

The Coming Failure of Manipulation Law?  
An Experimental Approach with Deep Reinforcement Learning

By Michael S. Barr,\* Gabriel Rauterberg,\*\* Megan Shearer,† and Michael P. Wellman†

*\*Gerald R. Ford School of Public Policy, \*\*Law School, and †Computer Science and Engineering, University of Michigan*

*Algorithms trading in financial markets increasingly deploy new and complex forms of artificial intelligence and machine learning. To study the potential trading strategies these algorithms may develop, we use an agent-based simulation to explore the behavior of algorithmic trading agents trained with deep reinforcement learning techniques. In our experimental setting, an agent trades directly in a market and also holds a portfolio of assets benchmarked to prices in that market. Although its reward function is simply to make profits, the algorithm trained through deep reinforcement learning autonomously develops trading strategies that are plausibly manipulative. In particular, the algorithm trades heavily and unprofitably in the market, but affects the benchmark's price, producing a net profit from its benchmark positions. If done intentionally, an individual engaging in such trading would have committed unlawful securities manipulation, but the algorithm was not designed to artificially affect prices, only to maximize profits. We use our experimental results to further underscore the need for reform of manipulation law, whose two core requirements are currently scienter (intent) and a "manipulative act." Demonstrating either of these elements for a reinforcement learning-trained algorithm will prove difficult both in concept and practice. Building on recent literature, we suggest ways in which the regulation of manipulation can become more robust against algorithmic challenges and in which the experimental study of algorithms can guide that agenda.*

## INTRODUCTION

Algorithms dominate trading in the stock market, and increasingly, they account for the majority of trading in other financial markets, such as those for options, futures, or treasuries.<sup>1</sup> Market participants use algorithmic trading to assimilate data, adjust orders, and execute transactions at scales and speeds that are not humanly possible. New forms of machine learning techniques, which allow algorithms to adapt in response to experience, are quickly being incorporated into firms' trading strategies.<sup>2</sup> A particular form of machine learning known as deep reinforcement learning has been successful in other domains, including computer vision, natural language processing, and video games. Its use in trading is only likely to become more common.

Commentators have worried that algorithms pose a variety of problems for the legal system. In securities regulation, this has included manipulation law.<sup>3</sup> The problem is baked into the definition of manipulation, which turns on a trader's intent.<sup>4</sup> The two key elements of a manipulation claim are generally "scienter" (i.e., a mental state approximating intent) and a "manipulative act." While the concept of manipulative act is elusive, it generally amounts to the creation of false price signals for a market.<sup>5</sup> Manipulation thus amounts to deliberately and artificially affecting market prices. The problem certain algorithmic trading strategies pose for manipulation law is that they can employ manipulative trading strategies without any individual ever intending that they do so and with effects on market prices that are opaque. Recognizing this possibility, commentators have called for significant changes in policy and law.<sup>6</sup>

But are these concerns about manipulative algorithms serious and imminent or speculative and improbable? Determining whether machine learning-trained algorithms are manipulating financial markets is no easy task. The most sophisticated algorithmic traders are

---

\*The authors acknowledge the generous support of the National Science Foundation (grant 1741190) [acknowledgments to come].

<sup>1</sup> See, e.g., COMMODITY FUTURES TRADING COMMISSION, *IMPACT OF AUTOMATED ORDERS IN FUTURES MARKETS* (2019).

<sup>2</sup> Ranko Masic, *Deep Reinforcement Learning Based Trading Application at JP Morgan Chase*, MEDIUM, <https://medium.com/@ranko.masic/reinforcement-learning-based-trading-application-at-jp-morgan-chase-f829b8ec54f2>.

<sup>3</sup> 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 (2015) (arguing that algorithmic trading poses a problem for manipulation in the futures market because the relevant claims turn on a requirement of intent and suggesting that CFTC rules create a failure to supervise claim that could be an effective tool in regulating manipulation); Daniel W. Slemmer, *Artificial Intelligence & Artificial Prices: Safeguarding Securities Markets from Manipulation by Non-Human Actors*, 14 BROOK. J. CORP. FIN. & COM. L. 149 (2019); Gina-Gail S. Fletcher, *Legitimate Yet Manipulative: The Conundrum of Open-Market Manipulation*, 68 DUKE L.J. 479 (2018) (arguing against the usefulness of an intent-based conception of market manipulation); Yesha Yadav, *The Failure of Liability in Modern Markets*, 102 VA. L. REV. 1031, 1073-74 (2016); see also Yavar Bathaee, *The Artificial Intelligence Black Box and the Failure of Intent and Causation*, 31 HARV. J. L. & TECH. 889 (2018); Tom C.W. Lin, *The New Market Manipulation*, 66 EMORY L.J. 1253 (2017) (discussing how the emergence of new financial technologies enabled novel forms of manipulation and identifying reform proposals).

