Ezrachi & Maurice E. Stucke, *Artificial Intelligence & Collusion: When Computers Inhibit Competition*, 2017 U. ILL. L. REV. 1775, 1783 (2017) (discussing the potential for independent machines to become interdependent on each other in the context of antitrust in markets where competitors are each using machines to set prices). The same potential in the securities market could lead to manipulative schemes where discrete machines each take a different role in a larger manipulation, which could lead to a scheme that is more complicated to detect.

37. See *ATSI Comm’ns, Inc. v. Shaar Fund, Ltd.*, 493 F.3d 87, 102 (2d Cir. 2007) (noting that “in some cases scienter is the only factor that distinguishes legitimate trading from improper manipulation.”).

will then be important to remember that while machines might be black boxes for any given decision, they are still ultimately only the product of their developer's training.³⁸

In either case, the question of what the A.I. "intended" to do might be impossibly convoluted. Moreover, both situations involve a risk of harm to securities markets, regardless of the developer's intent.³⁹ Although there have not been any formal allegations of A.I. market manipulation as of late, it is only a matter of time before courts will be tasked with adjudicating the legality of A.I. behavior given its potential utility in trading securities.⁴⁰ Furthermore, unlike far-off concerns regarding human-like and generally-skilled autonomous A.I.,⁴¹ the wash trade example demonstrates that the ability to trade securities, which could easily be assigned to A.I. given current technology, is all that is needed to manipulate securities prices.⁴²

II. CURRENT AVENUES FOR MANIPULATION LIABILITY

A. PRIMARY VIOLATIONS FOR MANIPULATION UNDER SECTION 10(B) OF THE ACT

First, it is important to understand why manipulative transactions, like wash sales, are prohibited. Section 10(b) of the Act makes it illegal "[t]o

38. The "Homunculus Fallacy" reminds us that "'there is no little person inside the program.' Instead algorithms act as they are programmed to act—no more, no less." See Karni Chagal-Feferkorn, *The Reasonable Algorithm*, 2018 U. ILL. J.L. TECH. & POL'Y 111, 132 (2018) (footnote omitted).

39. See Fox et al., *supra* note 35, at 115 (stating that wash sales "are socially harmful, however, in terms of the transaction reporting that they generate. Other traders do not know that the reported transactions involve wash or matched sales. They assume the reported prices reflect the valuations of persons who are engaged in genuine purchases and sales of the stock and they react accordingly to their disadvantage").

40. Both academics and those in industry have shown an interest in using machine learning to trade securities. See, e.g., Yadav, *supra* note 13, at 1066 n.125 (citing two instances where machine learning has been successfully applied to predict market changes in academia); Wall, *supra* note 12. Furthermore, industry leaders in securities trading have implemented effective A.I., and commentators suggest that competitors will have no choice but to implement their own A.I. methods or fall out of competition. See Calvello, *supra* note 2.

41. While an A.I. trader could potentially have more human-like abilities (e.g., communication), this Note only considers it practical to assess A.I. traders with abilities to trade securities given the current state of the art. This is because practical success with deep learning machines has been best achieved in narrow implementations (e.g., image recognition), and there is skepticism about the potential for deep learning machines to reach a "human-like" level of "general" intelligence. See Gary Marcus, *In Defense of Skepticism About Deep Learning*, MEDIUM, (Jan. 14, 2018), <https://medium.com/@GaryMarcus/in-defense-of-skepticism-about-deep-learning-6e8bfd5ae0f1>; see also Jason Pontin, *Greedy, Brittle, Opaque, and Shallow: The Downsides to Deep Learning*, WIRED (Feb. 2, 2018), <https://www.wired.com/story/greedy-brittle-opaque-and-shallow-the-downsides-to-deep-learning/> (discussing the ability of neural networks in performing narrow tasks but their ultimate difficulties with achieving more general intelligence).

42. See SEC v. Masri, 523 F. Supp. 2d 361, 372 (S.D.N.Y. 2007) ("[I]f an investor conducts an open-market transaction with the intent of artificially affecting the price of the security, and not for any legitimate economic reason, it can constitute market manipulation.").

use or employ, in connection with the purchase or sale of any security . . . any manipulative or deceptive device” in violation of rules promulgated by the Securities and Exchange Commission (SEC).⁴³ In turn, SEC Rule 10b-5(a) and (c) then make it unlawful “[t]o employ any device, scheme, or artifice to defraud” or “to engage in any act, practice, or course of business which operates or would operate as a fraud or deceit upon any person” in connection with the purchase or sale of securities.⁴⁴ Thus, acts like wash sales, which affect securities’ prices without any legitimate investment purpose, have become a well-established example of a manipulation.⁴⁵ Furthermore, the Supreme Court has recognized that the Act’s prohibitions are not limited to some predefined list of transactions; rather they are “meant to prohibit full range of ingenious devices that might be used to manipulate securities prices” and refer “generally to practices . . . that are intended to mislead investors by artificially affecting market activity.”⁴⁶ Note that not only manipulative acts but also material misstatements or omissions involving securities can give rise to Section 10(b) actions.⁴⁷

Rule 10b-5 requires plaintiffs bringing manipulation claims “to allege: (1) manipulative acts; (2) damage; (3) caused by reliance on an assumption of an efficient market free of manipulation; (4) scienter; (5) in connection with the purchase or sale of securities; (6) furthered by the defendant’s use of the mails or any facility of a national securities exchange.”⁴⁸ While scienter is required for liability, any conduct that misleads investors still runs counter to the purposes of the Act and goals of the SEC.⁴⁹

Because manipulation actions involve fraudulent conduct, they must be pled with particularity and specificity.⁵⁰ Courts have, however, relaxed the specificity standards for manipulation (as opposed to misstatement and omission) claims, as they have recognized that manipulation can involve facts solely in the defendant’s possession.⁵¹ Despite this recognition of an

43. 15 U.S.C. § 78j(b) (2012).

44. 17 C.F.R. § 240.10b-5 (2018).

45. *See* Ernst & Ernst v. Hochfelder, 425 U.S. 185, 205 (1976).

46. *Santa Fe Indus., Inc. v. Green*, 430 U.S. 462, 476–77 (1977).

47. *See* Cent. Bank of Denver, N.A. v. First Interstate Bank of Denver, N.A., 511 U.S. 164, 177 (1994) (stating that Section 10(b) “prohibits only the making of a material misstatement (or omission) or the commission of a manipulative act”).

48. *Sharette v. Credit Suisse Int’l*, 127 F. Supp. 3d 60, 77 (S.D.N.Y. 2015).

49. One of the goals of the Act, and thus the SEC, is to ensure that investors have access to the same information when making investment decisions. *See Santa Fe Indus.*, 430 U.S. at 477 (“[T]he fundamental purpose of the Act [is] ‘to substitute a philosophy of full disclosure for the philosophy of caveat emptor’” (quoting *Affiliated Ute Citizens of Utah v. United States*, 406 U.S. 128, 151 (1972))). Naturally, the idea of one party manipulating prices unbeknownst to other traders is difficult to square with this goal.

50. The Private Securities Litigation Reform Act (PSLRA) and Rule 9(b) of the Federal Rules of Civil Procedure require plaintiffs bringing actions for securities law violations and fraud to respectively plead with heightened particularity and specificity. *See* 15 U.S.C. § 78u-4(b)(1) (2012); FED. R. CIV. P. 9(b).

51. *ATSI Commc’ns, Inc. v. Shaar Fund, Ltd.*, 493 F.3d 87, 102 (2d Cir. 2007).

informational asymmetry, courts have still stressed the importance of the scienter requirement, with one court noting that “in some cases scienter is the only factor that distinguishes legitimate trading from improper manipulation.”⁵²

The pleading requirements and emphasis on scienter already make bringing market manipulation claims against human actors difficult,⁵³ but the involvement of black box A.I. will only further obfuscate the search for intent. Smoking-gun evidence is not always necessary as courts have recognized that in instances of classic manipulation, e.g., wash sales, “[t]he defendant’s manipulative intent can be inferred from the conduct itself.”⁵⁴ However, since manipulation can be effectuated through otherwise legitimate transactions, which do not reveal intent themselves, more complex forms of manipulation beyond the short list of classic examples might be more difficult to detect.⁵⁵ As discussed, A.I. systems will only compound this complexity problem.⁵⁶ In the absence of intent, it might then be fair to ask whether an A.I. developer knew there was a danger that their machine might trade manipulatively and failed to address it. This question essentially asks whether the developer was reckless.⁵⁷

It is settled law that recklessness can satisfy the scienter requirement for Section 10(b) claims brought for misstatements and omissions;⁵⁸ however, due to scienter’s critical role in differentiating manipulative from legitimate trading, some suggest that it can only be satisfied by a finding of intent and not mere recklessness.⁵⁹ If courts agree, the pleading requirements, informational asymmetries, and the black box problem together will likely

52. *Id.*

53. *See id.*

54. *SEC v. Masri*, 523 F. Supp. 2d 361, 367 (S.D.N.Y. 2007).

55. *See ATSI Commc’ns*, 493 F.3d, at 102 (noting that “in some cases scienter is the only factor that distinguishes legitimate trading from improper manipulation”).

56. This is not speculative, as high-frequency electronic trading systems have given rise to new forms of manipulation like “pinging” and “spoofing.” *See* Tom C.W. Lin, *The New Market Manipulation*, 66 *EMORY L.J.* 1253, 1288–90 (2017). Moreover, these systems are merely automated, and do not necessarily create the black box problem that autonomous traders will. *See* Jones, *supra* note 21.