<sup>4</sup> Sec. & Exch. Comm'n v. Lek Sec. Corp., 276 F. Supp. 3d 49, 60 (S.D.N.Y. 2017) (noting that "Market manipulation can be accomplished through otherwise legal means," and "in some cases scienter is the only factor that distinguishes legitimate trading from improper manipulation.") (citations omitted).

<sup>5</sup> See *infra* notes \_\_ and accompanying text.

<sup>6</sup> See, e.g., Scopino, *supra* note 3; Fletcher, *supra* note 3; Alessio Azzutti, Wolf-Georg Ringe, H. Siegfried Stiehl, *Machine Learning, Market Manipulation and Collusion on Capital Markets: Why the 'Black Box' Matters*, European Banking Institute Working Paper Series 2021 - no. 84.

notoriously secretive about the basic facts of their business model, let alone about the details of their most advanced trading techniques. Nor is advanced algorithmic trading likely to be revealed through disciplinary actions. Securities manipulation is already difficult to successfully prosecute when the manipulation is done by conventional means.

To gain traction on this issue, we use an experimental approach to understand the trading strategies that new machine learning techniques may produce. In this paper, we report on the results of our analysis, while aiming for accessibility to those outside computer science (which includes two of the authors).<sup>7</sup> We use our experimental analysis to motivate a specific agenda for reform and to provide guidance for tailoring legal solutions.

An experimental approach offers powerful advantages in studying algorithmic manipulation. The equity market involves a very large numbers of traders interacting in an institutionally complex setting. As a result, formal models may ignore important features of market microstructure or become intractable as larger numbers of participants and strategies are introduced.<sup>8</sup> Our use of a simulated market enables us to situate agents in a market microstructure that captures some of the basic features of the real-world equity market.

We study a specific form of manipulation—the manipulation of a financial benchmark that is calculated on the basis of market transaction prices. Many of the highest profile financial scandals of the last decade involved the manipulation of financial benchmarks. The London Interbank Offered Rate (“LIBOR”), an estimate of the rate at which banks can borrow from one another, was widely used across more than \$300 trillion in loans. Yet several banks were implicated in the manipulation of LIBOR, with regulatory and civil settlements amounting to over a billion dollars and criminal prosecutions of multiple individuals.<sup>9</sup> There have been alleged manipulations of transaction-based benchmarks, such as the Chicago Board Options Exchange’s Volatility Index (“VIX”) or World Markets/Reuters Closing Spot Rates in foreign exchange. Economic research has also explored the manipulation of securities’ closing prices, which are a kind of transaction-based benchmark.<sup>10</sup>

Our market consists of a single trader with external holdings dependent on a benchmark, who transacts with numerous background traders with no benchmark-linked holdings. The market is a continuous double auction in a standard limit order book and traders transact in a single security. The value of the benchmark is determined by the prices of transactions in the security. The background traders employ what is known as a zero-intelligence (ZI) strategy.<sup>11</sup> ZI traders submit limit orders shaded away from their valuations so as to produce a randomized positive surplus if executed. The ZI strategy has been found to reproduce stylized facts in

---

<sup>7</sup> For those interested in a full technical presentation of the results, see Megan Shearer, Gabriel Rauterberg & Michael P. Wellman, *An Agent-Based Model of Financial Benchmark Manipulation*, Thirty-sixth International Conference on Machine Learning: AI and Finance Workshop (2019), <https://par.nsf.gov/servlets/purl/10105527>; Megan Shearer et al., *An Agent-Based Model of Financial Benchmark Manipulation*, (2021).

<sup>8</sup> For instance, Duffie develops a theoretical model suggesting the optimality of a certain benchmark structure, but models a market with almost no microstructure detail. In a project exploring whether reinforcement learning algorithms autonomously develop collusive pricing strategies, one paper notes “On the theoretical side, the interaction among reinforcement-learning algorithms in pricing games generates stochastic dynamic systems so complex that analytical results seem currently out of reach.” Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolò & Sergio Pastorello, *Artificial Intelligence, Algorithmic Pricing, and Collusion*, 110 Am. Econ. Rev. 3267 (2020).

<sup>9</sup> Rauterberg & Verstein, *supra* note 27.

<sup>10</sup> Carole Comerton-Forde & Talis J. Putnins, *Stock Price Manipulation: Prevalence and Determinants*, 18 REV. FIN. 23 (2014).

<sup>11</sup> See Shearer, Rauterberg & Wellman, *supra* note 7.

simulated markets and is widely employed in agent-based studies. For the benchmark trader, we explore two strategy designs. The first is a hand-crafted variant of ZI, that offsets its bid to influence the benchmark in one direction or another. The second derives its trading strategies through deep reinforcement learning. Reinforcement learning (RL) is a machine learning approach that aims to derive a strategy (a mapping of observations of an environment to actions) that maximizes a given reward function. RL works by incrementally adjusting a strategy based on feedback received through repeated interaction with an environment. Deep RL employs multi-layer neural networks in its strategy representation. In principle, an RL algorithm can generate any strategy within its space of representable mappings as long as it tends to maximize the reward function.

The ZI trading strategy includes several tunable parameters, which we set using a method known as empirical game-theoretic analysis. EGTA employs agent-based simulation to identify Nash equilibria in games over a heuristic strategy space.<sup>12</sup> Simulation data provides the basis for inducing a normal-form game model, which is then analyzed using standard game-solving algorithms.<sup>13</sup> In our baseline setting, we identify equilibrium strategy profiles both without manipulation, and with a manipulator based on the ZI strategy. This manipulator trades unprofitably on the primary security in order to move the benchmark and thus increase its profit overall.

We further find in our experimental result that training with deep RL, without explicit direction, also produces a “manipulative” strategy that acts in the primary market so as to affect the value of its external holdings, resulting in net profits.

While our empirical analysis is in the spirit of a “proof of concept,” our results strongly suggest that trading algorithms can autonomously develop manipulative trading strategies. If this is true, there are immediate policy implications.

If an individual deliberately traded heavily but unprofitably in a market to affect a benchmark’s price, and thus increase the value of her benchmark-linked positions, it would almost certainly be unlawful securities manipulation. In our model, however, no individual intended the specific strategies the algorithm deployed, and the specific machine learning techniques involved, we will argue, make showing causation quite difficult.<sup>14</sup> We thus use our experimental results to motivate an agenda for reform of manipulation law.

Four main issues are relevant to potential changes to regulation:

- (1) Should the *mode* of regulation change (e.g., the composition of ex ante and ex post regulatory techniques)?
- (2) Should the identity of the *regulator* change?
- (3) Should the identity of the *regulated* change?

---

<sup>12</sup> Brinkman, E. and Wellman, M. P. Empirical mechanism design for optimizing clearing interval in frequent call markets. In 18th ACM Conference on Economics and Computation, pp. 205–221, 2017.

Wah, E., Lahaie, S., and Pennock, D. M. An empirical game-theoretic analysis of price discovery in prediction markets. In 26th International Joint Conference on Artificial Intelligence, pp. 510–516, 2016.

Wah, E., Wright, M., and Wellman, M. P. Welfare effects of market making in continuous double auctions. *Journal of Artificial Intelligence Research*, 59:613–650, 2017; Wellman, M. P. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems*, 30:1175–1189, 2016.

<sup>13</sup> See Shearer, Rauterberg & Wellman, *supra* note 7.

<sup>14</sup> Prior work has noted that machine learning algorithms will pose problems for proof of causation due to their complexity and the difficulty individuals sometimes face in interpreting their underlying processes. See, e.g., Azzutti, *supra* note 6.

(4) Should the substantive content of legal *claims* change?

In Table 1, we depict the main options for questions 1 and 2 along the vertical and horizontal axes, respectively, and provide examples regarding question 3 in the cells. Our main conclusion is that regulatory policy addressing manipulation should broaden out from the upper left-hand corner of Table 1, where it is currently focused, toward the bottom right-hand corner. It should shift from an emphasis on ex post criminal and administrative enforcement by governmental actors toward a larger supervisory ecology that uses ex ante regulation and oversight, as well as ex post tools with lower sanctions (and lower burdens of proof) to deter a broader range of undesirable trading strategies that fall short of traditional indicia of malice and intent. We are not arguing that criminal sanctions should not be used—they should when the requisite elements are demonstrable—but rather that criminal sanctions are unlikely to deter the range of conduct that could impair market integrity and harm other market participants.

Alongside criminal enforcement by the Department of Justice (DOJ) and administrative enforcement by the SEC and CFTC, the stock exchanges and national securities associations (such as the Financial Industry Regulatory Authority (FINRA)) function as self-regulatory organizations, while market intermediaries, such as broker-dealers, are also conscripted into serving a gate-keeping function. The future of trading regulation will need to rely on developing this interacting regulatory system in specific ways.

**Table 1. Manipulation Enforcement: Who Enforces and How?**

	DOJ	SEC/CFTC	FINRA	EXCHANGES	BROKER-DEALERS
<b>Ex post</b>	Criminal proceedings	Administrative proceedings	Disciplinary proceedings	Expulsion/Rule 611	Denial of market access
<b>Ex ante</b>	Forward-looking settlements	Market structure regulation	Surveillance	Market design	Post-trade surveillance

←

**Burden of proof and severity of sanction**

The difficulty of prosecuting complex, pervasive, and opaque algorithms also means that a wide variety of ex ante structural remedies that foreclose manipulation or make it less profitable are also worth considering.<sup>15</sup> For instance, we suggest repurposing existing tools, such as the Market Access Rule, to conscript broker-dealers into a gatekeeping function