57. While the circuits do not conform to one definition of recklessness, the Second Circuit defines it as “conduct which is highly unreasonable and which represents an extreme departure from the standards of ordinary care . . . to the extent that the danger was either known to the defendant or so obvious that the defendant must have been aware of it.” *Novak v. Kasaks*, 216 F.3d 300, 308 (2d Cir. 2000) (internal quotation marks omitted) (quoting *Rolf v. Blyth, Eastman Dillon & Co. Inc.*, 570 F.2d 38, 47 (2d Cir. 1978)).

58. *See Tellabs, Inc. v. Makor Issues & Rights, Ltd.*, 551 U.S. 308, 319 n.3 (2007) (acknowledging that each Circuit accepted that recklessness could satisfy the Section 10(b) scienter requirement but recognizing that the Circuits differed on the degree of recklessness required).

59. David Yeres et al., *U.S. Market Manipulation: Has Congress Given the CFTC Greater Latitude than the SEC to Prosecute Open Market Trading as Unlawful Manipulation? It’s Doubtful*, 38 *NO. 6 FUTURES & DERIVATIVES L. REP. NL* 1 (June 2018).

pose an insurmountable hurdle for plaintiffs attempting to allege manipulation by A.I. traders.

B. SECONDARY LIABILITY UNDER THE ACT

The poor fit of Section 10(b) to A.I. conduct raises a question as to the inherent fairness of employing an A.I. trader in the first place: if certain acts are manipulative, why should anybody besides the party who stood to gain from the A.I.'s efficiencies bear the cost of its actions when it impairs others in the market? Questions in this vein are not uncommon throughout the law,⁶⁰ and legislatures and judges have accordingly enacted and applied secondary liability to address these concerns in multiple contexts. This concern was the motivation behind Section 20(a) of the Act, which creates joint and several liability for “[e]very person who, directly or indirectly, controls any person liable” under the rules or regulations promulgated under the Act.⁶¹ Section 20(a) was enacted “to prevent people and entities from using straw parties, subsidiaries, or other agents acting on their behalf to accomplish ends that would be forbidden directly by the securities laws.”⁶² Section 20(a) requires only that the secondarily liable party have control over the primary actor, not scienter.⁶³ Thus, Section 20(a) seems like an attractive option for establishing liability without confronting the scienter problem for those who use manipulative machines as straw parties. However, there is one glaring issue—there must be a primary violation performed with scienter to establish secondary control person liability, which leads back to the initial problem of intent.⁶⁴

Moreover, the Supreme Court has foreclosed private causes of action for aiding and abetting primary Section 10(b) violators,⁶⁵ thus closing off any potential common-law avenue of relief that might be tried in order to

60. “[I]t is thought fair that one who benefits from the enterprise and has a right to control the physical activities of those who make the enterprise profitable, should pay for the physical harm resulting from the errors and derelictions of the servants while doing the kind of thing which makes the enterprise successful.” RESTATEMENT (SECOND) OF AGENCY § 161 cmt. a (AM. LAW INST. 1958) (comparing the master-servant and principal-agent relationships and noting that liability for masters and principals is based on a similar public policy).

61. 15 U.S.C. § 78t(a) (2012).

62. *Laperriere v. Vesta Ins. Grp., Inc.*, 526 F.3d 715, 721 (11th Cir. 2008) (citing H.R. REP. NO. 73-152, at 12 (1933)).

63. *See Paracor Fin., Inc. v. Gen. Elec. Capital Corp.*, 96 F.3d 1151, 1161 (9th Cir. 1996) (holding that “the plaintiff need not show the controlling person’s scienter or that they ‘culpably participated’ in the alleged wrongdoing”).

64. *Mizzaro v. Home Depot, Inc.*, 544 F.3d 1230, 1237 (11th Cir. 2008) (“Because a primary violation of the securities laws is an essential element of a § 20(a) derivative claim, we have held that a plaintiff adequately pleads a § 20(a) claim only if the primary violation is adequately pleaded.”).

65. *See Cent. Bank of Denver, N.A. v. First Interstate Bank of Denver, N.A.*, 511 U.S. 164, 191 (1994).

avoid the statutory requirements of Section 20(a).⁶⁶ And while courts have recognized that there may be multiple primary violators with smaller roles in a larger manipulative scheme,⁶⁷ each violator must still perform some “manipulative act” with scienter to be found liable under Section 10(b) and Rule 10b-5.⁶⁸ Thus, any attempts to establish liability under a secondary or multiple actors theory lead back to the scienter problem, and the manipulative A.I. slips through the statutory cracks.

C. OPTIONS BEYOND SECTIONS 10(B) AND 20(A)

Despite the established utility of Sections 10(b) and 20(a), manipulative activity might still be deterred or litigated under different regulations. Given the problems that human-coded automated trading systems have already caused,⁶⁹ the SEC,⁷⁰ Commodity Futures Trading Commission (CFTC),⁷¹ and Financial Industry Regulatory Authority (FINRA)⁷² have taken steps to regulate how market participants develop and use their trading technology. This allows regulators to “combat the new schemes of market manipulation while sidestepping the thorny issue of scienter.”⁷³ For instance, in actions brought under SEC Rule 15c3-5, the SEC need only prove that a defendant failed to “establish risk management controls that are *reasonably* designed

66. This is not to suggest that claims under aiding-and-abetting liability would not ultimately suffer the same fate for the fact that there is no primary violation or because aiding and abetting requires “knowing” provision of assistance. See *Rolf v. Blyth, Eastman Dillon & Co., Inc.*, 570 F.2d 38, 47 (2d Cir. 1978).

67. Section 10(b) makes it unlawful for people to violate the enumerated conduct “directly or indirectly.” 15 U.S.C. § 78j (2012); see *City of Providence v. BATS Glob. Mkts., Inc.*, 878 F.3d 36, 51 (2d Cir. 2017) (“[I]n any complex securities fraud . . . there are likely to be multiple violators, and even an entity that plays a secondary role in a securities fraud case may be held liable as a primary violator.”) (internal quotation marks and citations omitted).

68. See *Cent. Bank of Denver*, 511 U.S. at 191.

69. See, e.g., Lin, *supra* note 56, at 1260–61 (describing the “Flash Crash,” an event in 2010 in which an automated trading system executing orders without respect to price or time limits caused \$1 trillion in market value to vanish from the U.S. stock market in thirty minutes).

70. Merrill Lynch, Pierce, Fenner & Smith Inc., Exchange Act Release No. 78,929, Investment Company Act Release No. 4536, 2016 WL 5358114 (Sept. 26, 2016) (SEC action brought against party that allegedly failed to implement reasonable risk management systems under Rule 15c3-5).

71. Gregory Scopino, *Do Automated Trading Systems Dream of Manipulating the Price of Futures Contracts? Policing Markets for Improper Trading Practices by Algorithmic Robots*, 67 FLA. L. REV. 221, 235 (2015) (discussing how CFTC Rule 166.3 will potentially “be an effective weapon against [algorithmic trading] that disrupts or manipulates derivative markets because . . . decisions analyzing Regulation 166.3 appear to apply a reasonableness standard (as opposed to a scienter requirement)”).

72. In order to minimize risks of algorithmic manipulation, FINRA now requires those “primarily responsible for the design, development or significant modification of algorithmic trading strategies” to register as securities traders. See FINRA, Regulatory Notice 15-06: Registration of Associated Person Who Develop Algorithmic Trading Strategies (Mar. 2015), http://www.finra.org/sites/default/files/notice_doc_file_ref/Notice_Regulatory_15-06.pdf.

73. See Lin, *supra* note 56, at 1303.

to prevent” certain transactions.⁷⁴ While these regulations might prohibit some manipulative conduct, they are insufficient for several reasons.

First, these regulations only cover certain market participants that register with regulators, like broker-dealers, and not all potential manipulators need be registrants.⁷⁵ Normally, Section 10(b) is referred to as a “catch-all” because it applies to any person engaging in manipulative or deceptive conduct, but as discussed in Section II.A, it is not readily applicable to A.I. conduct.⁷⁶ Second, violations of Section 15(c) of the Act do not give rise to private causes of action leaving individuals without a remedy for losses due to manipulation.⁷⁷ Finally, standards based in reasonableness do not go far enough. Just as a black box may not provide any information about its developer’s intent, it may not provide any useful evidence regarding the developer’s reasonableness either.⁷⁸ Therefore, this Note contends that regulations should affirmatively require developers to generate meaningful feedback from their A.I. Consequently, the human response to that feedback—and not the actions or structure of the A.I. itself—will provide the basis for any potential liability analyses.

III. CREATING REGULATORY SOLUTIONS: THREE QUESTIONS

This Part discusses three steps that can be taken to prevent A.I. manipulation and enforce actions against manipulative A.I. traders. First, the Supreme Court must settle a circuit split as to what conduct actually constitutes manipulation. Second, an appropriate regulatory body and framework must be selected for creating A.I. development standards that will fit comfortably with the current Section 10(b) doctrine. Third, the

74. Merrill Lynch, Pierce, Fenner & Smith Inc., 2016 WL 5358114, at *1 (SEC action brought against party that allegedly failed to implement reasonable risk management systems under Rule 15c3-5).

75. While market manipulation actions often involve broker-dealers subject to rules promulgated under Section 15 of the Act, manipulations also often include other parties, like large shareholders, whose conduct might not necessarily be covered under SEC rules. *See* Rajesh K. Aggarwal & Guojun Wu, *Stock Market Manipulations*, 79 J. BUS. 1915, 1937 (2006) (finding that in 142 cases for manipulation brought by the SEC, large shareholders with at least 5% equity ownership were involved in 31.69% of the cases).

76. *See* Yadav, *supra* note 13, at 1043–44; discussion *supra* Section II.A.