<sup>15</sup> For important recent analyses of manipulation in financial markets, *see, e.g.*, Tom C.W. Lin, *The New Market Manipulation*, 66 EMORY L.J. 1253 (2017) (discussing how the emergence of new financial technologies enabled novel forms of manipulation and identifying reform proposals); Gina-Gail S. Fletcher, *Legitimate Yet Manipulative: The Conundrum of Open-Market Manipulation*, 68 DUKE L.J. 479 (2018) (arguing against the usefulness of an intent-based conception of market manipulation); Andrew Verstein, *Benchmark Manipulation*, 56 B.C. L. REV. 215, 220, 242-43 (2015) (discussing the extensive manipulation of financial benchmarks).

monitoring their customers' order flow.<sup>16</sup> Licensure requirements for algorithmic developers and traders should also incorporate ethical guidance on algorithmic development, and should require monitoring algorithms for their effects and periodic algorithmic impact assessments for the most heavily used or potentially problematic algorithms. On the ex ante, structural side of the ledger, speed bumps and frequent call markets could both reduce the current market-wide emphasis on speed and reduce certain forms of speed-based manipulation.<sup>17</sup> Certain forms of financial transaction taxes, such as a tax on canceled trades, could also be deployed to combat certain forms of manipulation known as "spoofing."<sup>18</sup> We also argue that manipulation law should transition toward a more purely functional or conduct-based definition of manipulation that makes no reference to human intent, and show that in some ways, steps in this direction have already been taken.<sup>19</sup>

We proceed as follows. In Part I we briefly describe the law governing manipulation across a number of financial trading markets. In Part II, we report on our experimental analysis, describing the nature of our approach, the techniques deployed, and the results obtained. In Part III, we discuss potential legal reforms.

## I. PROSECUTING ALGORITHMIC MANIPULATION

In Section A, we provide a brief overview of manipulation law. In Section B, we discuss the difficulties that an algorithm trained on deep reinforcement learning techniques might pose for manipulation law.

### A. *Manipulation Law*

Currently, the U.S. legal system does not provide a single claim for addressing manipulation in financial markets. Instead, two administrative agencies, the Securities and Exchange Commission (SEC) and Commodities Futures Trading Commission (CFTC), each regulate a broad set of financial instruments. New types of assets, such as initial coin offerings or crypto-currencies, are typically assimilated to one of the preexisting regimes. Only in recent years, have federal prosecutors made significant use of generic criminal provisions that could be used to sanction manipulation of any financial instrument.<sup>20</sup>

The concept of manipulation itself is notoriously hard to define, resulting in a body of law marked by conflict and confusion,<sup>21</sup> but while details vary, there is substantial overlap in the elements that a regulator must prove in a manipulation suit. The essence of manipulation law is a prohibition on deliberately and artificially affecting markets. Its key doctrinal elements are that a trader (1) intentionally or knowingly engaged in (2) a manipulative act.<sup>22</sup> A manipulative act consists of orders or transactions that send false price signals to the market, producing artificial prices that do not reflect the natural interplay of supply and demand. For

---

<sup>16</sup> See *infra* Section III.B.

<sup>17</sup> See *infra* Section III.A.

<sup>18</sup> See NSF Grant 1741190.

<sup>19</sup> See *infra* Part IV.

<sup>20</sup> *United States v. Blaszczyk*, 947 F.3d 19 (2d Cir. 2019).

<sup>21</sup> For analyses of different forms of manipulation, see, e.g., Merritt B. Fox, Lawrence R. Glosten & Gabriel V. Rauterberg, *Stock Market Manipulation and Its Regulation*, 35 YALE J. REG. 67 (2018); Tālis J. Putniņš, *Market Manipulation: A Survey*, 26 J. ECON. SURVS. 952 (2012).

<sup>22</sup> *ATSI Commc'ns, Inc. v. Shaar Fund, Ltd.*, 493 F.3d 87, 100 (2d Cir. 2007).

instance, successful criminal prosecutions of manipulation have involved traders who engaged in “spoofing”—posting orders to affect other market participants’ views of an asset’s value, when the manipulator had no intention that those orders ever execute.<sup>23</sup>

### B. Prosecuting Algorithmic Manipulation

Others have noted the basic problems machine learning-based algorithms pose for prosecuting manipulation and so we only briefly summarize them here.<sup>24</sup> The problem with scienter flows from manipulation law’s emphasis on intent but an autonomously developing algorithm’s lack of any mental state attributable directly to a designer. No substantive offense in securities law, and few elsewhere in the law, turn so decisively on individual intent as manipulation. As courts have often noted, “in some cases scienter is the only factor that distinguishes legitimate trading from improper manipulation.”<sup>25</sup>

There are problems with showing causation too. With a deep reinforcement learning approach, it will often not be clear what the algorithm’s trading strategy actually was and therefore it will be hard to prove that the machine sent “false price signals” to the market. In addition, DRL algorithms continuously adapt to their environment, including the other DRL algorithms against which they trade. Thus, prosecutors and regulators will have a hard time understanding why and how an algorithm affected market prices, or generating plausible counterfactuals without the alleged manipulation.