77. While no federal cases address causes of action under Rule 15c3-5, courts have found that other rules promulgated under Section 15(c) of the Act do not give rise to a private cause of action, noting that “the Supreme Court ‘presumes that, when a statute contains an enforcement mechanism but does not expressly provide a private remedy, Congress did not mean to permit private enforcement of the statute.’” *See* Harris v. TD Ameritrade, Inc., 805 F.3d 664, 666 (6th Cir. 2015) (quoting Mich. Corr. Org. v. Mich. Dep’t of Corr., 774 F.3d 895, 903 (6th Cir. 2014)).

78. Chagal-Feferkorn, *supra* note 38, at 135 (“[A]utonomous algorithms now make choices that are not programmed by their developers and that are not foreseeable by them, at least with respect to a large subset of all possible scenarios. Therefore, there is no perfect equivalence between programmers’ and algorithms’ choices, and the algorithms’ choices have an independent meaning.”).

regulatory body must inform itself of A.I. industry capabilities to create regulations that are both accurate and efficient in practice.

A. HOW SHOULD WE DEFINE MANIPULATION?

Before regulators can create standards for detecting A.I. manipulation, it will be necessary to define market manipulation for human and A.I. actors alike. Some conduct is classically manipulative, like the wash sales described in the Section I.B hypothetical.⁷⁹ For conduct that does not fit the classic mold, however, the subjective intent requirement has already led circuits to disagree over what human actions constitute manipulation.⁸⁰ As discussed above, A.I. traders will only complicate this problem when they inevitably participate in trading.⁸¹ Therefore, in resolving this split, the Supreme Court should not only consider intent-finding problems in the context of human action, but also in the context of A.I. action. Only then can regulators focus resources on detecting conduct that is actually prohibited, rather than conduct that is acceptable in some circuits but not others.

Currently, circuit courts generally agree that “[t]he gravamen of manipulation is deception of investors into believing that prices at which they purchase and sell securities are determined by the natural interplay of supply and demand, not rigged by manipulators.”⁸² Transactions affecting price are not necessarily manipulative, however, as Congress noted in passing the Act, “if a person is merely trying to acquire a large block of stock for investment, or desires to dispose of a big holding, his knowledge that in doing so he will affect the market price does not make his action unlawful.”⁸³

Beyond these general principles are where circuits split.⁸⁴ The Third Circuit has held that liability for legal transactions executed with manipulative intent also requires a showing that the accused “injected inaccurate information into the market.”⁸⁵ However, in *SEC v. Masri*, a court in the Southern District of New York found that Second Circuit precedent suggested a contrary result.⁸⁶ The acts in question in *Masri*

79. See *supra* Section I.B.

80. See *SEC v. Masri*, 523 F. Supp. 2d 361, 367 (S.D.N.Y. 2007) (discussing the circuit split).

81. Calvello, *supra* note 2 (claiming that “from this point forward, every other asset manager must either defend its current approach to active management or follow BlackRock’s lead,” after BlackRock announced a pivot towards using A.I. in its trading decisions).

82. Although the Third and Second Circuits are divided as to what types of conduct can constitute manipulation, they agree on a general definition of manipulation. See *GFL Advantage Fund, Ltd. v. Colkitt*, 272 F.3d 189, 205 (3d Cir. 2001) (quoting *Gurary v. Winehouse*, 190 F.3d 37, 45 (2d Cir. 1999)).

83. See S. REP. NO. 73-792, at 17 (1934).

84. Compare *Fezzani v. Bear, Stearns & Co. Inc.*, 777 F.3d 566 (2d Cir. 2015), with *GFL Advantage Fund*, 272 F.3d 189.

85. See *GFL Advantage Fund*, 272 F.3d at 205.

86. See *SEC v. Masri*, 523 F. Supp. 2d 361, 371–72 (S.D.N.Y. 2007).

involved “marking the close” or “the practice of repeatedly executing the last transaction of the day in a security in order to affect its closing price.”⁸⁷ The defendant allegedly had an incentive to force a security’s price upward in order to avoid the consequences of a detrimental contract option fixed to the price.⁸⁸ The court held that these otherwise legal trades could constitute manipulation if performed with manipulative intent.⁸⁹ In analyzing the circuit split, the court recognized that the Third Circuit’s “injection of inaccurate information” requirement was likely a result of judicial discomfort with violations based on intent alone.⁹⁰

Despite the inherent problems in evidencing intent, this Note contends that the Third Circuit’s rule requiring manipulative transactions to be accompanied by additional deception to be actionable is unnecessary and ultimately incorrect.⁹¹ To the extent that manipulative transactions alone can mislead investors, they are exactly the type of conduct the Act sought to prohibit.⁹² *Masri* itself acknowledged this principle and recognized that the injection of inaccurate information required by the Third Circuit could essentially be accomplished by manipulative transactions themselves.⁹³ Congress purposefully distinguished between manipulation and deceit in Section 10(b).⁹⁴ Even though manipulation actions might commonly involve further deception, manipulation must represent some distinctly prohibited conduct; otherwise, its inclusion in the statute would be superfluous.⁹⁵

Therefore, the Supreme Court should resolve the current circuit split by adopting the *Masri* holding. Furthermore, the Supreme Court should do so promptly because *Masri* noted that, despite the split, the SEC was nevertheless bringing actions solely for legitimate activities coupled with allegedly manipulative intent.⁹⁶ This divergence between enforcement and

87. *See id.* at 367.

88. *See id.* at 372.

89. *See id.*

90. *See id.* at 367.

91. *See id.* at 371–72.

92. *SEC v. Lek Sec. Corp.*, 276 F. Supp. 3d 49, 59 (S.D.N.Y. 2017) (“[O]ne of the fundamental goals of the federal securities laws is to promote transparency—that is, ‘to prevent practices that impair the function of stock markets in enabling people to buy and sell securities at prices that reflect undistorted (though not necessarily accurate) estimates of the underlying economic value of the securities traded.’”).

93. *See Masri*, 523 F. Supp. 2d at 372 n.17.

94. *See* 15 U.S.C. § 78j (2012) (“It shall be unlawful for any person, directly or indirectly . . . (b) [t]o use or employ, in connection with the purchase or sale of any security . . . any manipulative or deceptive device.”) (emphasis added).

95. *But see* *Schreiber v. Burlington N., Inc.*, 472 U.S. 1, 7–8 (1985) (“Congress used the phrase ‘manipulative or deceptive’ in § 10(b) . . . and we have interpreted ‘manipulative’ in that context to require misrepresentation.”). Given that the Circuits are still currently split, it is not clear whether otherwise legitimate trading activity could create this misrepresentation if performed with manipulative intent.

96. *See Masri*, 523 F. Supp. 2d at 367.

judicial interpretation provides poor notice of what behavior is acceptable,⁹⁷ and regardless of how the Supreme Court resolves the split, market participants are entitled to know where they stand with the SEC.

The Supreme Court should also adopt the remainder of the *Masri* court's holding that "in order to impose liability for an open market transaction, the SEC must prove that *but for* the manipulative intent, the defendant would not have conducted the transaction."⁹⁸ The court reasoned, inter alia, that even if a transaction has the foreseeable effect of increasing or decreasing a security's price in a way that favors a trader, if the trader has some other legitimate investment purpose for transacting, it cannot be said that they are "artificially" affecting the price.⁹⁹ In addition to the fact that this flows logically from the *Masri* holding and reflects legislative intent, the Supreme Court should also adopt this test because it will be conducive to determining the intent of the A.I. traders that will inevitably participate in securities markets.

As discussed, generally asking "why" an A.I. machine made a given decision is impracticable.¹⁰⁰ Instead, the "but for" *Masri* formulation sets parameters for a narrower, hypothetical machine test: if machine inputs that are predictive of a significant price impact are adjusted to instead represent a negligible impact, will the machine still transact? While this test (the *Masri* test) makes it no easier to find human intent, current methods of machine interpretability might actually be able to provide insight into this narrower question for A.I. decisions.

For instance, the *Masri* test might be executed with methods like Local Interpretable Model-Agnostic Explanations (LIME). LIME operates by modifying "a single data sample by tweaking the feature values and observ[ing] the resulting impact on the output,"¹⁰¹ and returning a list of variables that were most determinative of the output. In terms of interpretability, this is often "related to what humans are interested in when observing the output of a model."¹⁰² In the words of LIME's creator, this simplifies the task of generating an explanation because "even really

97. Fox et al., *supra* note 35, at 122 (recognizing that petitioners to the Supreme Court have accurately stated that "the current split among the federal circuits creates significant confusion for market participants and disparities of outcome for those accused").

98. See *Masri*, 523 F. Supp. 2d at 372.

99. *Id.* at 373. For instance, in *Masri*, the court noted that increased trading before market close could simply be due to the fact that "traders monitor activity and their positions throughout the day before conducting their trades." *Id.* at 370.

100. See Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a 'Right to an Explanation' is Probably Not the Remedy You Are Looking For*, 16 DUKE L. & TECH. REV. 18, 65–67 (2017) (discussing the "transparency fallacy" and why certain information that pertains to how a machine makes a decision is actually unlikely to be useful for a person seeking redress for an improper machine decision).

101. Lars Hulstaert, *Understanding model predictions with LIME*, MEDIUM (July 11, 2018), <https://towardsdatascience.com/understanding-model-predictions-with-lime-a582fdff3a3b>.

102. *Id.*

complicated models are not that complicated in a fixed neighborhood,” or the space surrounding a single data sample.¹⁰³ In the case of market manipulation, this would involve a focus on the variables related to the anticipated change in market price.¹⁰⁴ Additionally, insofar as LIME only perturbs inputs to observe outputs, it is “model-agnostic,” or pliable to a range of machine structures.¹⁰⁵ Additionally, LIME requires no access to training data to return an explanation.¹⁰⁶

Others might recognize that the *Masri* test could potentially be answered with “closest-world” counterfactual questions, which ask “how the world would have to be different for a desirable outcome to occur.”¹⁰⁷ Academics have noted how answers to “closest-world” counterfactual questions can be “efficiently and effectively computed by applying standard techniques, even to cutting-edge architectures” in A.I. machines.¹⁰⁸ Rather than generating a list of variables ranked by importance like LIME, this method returns a hypothetical scenario in which the slightest change of inputs would yield a “closest possible world” where the outcome is different.¹⁰⁹ Once again, a result altering the variables affecting market price impact would be relevant to a manipulation determination. Moreover, like LIME, this method does not require an understanding of the machine’s innerworkings.¹¹⁰

These methods may not yield perfect results, but they demonstrate that logical formulations of legal standards, like the *Masri* test, might provide more meaningful information about A.I. decisions. In any event, they go

103. Kyle Polich, *Trusting Machine Learning Models with LIME*, DATA SKEPTIC, <https://dataskeptic.com/blog/transcripts/2016/trusting-machine-learning-models-with-lime> (podcast and transcript containing interview with Marco Riberio, LIME creator).

104. This Note assumes that market conditions and the anticipated trade characteristics can be used to calculate impact on market price with some accuracy. This is not considered to be unrealistic, as the National Association of Securities Dealers (NASD) and the New York Stock Exchange (NYSE) have required investment advisers to consider the market impact of a given transaction as part of their duty of “best execution.” See *Kurz v. Fid. Mgmt. & Research Co.*, No. 07–CV–709–JPG, 2008 WL 2397582, at *2 (S.D. Ill. June 10, 2008) (explaining factors relevant to the duty of best execution).

105. Hulstaert, *supra* note 101.

106. Polich, *supra* note 103.

107. Consider a machine that makes a decision based on a number of input factors. Counterfactual questions essentially “ask” the machine what change in inputs would have yielded a different decision. While drastically changing all of the inputs might yield a different result, as an answer this does not signal much about how the machine “thinks.” By instead seeking out the “minimal amount of information capable of altering a decision,” or the “closest-world” in which a different decision would be made, it is possible to see which factors were most critical to a decision. See Sandra Wachtera, Brent Mittelstadt & Chris Russell, *Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR*, 31 HARV. J.L. & TECH. 841, 844–45 (2018).

108. At the cost of receiving overly restrictive explanations, proponents of the “closest world” counterfactual explanations offer their method as an alternative to methods like LIME, claiming that their counterfactuals provide more interpretable results. See *id.* at 852.

109. *Id.* at 844.

110. *Id.* at 860.

beyond asking what the A.I. intended to do, generally, and avoid an unnecessary and burdensome investigation into the A.I.'s training history and architecture.¹¹¹ This Note suggests that in fields where A.I. is replacing or supplementing human decision making, courts should shape human-centric inquiries into more logical formulations. The Supreme Court can do this in the manipulation context by adopting the *Masri* test and take the first step towards transforming what has been a hypothetical question of intent for humans into a demonstrable matter of fact for future A.I. actors. This will provide greater clarity in determining if machine conduct is manipulative, whether by diligent developers during use or by factfinders examining conduct in a courtroom after the fact. Additionally, regulators can use these capabilities as the basis for practical expectations in industry regulation.

B. WHAT REGULATORY FRAMEWORK SHOULD BE USED TO OVERSEE A.I. TRADING?

While the *Masri* test might be useful for determining machine intent, A.I. developers will still have no specific obligations to *Masri* test their machines without additional regulations. Ideally, Congress could both codify the *Masri* definition and establish causes of action against developers of manipulative A.I. independent of Section 10(b). Alternatively, the SEC could enact the same A.I. standards as rules and enforce them against certain parties based on existing statutory grants of power from Congress.¹¹²

Instead, this Note considers how industry standards for A.I. imposed by Self-Regulatory Organizations (SROs)¹¹³ or other industry authorities could deter manipulation while also supporting enforcement of Section 10(b) actions.¹¹⁴ These standards would not replace Section 10(b), but they would provide additional affirmative obligations for market participants and generate a paper trail to provide more relevant facts for potential Section

111. See Edwards & Veale, *supra* note 100, at 66–67 (discussing the “transparency fallacy” and why certain information that pertains to “how” a machine makes a decision is actually unlikely to be useful for a person seeking redress for an improper machine decision).

112. The SEC already requires broker-dealers to establish systems to manage financial and regulatory risks which, in effect, provide a mechanism to stop some manipulative behavior. See Yadav, *supra* note 13, at 1057–58; 17 C.F.R. § 240.15c3-5 (2018).

113. SROs include organizations like FINRA, stock exchanges, and clearing houses. See *Self-Regulatory Organization – SRO Definition*, INVESTOPEDIA, <https://www.investopedia.com/terms/s/sro.asp> (last visited Sept. 21, 2019).

114. For example, GAAP violations can be used to establish an inference of scienter when other “red flags” are ignored. See *In re Oxford Health Plans, Inc. Sec. Litig.*, 51 F. Supp. 2d 290, 295–96 (S.D.N.Y. 1999). Courts have also held that violation of SRO rules can provide relevant evidence for certain Section 10(b) actions. See *GMS Group, LLC v. Benderson*, 326 F.3d 75, 82 (2d Cir. 2003) (finding that violations of NASD (the predecessor to FINRA) may be considered relevant for purposes of Section 10(b) unsuitability claims).

10(b) manipulation claims. The analysis of industry, as opposed to governmental, regulation is the focus of this Note for multiple reasons.

First, Section 10(b) has been the basis for manipulation actions for the past eighty years, so judges are likely dependent on the Section 10(b) doctrine in the manipulation context.¹¹⁵ Relatedly, Congress is unlikely to create legislation for a problem that is currently speculative in nature.

Second, the problem presented by A.I. traders is evidentiary and not so distinct from the evidentiary problem of proving human intent as to require an entirely new statute. Indeed, to the extent that machine actions *could* be explained, there would be no reason to cast aside the large body of instructive Section 10(b) doctrine just because a computer made the decisions in question.¹¹⁶ Moreover, for the same reason that the longstanding body of Section 10(b) case law provides some guide for legality, new legislation might inversely suffer from judicial uncertainty when the first actions are brought, despite whatever contemporary and incisive language drafters would use in a statute.¹¹⁷

Third, there is no single market for securities, and trading involves many intermediaries. Thus, commentators have recognized that some SROs, like exchanges, are advantageously situated to regulate the transactions they facilitate.¹¹⁸ Also, without a wider grant of authority from Congress, the SEC may not be able to impose affirmative obligations on all would-be manipulators. Therefore, industry standards might regulate a wider range of potential manipulators, while still complementing Section 10(b)'s function as a "catch-all" provision. This assumes that Section 10(b) plaintiffs may put forth evidence of a defendant's noncompliance with certain industry standards in order to help establish scienter.¹¹⁹

Fourth, regulating via SROs is not only promising in theory but it is also realistic, as SROs have a personal incentive in stopping manipulation. Plaintiffs have had recent success in actions against exchanges for facilitating alleged market manipulations,¹²⁰ and SROs seeking to avoid a similar fate could benefit by proactively creating their own regulations.

115. See Fox et al., *supra* note 35, at 123 ("[T]he common law jurisprudence of securities manipulation has overwhelmingly developed around Rule 10b-5. While the ambit of § 9(a)(2) is now wide, path dependency seems to have resulted in § 10(b), rather than § 9(a)(2), continuing to be the regulators' and private litigants' statute of choice.")

116. See, e.g., *In re Facebook, Inc., IPO Sec. & Derivative Litig.*, 986 F. Supp. 2d 428 (S.D.N.Y. 2013) (finding that the failure of an automated trading platform created liability for the exchange that developed the platform, when the exchange ignored testing that suggested the automated system was not prepared for the order traffic that the exchange was expecting).

117. See, e.g., Lin, *supra* note 56, at 1301 (discussing how regulators using new grants of authority "will be lacking in meaningful precedent in the near term, particularly on the issue of scienter").

118. See, e.g., *id.* at 1304.

119. See cases cited *supra* note 114.

120. See *City of Providence v. BATS Glob. Mkts., Inc.*, 878 F.3d 36, 51 (2d Cir. 2017) (holding that plaintiff-investors properly pled claims for market manipulation against the defendant stock

Lastly, whether the standards discussed herein are codified in a statute or regulated via industry standards, plaintiffs would present the same evidence in each case. The critical difference would be that a plaintiff bringing a Section 10(b) action would highlight a defendant's departure from the industry standards to prove intent or recklessness, while a new statute might allow for determinations under a negligence standard.¹²¹ In other words, plaintiffs who could prove scienter in a Section 10(b) case would have no difficulty proving negligence under customized statutes. Therefore, considering how scienter issues would be resolved in a Section 10(b) action provides a more comprehensive analysis.

Considering the potential harm that unregulated A.I. could cause in certain contexts, some have suggested that an agency akin to the Food and Drug Administration (FDA) should be charged with granting pre-market approval for algorithms.¹²² Yet, in the context of A.I. traders, pre-market approval would provide little value. Unlike drugs that are tested for predictability, A.I.'s inherent utility comes from the fact that it makes unpredictable decisions that humans cannot anticipate.¹²³ A.I. algorithms also evolve as they trade, which could quickly render prior approval meaningless.¹²⁴ Moreover, many machines might not ever encounter opportunities for profitable manipulation.¹²⁵ Thus, the benefits of screening each A.I. trading system would be far outweighed by the administrative burden of a process mandating affirmative, pre-market approval.