Showing that an algorithm trained through deep reinforcement learning has scienter or engaged in a manipulative act can thus prove extremely difficult. Machine learning algorithms that use deep reinforcement learning adaptively generate their own strategies to achieve objectives, developing in ways that their designers may not have intended, nor be able to understand or explain. Yet, these algorithms can generate highly profitable trading strategies, potentially executing the same trades that would be unlawful if done deliberately by an individual trader. The problem particular to machine learning techniques is that they can generate algorithms that iteratively and adaptively respond to their environment, developing strategies that their designers did not intend; those algorithms can also generate strategies that no one is able to interpret or explain, even *ex post*.<sup>26</sup> What is thus transformational about these algorithms is the extent to which their ultimate strategies may not be the product of human intent and the difficulties their designers (or prosecutors) will face in understanding precisely what trading strategy the algorithm developed and why.

## II. EMPIRICAL ANALYSIS

In Section A, we explain our choice of modeling the manipulation of a financial benchmark. In Section B, we explain our simulation approach, known as empirical game-

---

<sup>23</sup> United States v. Coscia, 866 F.3d 782 (7th Cir. 2017).

<sup>24</sup> See Azzutti, Ringe & Stiehl, *supra* note 6; Scopino, *supra* note 3.

<sup>25</sup> ATSI Commc’ns, Inc. v. Shaar Fund, Ltd., 493 F.3d 87, 102 (2d Cir. 2007); *see also* Sec. & Exch. Comm’n v. Lek Sec. Corp., 276 F. Supp. 3d 49, 60 (S.D.N.Y. 2017) (“Market manipulation can be accomplished through otherwise legal means.”).

<sup>26</sup> Both scholars and media commentators have noted that algorithmic trading challenges manipulation law’s reliance on intent. *See e.g.*, Daniel W. Slemmer, *Artificial Intelligence & Artificial Prices: Safeguarding Securities Markets from Manipulation by Non-Human Actors*, 14 BROOK. J. CORP. FIN. & COM. L. 149 (2019) (arguing that creators of artificial intelligences should garner feedback from them and maintain evidentiary records).

theoretic analysis, as well as the machine learning techniques we use to train our trading algorithm.

### *A. Benchmark Manipulation*

A financial benchmark is a summary statistic designed to convey information about a financial reality. For instance, the value of the S&P 500 index is commonly used as shorthand for the performance of the United States stock market as a whole. Benchmarks serve a large number of valuable market functions,<sup>27</sup> including serving as price terms in a large number of financial contracts.

Recall that LIBOR, a single benchmark for short-term interest rates, supported more than \$300 trillion in global loan volume. In particular, the LIBOR rate was used as a term in contracts from mortgage loans to interest rate swaps.<sup>28</sup> In 2011, however, credible allegations emerged that banks had manipulated the LIBOR rate before and during the financial crisis. Since then, regulators have investigated other benchmarks for manipulation and imposed some of the largest penalties ever paid by financial institutions. Reforming benchmarks has become a focus of regulatory and academic energy, and LIBOR is being abandoned.<sup>29</sup> One central question is whether alternatives to LIBOR could also be manipulated or used to manipulate other markets. The proposed replacement for LIBOR is the transaction-based benchmark Secured Overnight Finance Rate (SOFR), which would be difficult to manipulate in a manner resembling the manipulation of LIBOR. However, a trader using DRL could include its benefits from SOFR in its reward function while trading in the Treasury market and its policy could learn to submit orders to impact SOFR and increase its profits from the benchmark.

Precisely because benchmarks provide a concise statement of market realities, and are used in contracts, market actors have incentives to move the price of benchmarks, potentially reducing the informativeness of benchmarks' prices. Indeed, many of the highest profile financial manipulations of the last decade have involved benchmarks, resulting in some of the largest fines ever levied against financial institutions.<sup>30</sup> Besides the practical relevance of the setting, there is also a small but rich theoretical literature studying the optimal structure of financial markets to which we hope to contribute.<sup>31</sup>

### *B. Empirical Game-Theoretic Analysis and Deep Reinforcement Learning*

We model a market environment in which a potentially manipulative algorithm interacts with a population of non-manipulative agents. The manipulator trades directly in a market that determines a benchmark's price and it also has external holdings whose value is based on the benchmark. We calculate the benchmark as the volume-weighted average price ("VWAP") of all transactions that occur during each trading period we model. We use VWAP because it is

---

<sup>27</sup> Gabriel V. Rauterberg & Andrew Verstein, *Index Theory: The Law, Promise, and Failure of Financial Indices*, 30 YALE J. REG. 1 (2013).