The FDA suggestion is nevertheless insightful for recognizing that regulators should be experts with an acute understanding of the industry. With respect to government regulation, however, a new agency would not be necessary for regulating A.I. traders, since the SEC would already have the expertise and capability to do so. Moreover, this Note contends that

exchange, in alleging that the stock exchange had committed a manipulative act by permitting high-frequency traders to manipulate prices via the structure of the exchange).

121. See, e.g., Merrill Lynch, Pierce, Fenner & Smith Inc., Exchange Act Release No. 78,929, Investment Advisors Act No. 4536, 2016 WL 5358114 (Sept. 26, 2016) (bringing an enforcement action against a defendant for failure to implement *reasonable* safeguards in accordance with Rule 15c3-5).

122. See generally Andrew Tutt, *An FDA for Algorithms*, 69 ADMIN. L. REV. 83 (2017).

123. See Bathae, *supra* note 4, at 907 (“[W]e may be able to tell what the AI’s overarching goal was, but black-box AI may do things in ways the creators of the AI may not understand or be able to predict.”).

124. As an A.I. machine encounters more data, it internally adjusts its own algorithm in order to make more accurate categorizations or predictions about future data. See McClelland, *supra* note 1.

125. See Daniel R. Fischel & David J. Ross, *Should the Law Prohibit “Manipulation” in Financial Markets?*, 105 HARV. L. REV. 503, 512–19 (1991) (arguing that in practice, there is a low probability that solely trade-based manipulations will ever be successful or profitable); Fox et al., *supra* note 35, at 77 (discussing studies showing the existence of certain types of manipulative behavior but commenting that overall there is little data showing how often market manipulation occurs).

non-governmental entities are even better suited to the task given their own expertise, incentives, and positions in the marketplace.

As shown by the impact of GAAP¹²⁶ in financial reporting, standards established by independent non-governmental organizations can substantially affect the way companies exercise diligence.¹²⁷ Additionally, these regulatory benefits accrue without the administrative burden associated with an affirmative regulatory approval process.¹²⁸ And while standards like GAAP do not have the force of law, they still incentivize diligence, since non-compliance can cause negative business repercussions¹²⁹ and can even be used as evidence in Section 10(b) actions.¹³⁰

To the extent that an industry standard for generating A.I. explanations could be effective, it could realize these same benefits that GAAP does in the accounting context. While A.I. does present a challenge in that different enterprises will use different systems to meet different trading objectives, an expert body could feasibly stay informed as to the currently available methods for A.I. explanations, like LIME or “closest world” counterfactuals. An SRO or industry authority could then require developers to devise their own specialized methods for performing the *Masri* test on their machines. GAAP shows similar flexibility can be practicable; often, there are a range of reasonable choices in financial reporting, and what might be improper in one instance might be acceptable in another.¹³¹ Moreover, while this Note does not specifically address other problems financial institutions might encounter when using A.I.,

126. “GAAP” is an acronym for “generally accepted accounting principles.” GAAP is issued by the Financial Accounting Standards Board and refers to the set of standards companies must comply with when compiling financial statements. *Generally Accepted Accounting Principles (GAAP)*, INVESTOPEDIA, <https://www.investopedia.com/terms/g/gaap.asp> (last visited Sept. 21, 2019).

127. See Sean Ross, *Who Enforces GAAP?*, INVESTOPEDIA, <https://www.investopedia.com/ask/answers/020915/who-enforces-gaap.asp> (last updated Nov. 21, 2018) (“[M]ost companies follow GAAP as though they were law. This is one of the chief examples of private businesses regulating themselves to help promote credibility within an industry.”).

128. When corporations file SEC mandated financial disclosures, the SEC does not give affirmative approval to filings. *Division of Corporation Finance: Filing Review Process*, U.S. SEC. & EXCH. COMM’N. (Jan. 29, 2019), <https://www.sec.gov/divisions/corpfin/cffilingreview.htm> (“The Division’s review process is not a guarantee that the disclosure is complete and accurate—responsibility for complete and accurate disclosure lies with the company and others involved in the preparation of a company’s filings.”).

129. See Ross, *supra* note 127 (“Failure to [comply with GAAP] could violate lenders’ agreements, cause stock prices to drop or ruin business deals. These auditing requirements create useful leverage for the FASB and GAAP.”).

130. See *In re Oxford Health Plans, Inc. Sec. Litig.*, 51 F. Supp. 2d 290, 295–96 (S.D.N.Y. 1999).

131. *Thor Power Tool Co. v. Comm’r*, 439 U.S. 522, 544 (1979) (“‘Generally accepted accounting principles’ . . . tolerate a range of ‘reasonable’ treatments, leaving the choice among alternatives to management.”).

independent organizations would be well-situated to address other A.I. issues in securities markets as they arise.¹³²

C. WHAT RULES SHOULD REGULATORS ENACT TO DETER AND DETECT A.I. MANIPULATION?

At a high level, industry standards would define both the “how” and the “when” for generating A.I. explanations. The “how” refers to defining parameters for how explanations should be generated, so that the results are accurate. The “when” refers to defining when explanations should be generated, so that testing is efficient and not overly burdensome.¹³³ In actions for seemingly manipulative conduct, there would then be clearer evidence to guide judicial inquiries.¹³⁴ Instead of asking how A.I. made a decision, the inquiry would focus on whether a developer recklessly disregarded industry standards for how and when tests should be performed, giving rise to a manipulation risk.¹³⁵

Successfully defining the “how” and the “when” for *Masri* tests will depend on regulators’ ability to meet two respective challenges. First, regulators will need to keep pace with the current state of A.I. and explanation methods to ensure that *Masri* tests accurately predict market impact. This is because there is no single standard for “interpretability,” especially across different A.I. models,¹³⁶ and different interpretability standards might be warranted as A.I. continues to develop. Second, so as not to burden useful and lawful implementations of A.I., regulators will

132. For instance, others have hypothesized about A.I. that can not only trade but can also communicate information on a person or entity’s behalf, creating the potential for fraudulent misstatements. See Bathaee, *supra* note 4, at 912. Were this to become a common A.I. implementation, SROs or an industry authority would already be informed as to industry capabilities for testing and monitoring A.I. and would be well situated to address these types of problems.

133. This would conceivably take into account some of the factors that would be used by investment advisors in calculating the market impact of their trades. See *Kurz v. Fid. Mgmt. & Research Co.*, No. 07–CV–709–JPG, 2008 WL 2397582, at *2 (S.D. Ill. June 10, 2008) (explaining that the duty of best execution involves analyzing how a trade will impact a market).

134. For instance, when a defendant fails to follow an accepted industry standard, a court might ask whether they ignored other “red flags.” See *Oxford Health Plans*, 51 F. Supp. 2d at 295–96 (S.D.N.Y. 1999).

135. This assumes that manipulation actions could be brought by alleging that A.I. developers were reckless. See discussion *infra* Section IV.A. In a recklessness analysis, the industry standard would form the basis for the duty to monitor and would provide a benchmark for “ordinary care.” See *Novak v. Kasaks*, 216 F.3d 300, 308 (2d Cir. 2000) (defining recklessness as “conduct which is highly unreasonable and which represents an extreme departure from the standards of ordinary care . . . to the extent that the danger was either known to the defendant or so obvious that the defendant must have been aware of it”) (internal quotation marks omitted) (quoting *Rolf v. Blyth, Eastman Dillon & Co., Inc.*, 570 F.2d 38, 47 (2d Cir. 1978)).

136. See FINALE DOSHI-VELEZ & BEEN KIM, TOWARDS A RIGOROUS SCIENCE OF INTERPRETABLE MACHINE LEARNING 1 (Mar. 2, 2017), <https://arxiv.org/pdf/1702.08608.pdf> (“Unfortunately, there is little consensus on what interpretability in machine learning *is* and how to *evaluate* it for benchmarking.”).

need to define the scenarios when manipulation is truly a risk, and only invoke *Masri* testing requirements in those instances.

Regarding the first regulatory challenge, defining parameters for accurate *Masri* tests to differentiate between manipulative and legitimate conduct will require regulators to apprise themselves of the limitations of interpretability methods for A.I. decision-making. Specifically, since there is no single independent variable representing the impact of a transaction on a security's price, market impact can only be roughly anticipated by analyzing multiple input variables.¹³⁷ But when market impact variables also bear on legitimate investment decisions, there is a risk that *Masri* test results might be obscured.¹³⁸

The variable control problem is effectively illustrated in a non-financial example. Developers created an A.I. model that calculated risk of hospital readmission for pneumonia patients based on various factors.¹³⁹ The developers naturally expected that patients with asthma would be more likely to be readmitted, but developers were surprised to find that the model associated asthma with a lower risk of readmission.¹⁴⁰ Upon further inspection the developers realized that they did not account for the fact that asthma patients received better care when first admitted, and thus asthma became a statistical proxy for a factor that correlated with a lowered risk of readmission.¹⁴¹

This problem demonstrates that A.I. does not necessarily map for human ideas of causation; rather, it finds correlative patterns in data.¹⁴² When “asking” a machine whether it would have transacted but for an anticipated change in price, it will be necessary to permute variables associated with market impact.¹⁴³ But, if the machine had a legitimate

137. See J.P. BOUCHAUD, PRICE IMPACT 1 (Aug. 24, 2017), <https://arxiv.org/pdf/0903.2428.pdf> (“The most important questions are related to the volume dependence of impact (do larger trades impact prices more?), and the temporal behaviour of impact (is the impact of a trade immediate and permanent, or is there some lag dependence of the impact?).”).

138. For example, in the context of human trading behavior, courts have recognized that the timing of trades closer to the end of the day could be evidence of intent to affect closing price, or alternatively, could be related to traders closing out their legitimate positions after a day of careful observations. See *SEC v. Masri*, 523 F. Supp. 2d 361, 369–70 (S.D.N.Y. 2007).

139. Andrew Slavin Ross, Michael C. Hughes & Finale Doshi-Velez, *Right for the Right Reasons: Training Differentiable Models by Constraining Their Explanations*, in PROCEEDINGS OF THE TWENTY-SIXTH INT’L JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE 2662, 2662 (2017), <https://www.ijcai.org/proceedings/2017/0371.pdf>.

140. *Id.*

141. *Id.*

142. See Edwards & Veale, *supra* note 100, at 59 (“[I]n some systems there is no theory correlating input variables to things humans understand as causal or even as ‘things.’”).

143. For instance, a LIME explanation would minimize variables related to price impact to determine if those variables were determinative of the original decision in the real data case. See Polich, *supra* note 103.

investment purpose, it might have also depended on these same variables.¹⁴⁴ Thus, even if the machine returns an alternate answer after the variables are changed in a *Masri* test, it still cannot absolutely be said that the price impact was the but for cause.

These issues are likely manageable. Indeed, to the extent that a developer allows an A.I. machine to trade autonomously, the developer effectively wagers that the machine can make accurate predictions about organic market activity. Thus, it is not thought to be unreasonable to require A.I. to occasionally make predictions about the market impact of its *own* trades with a similar degree of accuracy. For developers who nevertheless claim that computing accurate *Masri* tests is too difficult, there might be alternatives. For instance, human approval for A.I. transactions beyond a certain dollar or share threshold might be required.¹⁴⁵ Alternatively, others have suggested that A.I. developers should periodically log versions of their machine.¹⁴⁶ In this scenario, an interested plaintiff might later retrieve a prior machine version during discovery and return a machine explanation that a defendant refused to provide.

Even if regulators can establish guidelines for accurate *Masri* tests, they will still face the second regulatory challenge: defining instances when *Masri* tests will be required that are not overly burdensome on lawful A.I. operation. For instance, a blunt regulation might require A.I. traders to return an explanation before every transaction to minimize the risk of manipulation. This, however, would effectively increase the time required to execute a transaction. If the A.I. is employed in a capacity where the speed of its computations is critical, as is already the case for many automated trading systems,¹⁴⁷ this delay might render using the A.I. inefficient.

Regulations might also require explanations for each training transaction to discourage manipulative behavior before there are any real-time financial stakes. Using current methods, this might require creating and storing manipulation test data for each real data case.¹⁴⁸ Note, however, that A.I. is normally employed to make use of incomprehensibly large sets

144. In the context of human trading behavior, courts have recognized that the timing of trades at the end of the day could be evidence of intent to affect price, or alternatively, could be related to traders closing out their legitimate positions after a day of careful observations. *See* SEC v. Masri, 523 F. Supp. 2d 361, 369-72 (S.D.N.Y. 2007).

145. *See* Chagal-Feferkorn, *supra* note 38, at 134 (“Unless the ‘person in the loop’ needs to authorize each and every decision by the algorithm, the algorithm would reach conclusions based on new information that the human programmer would not have had a chance to consider. Granted, the programmer can certainly set boundaries that the algorithm may never cross”) (footnote omitted).

146. Wachtera et al., *supra* note 107, at 882 n.213.

147. *See* Tara E. Levens, *Too Fast, Too Frequent? High-Frequency Trading and Securities Class Actions*, 82 U. CHI. L. REV. 1511, 1529 (2015).

148. For instance, LIME requires creating a data case that slightly alters inputs of a real data case to generate explanations. *See* Polich, *supra* note 103.

of data¹⁴⁹ and can require large sets of data for training.¹⁵⁰ Requiring developers to generate even larger datasets for manipulation training could simply burden them with too much information to manage. This might be especially burdensome considering that there are only limited instances in which manipulation will be viable.¹⁵¹ Ultimately, tests would likely be required during training and real-time use, but only upon triggering of some trade volume threshold, and regulators' expertise will be critical in defining the specifics of these requirements.

There are no perfect solutions to these regulatory challenges. Rather, these challenges demonstrate that regulation will involve a balance between both allowing lawful A.I. users to benefit from their innovation and requiring A.I. users to generate accurate explanations in a limited set of suspect circumstances. Developing regulations with "model-agnosticism" in mind can help regulators retain flexibility and allow for standards that can be applied to different types of A.I. machines across the industry.¹⁵² Then the burden will rightfully rest on the developers, who know their machines best, to implement safeguards against manipulation. Like accountants following GAAP, developers may select one of many reasonable methods, in the shadow of litigation that looms over them if they are reckless in their methodology.¹⁵³

In enacting these regulations, regulators may have to draw bright lines so that *Masri* tests result in quantitatively meaningful results. This is vulnerable to classic criticisms that bright line definitions of manipulation will either be too restrictive and deter legitimate conduct, or will fail to capture all undesirable conduct due to the amorphous nature of manipulation.¹⁵⁴ These criticisms, however, do not consider the fact that A.I. capabilities can be channeled to meet regulatory obligations in ways that would not be practicable for human action.

Regarding the first criticism, defining transactions that pose a manipulation risk will not necessarily prohibit those transactions. Rather, it will merely require A.I. developers to return an explanation of the *Masri* test from their machine. Therefore, if the machine was not in fact transacting to manipulate, this will be shown in an accurate test.

149. See Crawford, *supra* note 8.

150. See Hardesty, *supra* note 9.

151. Fischel & Ross, *supra* note 125, at 512–19 (discussing that, in practice, there is a low probability that solely trade-based manipulations will ever be successful or profitable).

152. Hulstaert, *supra* note 101.

153. See *In re Oxford Health Plans, Inc. Sec. Litig.*, 51 F. Supp. 2d 290, 295–96 (S.D.N.Y. 1999) (holding that auditors' non-compliance with accepted accounting standards and ignorance of other "red flags" created a strong inference that they acted recklessly).

154. Fischel & Ross, *supra* note 125, at 523 ("Even a seemingly narrow legal rule focusing on 'clear' evidence of manipulative intent (assuming such 'clear' evidence exists) is unlikely to provide net social benefits.").

Additionally, developers will be allowed to customize *Masri* tests for their own machines, foreclosing arguments that regulations are unworkable.

This leads to the second criticism that allowing developers to set their own test parameters based on arbitrary limits will allow some manipulative conduct to go undetected. Indeed, arbitrary limits might not capture “the full range of ingenious devices that might be used to manipulate securities prices.”¹⁵⁵ But regulations will still deter manipulation for two reasons. First, if evidence can be put forth that certain conduct was in fact manipulative, merely technical or formalistic compliance with an industry standard will not excuse a developer from liability.¹⁵⁶ Second, any manipulative conduct that falls beyond the prohibitions of bright-line rules will be riskier for potential manipulators. Critics of market manipulation regulation already acknowledge that manipulative conduct is “self-detering” given the rarity of instances where it is profitable and the risk of lost investment if other market forces foil the manipulation.¹⁵⁷ Naturally, pouring more capital into the manipulation can lessen the impact of any disruptions, and therefore increase the likelihood of success for the manipulation, via a “price pressure” effect.¹⁵⁸ So while requiring explanations for certain transactions might not detect all potentially manipulative behavior, it can provide a deterrent in the range of manipulations that are the safest bets. And while A.I. traders might still be *capable* of some riskier manipulations, they will necessarily have some pre-programmed risk appetite,¹⁵⁹ and they may be prevented from doing so. Other provisions in the Act like Section 16(b) operate similarly; conduct that on its face may either be nefarious or innocent cannot be prohibited

155. *Santa Fe Indus., Inc. v. Green*, 430 U.S. 462, 476–77 (1977).

156. To continue the analogy to GAAP, courts have recognized compliance with GAAP in financial reporting does not excuse companies that nevertheless prepare misleading financial statements. See *U.S. v. Ebbers*, 458 F.3d 110, 126 (2d Cir. 2006).

157. Fischel & Ross, *supra* note 125, at 512–19 (discussing that in practice, there is a low probability that trade-based manipulations will ever be successful or profitable).

158. While the existence of the effect is debated, even those skeptical of the effect recognize that to the extent it exists, “securities have supply and demand elasticities no smaller in magnitude than 1. This means that a prospective manipulator must purchase at least one percent of the firm’s outstanding shares to cause a one percent change in price. Such a large capital outlay exposes the would-be manipulator to tremendous risk” Fischel & Ross, *supra* note 125, at 518–19 (footnote omitted). Such risks support the point that *Masri* testing for these types of transactions would provide an additional deterrent effect, as it would force would-be manipulators into even riskier transactions to avoid detection.

159. Even though an A.I. trader will constantly make predictions, developers are likely to only allow the machine to trade on predictions that are supported by some level of confidence. See Kuttruf, *supra* note 24 (discussing that developers will have to manually input parameters like “a threshold for the model confidence of a given prediction” in order to guide how a machine acts on its market predictions).

outright, so a bright-line rule is drawn to prohibit the conduct in instances where it is more likely to be nefarious.¹⁶⁰

IV. JUDICIAL IMPLEMENTATION OF REGULATIONS

As discussed, this Note only considers how an industry standard would regulate conduct against the backdrop of the current Section 10(b) doctrine. This is not necessarily a negative limitation. Judges have traditionally retained flexibility in meeting the legislative purposes of the Act as the nature of fraudulent activity has evolved in the past eighty years.¹⁶¹ In this tradition of flexibility, modern realities will require courts to make two legal realizations to account for the fact that the Section 10(b) doctrine has yet to encounter A.I.¹⁶²: (1) reckless conduct can satisfy scienter in appropriate A.I. manipulation cases; and (2) deviations from industry standards for A.I. development can provide circumstantial evidence of recklessness.¹⁶³

A. RECKLESSNESS MAY BE A SUFFICIENT SCIENTER IN MARKET MANIPULATION CASES

As discussed, the importance of scienter in market manipulation actions has led some to conclude that reckless, as opposed to intentional conduct, cannot create liability.¹⁶⁴ However, it is worthwhile to consider whether courts have rejected recklessness in manipulation cases, or whether it is just an unlikely allegation given the typical manipulation fact patterns.¹⁶⁵

160. Section 16(b) of the Act permits recovery of any combination of purchases and sales of equity securities within any six-month period by any officer, director, or 10% shareholder of a company with securities registered under the Act. *See Credit Suisse Sec. (USA) LLC v. Simmonds*, 566 U.S. 221, 223 (2012). Thus, a temporally extended form of insider trading is not expressly prohibited, but insiders will incur the risk of forfeiting gains within the six-month period if they choose to act on inside information.