<sup>28</sup> *Id.*

<sup>29</sup> Darrell Duffie & Piotr Dworczak, *Robust Benchmark Design*, NBER WORKING PAPERS 20540.

<sup>30</sup> See, e.g., Carole Comerton-Forde & Talis J. Putnins, *Stock Price Manipulation: Prevalence and Determinants*, 18 REV. FIN. 23 (2014); see also Andrew Verstein, *Benchmark Manipulation*, 56 B.C. L. REV. 215, 220, 242-43 (2015).

<sup>31</sup> Rauterberg & Verstein, *supra* note 27.



widely relied on in practice and has prior theoretical support as being relatively robust to manipulation and thus providing a difficult test-case for successful manipulation.<sup>32</sup>

The manipulator adjusts its orders systematically to influence the benchmark in a specific direction. While it may lose in the market through its trades, it can earn a net profit if it can sufficiently shift the benchmark price to alter the value of its external holdings.

Our market, like the stock market, operates as a continuous double auction, where participants submit limit orders to trade a security with a fundamental value that varies over time. Traders' valuation of the security combines their estimate of the fundamental value of the security based on their noisy observations with their individual preferences over being long and short (due to consumption, liquidity needs, etc.) plus market information based on the actions of other traders. Non-manipulative traders arrive to the market according to a random process and trade so as to improve their surplus. The manipulator similarly has a noisy view of the fundamental value and a private valuation. It submits orders to maximize the joint sum of its trading profits and external holdings. As noted, at the end of trading, the benchmark is calculated based on the VWAP of transaction prices in the trading market.

To determine the effects of benchmark manipulation in equilibrium settings, we employ a method known as empirical game-theoretic analysis (EGTA) over a variety of market environments. EGTA constructs a game-theoretic model through agent-based simulation, and identifies Nash equilibria over representative sets of agent strategies.<sup>33</sup> An iterative computational process is used to identify potential equilibria, and then confirm or refute them by examining deviations until quiescence.<sup>34</sup>

We first conduct EGTA over background traders and manipulators using ZI strategies. To account for stochastic effects, we sample at least 50,000 simulation runs to estimate payoffs of relevant strategy profiles. We then calculate the equilibrium surplus for agents and the market in a variety of settings where there is and is not a manipulator. Intuitively, the manipulator loses money in market trading—precisely because it trades at prices it does not believe represent the fundamental value—enriching the other trading agents, but successfully manipulates the benchmark at the expense of benchmark counterparties. In sum, the manipulator's total surplus (combining its trading outcome and external holdings) increases with manipulation, although its market trading surplus decreases. The allocation in the manipulated market is also less efficient. Figure 1 displays (a) the total surplus of the manipulative trading agent with and without a manipulation, and (b) the surplus of the manipulative agent exclusively from trading in the market (i.e., excluding the value of its benchmark holdings) with and without manipulation.<sup>35</sup>

---

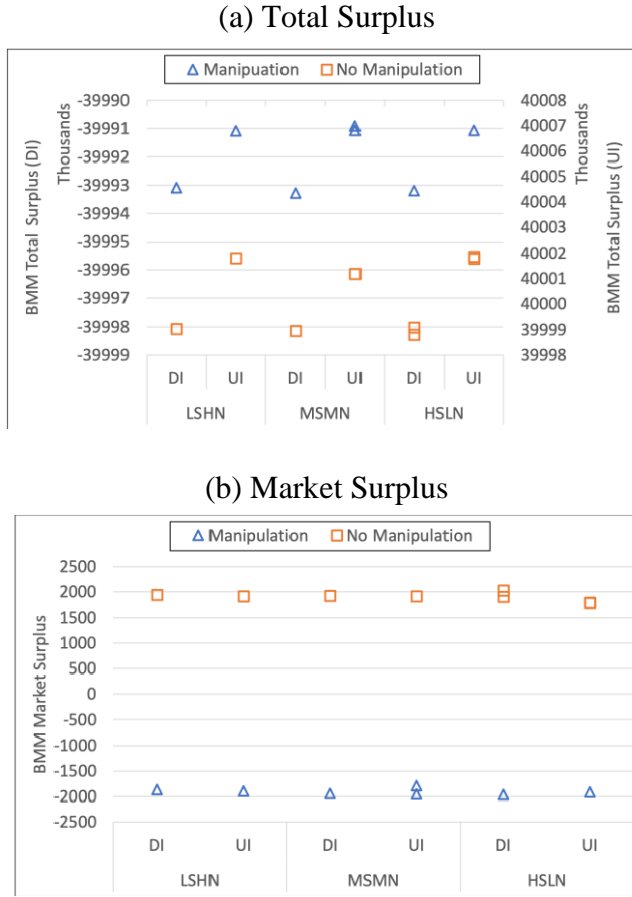
<sup>32</sup> See Duffie, D. and Dworzak, P. Robust benchmark design. (3175), 2018; see also Duffie, D. and Stein, J. C. Reforming LIBOR and other financial benchmarks. *Journal of Economic Perspectives*, 29(2):191–212, 2015.