161. While not directly addressing market manipulation, the Supreme Court has acknowledged that “[w]hen we deal with private actions under Rule 10b-5, we deal with a judicial oak which has grown from little more than a legislative acorn” and that with respect to Rule 10b-5, the court will take into account policy considerations when “neither the congressional enactment nor the administrative regulations offer conclusive guidance.” *Blue Chip Stamps v. Manor Drug Stores*, 421 U.S. 723, 737 (1975).

162. *See* John C. Coffee, Jr., *Introduction: Mapping the Future of Insider Trading Law: Of Boundaries, Gaps, and Strategies*, 2013 COLUM. BUS. L. REV. 281, 317 (“Rule 10b-5 was intended to evolve to keep pace with the ingenuity of fraudsters.”).

163. Despite the opacity and autonomy of any given A.I. machine, courts should feel comfortable judging the conduct of the developers when they are in control of the data being fed into a machine and monitoring its operation. *See Crawford, supra* note 8 (“The great reveal about [machine learning algorithms] is that they aren’t all that smart—they’re basically just feeling around, through trial and error, to try and find the relationships in your data.”).

164. *See Yeres et al., supra* note 59.

165. *See, e.g., SEC v. Lek Sec. Corp.*, 276 F. Supp. 3d 49, 60 (S.D.N.Y. 2017) (discussing the sufficiency of recklessness to establish scienter in Section 10(b) actions, in an action brought specifically for market manipulation).

For instance, *ATSI* has been cited for the proposition that recklessness cannot satisfy market manipulation scienter¹⁶⁶ because of the critical nature of intent.¹⁶⁷ There are three problems with this contention. First, recklessness was not alleged in *ATSI*—intent was; and the court simply found that the evidence did not support the allegation. Second, before discussing the importance of scienter, the court noted at the outset of its analysis that plaintiffs can allege scienter with “strong circumstantial evidence of conscious misbehavior or recklessness,”¹⁶⁸ for all Section 10(b) actions. Third, the Second Circuit has elsewhere stated that “recklessness is a sufficiently culpable mental state in the securities fraud context,”¹⁶⁹ and has found that plaintiffs not alleging specific intent may still prove “conscious recklessness—i.e., a state of mind *approximating actual intent*”¹⁷⁰ Therefore, while recklessness might require a stronger evidentiary showing,¹⁷¹ courts have recognized that it is merely an alternative means to satisfying Section 10(b) scienter and not an insufficient form of culpability.

For the same reason that the current manipulation doctrine provides no home for A.I., it is also historically unlikely that plaintiffs would need to allege recklessness in manipulation cases. In instances where manipulation occurs because of a defendant’s failure to monitor another, control-person liability, which requires no proof of scienter,¹⁷² is a more favorable option. Additionally, while misstatements and omissions often involve avoidance or ignorance of information, for which recklessness provides a better analytical framework, manipulation has up until now always required affirmative acts.¹⁷³ For these reasons, any existing disparity between the frequency of successful recklessness allegations in misstatement and omissions cases compared to manipulation cases is likely a product of this practical dynamic. Looking to the statute itself, Section 10(b) discusses

166. See Yeres et al., *supra* note 59, at 7.

167. See *ATSI Commc’ns, Inc. v. Shaar Fund, Ltd.*, 493 F.3d 87, 102 (2d Cir. 2007).

168. See *id.* at 99.

169. *Teamsters Local 445 Freight Div. Pension Fund v. Dynex Capital Inc.*, 531 F.3d 190, 194 (2d Cir. 2008).

170. *S. Cherry St., LLC v. Hennessie Grp. LLC*, 573 F.3d 98, 109 (2d Cir. 2009).

171. See *ECA v. JP Morgan Chase Co.*, 553 F.3d 187, 199 (2d Cir. 2009) (recognizing that since recklessness does not involve a specific motive by the defendant, the “strength of the circumstantial allegations must be correspondingly greater” to prove scienter).

172. See *Paracor Fin., Inc. v. Gen. Elec. Capital Corp.*, 96 F.3d 1151, 1161 (9th Cir. 1996) (“The plaintiff need not show the controlling person’s scienter or that they ‘culpably participated’ in the alleged wrongdoing.”).

173. See Fox et al., *supra* note 35, at 124 (“[A] manipulation claim under Rule 10b-5, even when it involves misrepresentations, also requires ‘manipulative acts.’”) (citing *ATSI Commc’ns*, 493 F.3d 87); Fischel & Ross, *supra* note 125, at 512 (“[T]he analogy [between manipulation and fraudulent nondisclosure] breaks down upon closer analysis. Nondisclosure in law is actionable only when there is a duty to speak arising from objectively observable criteria, such as the relationship between the parties.”).

deceit and manipulation together, and they are similarly prohibited.¹⁷⁴ There is no indication that Congress intended to impose different culpability requirements for a misstatement as opposed to a manipulative act.

Even if recklessness has historically been an unlikely candidate for market manipulation, the existing recklessness doctrine for misrepresentations will apply just as well to reckless developers of A.I. This is because recklessness can be proven by showing that defendants “failed to review or check information that they had duty to monitor or ignored obvious signs of fraud,”¹⁷⁵ and industry standards could conceivably outline this duty for manipulation.

Two examples discussed below involve instances where courts have found liability under Section 10(b) for deceit in instances where defendants failed to monitor certain activity. Liability in the first case involved a failure of modern technology, and liability in the second turned only on the actions and inaction of human actors. Together, these cases demonstrate the flexibility of the recklessness doctrine and why it will provide a good fit for reckless use of manipulative A.I.

First, in *Facebook*, a stock exchange’s failure to properly test an automated stock trading platform constituted recklessness when the exchange anticipated a high trade volume and touted the platform’s potential, but the system ultimately failed to execute trades in a timely manner.¹⁷⁶ While nothing in the *Facebook* record suggests that the system was autonomous or a black box, details pertaining to *how* the system failed were not relevant to the legal determination. Simply stated, the exchange developed an automated system for its own benefit and understood that its failure could adversely affect traders on the exchange yet disregarded that risk to deploy an inadequate system anyways. While the recklessly disregarded risk in *Facebook* was of misrepresentation to traders, this Note posits that an analogous fact-pattern will eventually emerge where the risk is of A.I. manipulation instead. At that point, it will again be important to ask whether there is a principled reason for finding that the former case can create liability, while the latter cannot. Because once this artificial distinction is cast aside, the *Facebook* variety of recklessness will provide a useful framework for holding A.I. developers accountable.

It is fair to observe that *Facebook* involved a case where there was evidence of the defendants’ knowledge of the system’s potential

174. Section 10(b) of the Act makes it illegal “[t]o use or employ, in connection with the purchase or sale of any security . . . any manipulative or deceptive device.” 15 U.S.C. § 78j(b) (2012).

175. *Cornwell v. Credit Suisse Grp.*, 689 F. Supp. 2d 629, 637 (S.D.N.Y. 2010) (citing *Novak v. Kasaks*, 216 F.3d 300, 308 (2d Cir. 2000)).

176. *See In re Facebook, Inc., IPO Sec. & Derivative Litig.*, 986 F. Supp. 2d 428 (S.D.N.Y. 2013).

inadequacies.¹⁷⁷ Future defendants in A.I. manipulation cases might argue that since A.I. actions cannot be anticipated, then holding developers liable for any suspicious, unexplainable transaction will effectively impose strict liability. But this argument finds no sympathy under the existing Section 10(b) recklessness doctrine. A second example demonstrates that “knowledge” does not necessarily require a defendant to know or be able to explain exactly how fraudulent events transpired to become liable, *if* they ignored a duty to monitor the source of the fraud.

In *Rolf*, a plaintiff whose investments were squandered by a third-party advisor brought a Section 10(b) claim against their intermediary broker because the broker instilled confidence in the plaintiff about their investments with the advisor.¹⁷⁸ The court found that the defendant had acted recklessly in violation of Section 10(b) because he knew that the advisor had made unsuccessful investments in the past and failed to adequately supervise the advisor.¹⁷⁹ The court made no finding that the defendant intentionally encouraged the third-party advisor to fraudulently mismanage his client’s money.¹⁸⁰ Indeed, the defendant did not need to instruct, participate in, or specifically understand how the advisor mismanaged the money to be held liable. The advisor’s thoughts and intentions might as well have been a “black box” beyond the defendant’s comprehension, but the defendant’s failure to monitor and restrict the advisor’s fraudulent actions nevertheless constituted recklessness.¹⁸¹

This analogy of the human mind to an A.I. trader does not suggest that A.I. will be so human-like in its autonomy that straightforward applications of agency or vicarious liability are appropriate. It is simply a comparison to show that humans can be as opaque and complex as any A.I. machine, and this has not prevented courts from following circumstantial evidence of human recklessness. Therefore, if industry standards can provide a benchmark for A.I. oversight, courts will be equipped to similarly analyze evidence of human recklessness, where it exists.

177. *See id.* at 467.

178. *Rolf v. Blyth, Eastman Dillon & Co., Inc.*, 570 F.2d 38, 43 (2d Cir. 1978).

179. *Id.* at 47.

180. This is not surprising; the liability of the defendant was predicated on the investment advisor’s fraudulent acts, and had the investment advisor not acted fraudulently, the defendant could have avoided suit. *See id.*

181. The liability in *Rolf* was based on an aiding-and-abetting theory, before it was eliminated as a basis for liability by the Supreme Court in *Central Bank of Denver*. One court, however, has recognized that at the time, there were no grounds for distinction between primary violators with secondary roles and aiders and abettors. That same court concluded that if the same action were brought today, it also would most likely constitute a primary violation as well. *See In re Parmalat Sec. Litig.*, 376 F. Supp. 2d 472, 495 (S.D.N.Y. 2005).

B. NON-COMPLIANCE WITH A.I. INDUSTRY STANDARDS IS EVIDENCE OF RECKLESSNESS

Recklessness involves conduct that is “highly unreasonable and which represents an extreme departure from the standards of ordinary care to the extent that the danger was either known to the defendant or so obvious that the defendant must have been aware of it.”¹⁸² An industry standard for A.I. development will be useful in defining standards of ordinary care. However, courts will have the difficult task of determining when defendant developers must have been aware of risks and made an extreme departure from those ordinary standards.

This difficulty in showing scienter has existed in the Section 10(b) doctrine long before the use of computers.¹⁸³ Direct evidence of intent is rare, and correspondingly, Section 10(b) actions are commonly brought on the basis of circumstantial evidence.¹⁸⁴ To begin at the pleading stage, the Supreme Court has held that Section 10(b) complaints will only succeed “if a reasonable person would deem the inference of scienter cogent and at least as compelling as any opposing inference one could draw from the facts alleged.”¹⁸⁵ The Court noted that inferences need not be of the “‘smoking-gun’ genre, or even the ‘most plausible of competing inferences.’”¹⁸⁶ Nevertheless, this still presents a substantial burden to plaintiffs. For example, courts have dismissed claims when opposing inferences of mere “mismanagement” did not rise to the level of recklessness.¹⁸⁷ Indeed, in cases where claims ultimately allege that *inaction* resulted in a Section 10(b) violation, it is difficult for plaintiffs to plead facts distinguishing recklessness from opposing inferences of mere negligence.¹⁸⁸

At the pleading stage, however, courts should recognize that the black box problem could actually tip in favor of plaintiffs bringing A.I. manipulation claims. Because while the black box problem might complicate plaintiff’s success on the merits, all that is necessary at the pleading stage is an inference that is at least as strong as a non-culpable

182. *Kalnit v. Eichler*, 264 F.3d 131, 142 (2d Cir. 2001).

183. *See* *The Federal Corp.*, Exchange Act Release No. 34-3909, 25 S.E.C. 227, 1947 WL 24489, at *3 (Jan. 29, 1947) (“Since it is impossible to probe into the depths of a man’s mind, it is necessary in the usual case . . . that the finding of manipulative purpose be based on inferences drawn from circumstantial evidence.”).

184. This is evidenced by the fact that pleading standards for alleging Section 10(b) scienter are framed around circumstantial evidence, presumably because that is all plaintiffs have access to in most cases. *See* *ATSI Commc’ns, Inc. v. Shaar Fund, Ltd.*, 493 F.3d 87, 99 (2d Cir. 2007) (holding that a plaintiff may satisfy pleading requirements by alleging facts “constituting strong circumstantial evidence of conscious misbehavior or recklessness”).

185. *Tellabs, Inc. v. Makor Issues & Rights, Ltd.*, 551 U.S. 308, 324 (2007).

186. *Id.*

187. *See* *Westchester Teamsters Pension Fund v. UBS AG*, 604 F. App’x 5, 8 (2d Cir. 2015).

188. *See, e.g., S. Cherry St., LLC v. Hennessee Group LLC*, 573 F.3d 98 (2d Cir. 2009) (finding it more likely that the defendant financial firm was merely negligent, and not intentionally fraudulent, in failing to perform due diligence).

competing inference.¹⁸⁹ Courts have found that certain conduct on its face can give rise to an inference of manipulation.¹⁹⁰ Therefore, plaintiffs who specifically plead that a profit-motivated machine made a facially manipulative transaction might raise a plausible inference, forcing defendants to raise a competing one that is more compelling.

Note, it is not explicitly the defendant's burden to raise opposing inferences and phrasing it as such could be considered a mischaracterization.¹⁹¹ It is correct that at the pleading stage a defendant often does not even have an opportunity,¹⁹² let alone a burden, to put forth evidence refuting a claim. Accordingly, *Tellabs* held that trial courts should determine what competing inferences are raised in the complaint.¹⁹³ In some cases, however, courts discuss competing inferences gleaned from the complaint as arguments raised by defendants.¹⁹⁴ Overall, it could be said that defendants have an informal burden to raise a competing inference from the complaint if the court cannot locate one on its own initiative. In practice, extracting competing inferences is not really much of a burden to anyone, since it simply requires courts to draw on their own human experiences when considering alternate reasons for the alleged conduct.¹⁹⁵

When complaints allege manipulation by an A.I. black box, however, courts must realize that they cannot draw from this same well of human experience and common sense, and a *de facto* burden will fall on the defendant to raise a competing inference. Subsequently, when defendants cannot explain their A.I.'s intent beyond general profitability, courts should reject competing inferences of mere negligence as less compelling. Even when defendants actually have some explanation methods in place for monitoring, they still might only raise a factual dispute inappropriate for a motion to dismiss the complaint.¹⁹⁶ In both scenarios, plaintiffs would be able to push forward into discovery. While defendants might resent the

189. *Tellabs, Inc. v. Makor Issues & Rights, Ltd.*, 551 U.S. 308, 324 (2007).

190. *SEC v. Masri*, 523 F. Supp. 2d 361, 367 (2007) (acknowledging that in some cases “[t]he defendant’s manipulative intent can be inferred from the conduct itself”).

191. Since *Tellabs* describes a pleading standard, unsurprisingly, it is only focused on the plaintiff’s allegations in the complaint. *See Tellabs*, 551 U.S. at 324, 326 (holding that the inference of scienter should be at least as compelling as “any opposing inference one could draw from the facts alleged” and that courts must review “all the allegations holistically”).

192. *See id.* at 322 (“[F]aced with a Rule 12(b)(6) motion to dismiss a § 10(b) action, courts must, as with any motion to dismiss for failure to plead a claim on which relief can be granted, accept all factual allegations in the complaint as true.”).

193. *See id.* at 324, 326 (holding that the inference of scienter should be at least as compelling as “any opposing inference one could draw from the facts alleged” and that courts must review “all the allegations holistically”).

194. *See Matrixx Initiatives, Inc. v. Siracusano*, 563 U.S. 27, 48–49 (2011).

195. The standard for assessing inferences drawn from the complaint is from the point of view of a “reasonable person.” *Tellabs*, 551 U.S. at 324, 326.

196. *See id.* at 322 (“[F]aced with a Rule 12(b)(6) motion to dismiss a § 10(b) action, courts must, as with any motion to dismiss for failure to plead a claim on which relief can be granted, accept all factual allegations in the complaint as true.”).

progression of litigation, given the current absence of regulations, they would probably still feel confident about succeeding on the merits. Because with no obligation on the defendant to document machine behavior, the plaintiff would likely be unable to gather facts to prove the machine's manipulative intent beyond a preponderance of the evidence.¹⁹⁷

As discussed earlier, this will be the point at which industry standards will provide value. Defendants that are diligent in their adherence to A.I. development standards will be able to produce records and swiftly force summary judgment.¹⁹⁸ However, when defendants do not comply with all industry standards, a plaintiff with a meritorious claim can use the discovery process to show that deviation from A.I. industry standards was not just negligent, but reckless. This is where regulatory creation of paper trails will aid the plaintiff—not by carrying the weight of a statute, but by providing evidence of the conduct itself. Obviously, a defendant who does not adhere to standards and fails to keep records will weigh under heavier evidence of recklessness. Conversely, a defendant who cannot explain their machine in complete compliance with industry-best practice but puts other limitations in place to prevent against manipulation might fare better in proving away liability, despite their lack of compliance. In any event, industry standards will generate more evidence reflective of human scienter. At this point, courts can use the body of Section 10(b) doctrine to make adjudications as they have for the better part of a century.

CONCLUSION

In conclusion, there is a potent risk that A.I. will eventually participate in securities markets and will complicate enforcement of market manipulation regulations that only contemplate human actors. Therefore, carefully defining market manipulation and establishing effective industry standards will both deter A.I. manipulation and allow for enforcement of actions under the traditional Section 10(b) framework.

*Daniel W. Slemmer**

197. *See id.* at 328–29 (“A plaintiff alleging fraud in a § 10(b) action, we hold today, must plead facts rendering an inference of scienter *at least as likely as* any plausible opposing inference. At trial, she must then prove her case by a ‘preponderance of the evidence.’ Stated otherwise, she must demonstrate that it is *more likely* than not that the defendant acted with scienter.”).

198. *See Celotex Corp. v. Catrett*, 477 U.S. 317, 322 (1986) (“[E]ntry of summary judgment is mandated, after adequate time for discovery and upon motion, against a party who fails to make a showing sufficient to establish that existence of an element essential to that party’s case and on which that party will bear the burden of proof at trial.”).

* B.S.E.E., SUNY Buffalo, 2014; J.D. Candidate, Brooklyn Law School, 2020. Thank you to Bill Williams, Wilson Chow, Ian Kitts, and the entire Brooklyn Journal of Corporate, Financial & Commercial Law for their thoughtfulness and diligence in preparing this Note for publication.