<sup>33</sup> Michael P. Wellman, *Putting the Agent in Agent-Based Modeling*, 30 AUTONOMOUS AGENTS AND MULTI-AGENT SYSTEMS 1175 (2016).

<sup>34</sup> See Shearer, Rauterberg & Wellman, *supra* note 7.

<sup>35</sup> For more detail, see Shearer, Rauterberg & Wellman, *supra* note 7, at 5. On the primary y-axis is total surplus when the manipulator attempts to shift the benchmark down, while the secondary y-axis shows total surplus for attempts to shift the benchmark up. In both figures, the x-axis is different market environments we model in which material parameters (e.g., fundamental value, observation variance, benchmark impact) vary.

Figure 1. Surplus from Manipulation by Zero Intelligence Trading Agent



We then enrich our analysis by introducing a deep reinforcement learning (“DRL”) algorithm. In particular, we use a DRL technique known as a deep deterministic policy gradient, to learn a trading strategy for this environment with external holdings based on the benchmark.<sup>36</sup> The DRL-trained agent’s goal is to maximize its combined utility from the market and benchmark; the reward function includes incremental payoffs from changes to the benchmark holdings, while the algorithm is trading in the market. The state space for this agent is too large to be represented by a Q-table—a lookup table with combinations of states, actions, and rewards—and so deep reinforcement learning is a necessity. A challenge that arises from this environment is that the action space, prices to submit, is continuous, so we use deep deterministic policy gradient to train our agent. We study whether the agent learns to affect the benchmark by trading directly in the market.

The DRL algorithm learns a trading strategy that outperforms the hand-coded ZI manipulation strategy. Importantly, the DRL algorithm learns to alter the benchmark’s value by submitting very unprofitable market orders to increase the size of its external holdings from the benchmark. In short, it manipulates the benchmark to maximize its profits, despite the fact that it was not designed or intended to engage in any manipulation. The inputs necessary for

<sup>36</sup> Deep deterministic policy gradient is a model-free, off-policy actor-critic algorithm that uses a neural network to learn a continuous action space.

manipulation were simply a strategy space that included the ordinary submission of orders and a reward function to maximize utility that included both profits from trading in the market and profits from the external holding.

While relatively simple in concept, our manipulative DRL algorithm has immediate and important consequences for policy beyond highlighting the lack of individual intent at the level of trading strategy. The DRL algorithm learns to manipulate the benchmark only because its reward function includes the external holdings. Thus, one way to prevent this strategy would be to limit the scope of the reward function. When should a trading algorithm be made aware of a firm's benchmark-linked positions? Does it matter if the algorithm is deciding to trade (e.g., increasing the firm's purchasing or selling of a security at the close of trading when a contract price is fixed?) or *not* to trade (e.g., terminating trades that would otherwise have executed at the close of trading because they would disadvantageously affect the contract price of a security)? These are perennial and puzzling questions for manipulation law, but they are posed anew with sharp immediacy by algorithmic trading agents.

The results also illustrate the distinctive advantages of our approach. To our knowledge, this study is the first to train a trading strategy to manipulate benchmarks without design. It adds additional realism to the existing theoretical literature, which generally features markets without any microstructure (i.e., without any details regarding the mechanics of trade). A simulation approach also allows for the complexity of many agents' interaction in a trading environment where purely analytic game theoretic approaches are unlikely to be tractable.

### III. REORIENTING POLICY

Designing an effective regulatory system for manipulation requires both reconceptualizing what it is and recalibrating its enforcement regime. The rise of ubiquitous, automatic, and essentially instantaneous algorithms, trained through machine learning, calls for a new approach. In particular, we argue it calls for reorienting our regulatory approach away from its current emphasis on ex post governmental enforcement actions with high sanctions and burdens of proof, toward a supervisory ecology that relies more on ex ante structural solutions and ex post enforcement with lower sanctions, weaker burdens of proof, and rapid timetables.<sup>37</sup>

In general, several features of algorithmic trading favor reorienting policy toward ex ante and structural remedies to manipulation, and when enforcement occurs ex post, favor lower sanction, lower proof claims over those with higher sanctions and higher burdens of proof. Changes to the regulatory environment will likely also mean that there should be changes to which regulators are involved and potentially also to which persons should be regulated, as illustrated in Figure 1. One core reason why legal reform is advisable has been developed in the preceding parts—algorithms trained through deep reinforcement learning can harm markets in the same ways and through the same trading behavior as an individual's intentional manipulation, but may be prohibitively difficult to prosecute under existing law. Four other features of algorithmic trading also heighten the need for legal reform.

First, algorithmic trading is not only part of most financial trading markets, it is increasingly pervasive, and in many markets, human trading is now the exception. Trading markets involve a vast array of market participants, interacting in continuous time at levels of

---

<sup>37</sup> For other reasons, such as the errors and contagion possible in algorithmic markets, some call for wholesale changes to the nature of liability in financial markets. For an important example, *see* Yadav, *supra* note \_\_, at \_\_.

granularity where microseconds (millionths of a second) matter. The pervasive role of algorithms reflects the transformation of how finance works. Algorithmic trading is a basic feature of trading markets for the foreseeable future. In other contexts, scholars have called for disclosure to reduce problems with the use of algorithms, for example, in sentencing.<sup>38</sup> When it comes to courts' use of algorithms to assist judicial sentencing decisions, it may be desirable to require the government to retain a complete record of the algorithm's code, and even to require its ex ante disclosure to a regulatory agency or a broader set of stakeholders. Any ex ante disclosure mandate in trading, however, would not likely be effective, given that trading has essentially become an algorithmic business, with thousands of quickly-evolving algorithms continuously interacting with each other.

A second important feature is that the individual victims of a manipulation may be highly diffuse and each may have incurred only limited damages. In some cases, as in our benchmark model, background traders in the underlying asset benefited from the manipulation, and so lack any incentive to police it. Many financial wrongs, and certainly manipulation, have always posed this problem of diffuse harms. The algorithmic character of trading only heightens it, however, because it increases the speed and scale at which even an individual trader can act. This feature favors a greater role for regulators (as opposed to private plaintiffs) or at least the greater use of aggregative claims in private litigation. It may be difficult for plaintiffs' lawyers to identify any victim, however, and thus locate a prospective plaintiff for a suit to represent the class of individuals who were injured.

A third feature involves the wrongfulness of the kind of manipulation on which this article has focused. We have been emphasizing algorithms that develop manipulative trading strategies in ways that their designers did not intend. Because of this lack of intentional wrongdoing, legal claims with severe criminal sanctions may be normatively inappropriate.

A fourth feature is their technological complexity. Understanding machine learning algorithms require significant technical competence as well as costly resources to analyze the algorithm at issue. These features mean there are likely to be economies of scale and scope in regulating algorithms.<sup>39</sup> It also means that regulators closer to actual trading and with greater knowledge of current algorithmic developments will be better suited to bringing effective enforcement actions.

In Section A, we suggest clarifying manipulation law. In Section B, we consider ex ante changes to law and policy. Lastly, in Section C, we consider reforms to ex post enforcement.

### **A. *Clarifying Manipulation Law***

Manipulation law should be clarified in two ways. First, the Supreme Court should bring clarity, certainty, and predictability to black-letter law about the meaning of the central prohibitions on securities manipulation by resolving federal circuit splits. Second, FINRA could promulgate clearer guidance as to what constitutes manipulative activity for borderline cases.

---

<sup>38</sup> See, e.g., Aziz Z. Huq, *Constitutional Rights in the Machine-Learning State*, 105 CORNELL L. REV. 1875, 1947 (2020) (suggesting as a potential disclosure mandate that “an algorithmic decision should be accompanied by a ‘datasheet’ that records the choices and manipulations of training data”); see also Aziz Z. Huq, *A Right to A Human Decision*, 106 VA. L. REV. 611 (2020).

<sup>39</sup> See NSF Grant 1741190

Prohibiting the manipulation of securities was a significant motivation for enacting the federal securities laws almost a century ago.<sup>40</sup> A striking fact is that federal courts still disagree about the central question of manipulation law—whether open market manipulation is unlawful. While some courts have held that it is, others have been impressed by reasonable, but ultimately false arguments that there is no need for manipulation law. For instance, one famous argument concludes that profitable open-market manipulation is not possible, and that the practice is thus self-detering.<sup>41</sup> Skepticism about the coherency of open market manipulation as a concept may have also contributed to the doctrinally unfounded demand by some courts that manipulation involve the dissemination of false information by the manipulator,<sup>42</sup> or skepticism that commission of a crime can turn solely on a trader’s intent.<sup>43</sup> This confusion almost certainly has downstream consequences for prosecutors’ willingness to bring enforcement actions.

As our model suggests, algorithmic manipulation underlines that it is a mistake to be skeptical of the profitability of market manipulation. Open market manipulation can be profitable and an algorithm can easily develop profitable manipulative strategies without even being designed to do so. The Supreme Court should clarify that Section 10(b) prohibits open market manipulation and that trading alone can constitute a manipulative act.<sup>44</sup>

Second, FINRA could further clarify forms of manipulation prohibited under its broad rules. As a self-regulatory organization, FINRA enjoys the freedom to promulgate relatively broad mandates for its members conduct, which might be problematic as part of the criminal law.<sup>45</sup> For instance, FINRA Rule 2010 requires a member to “observe high standards of commercial honor and just and equitable principles of trade” when conducting its business.<sup>46</sup> Rule 2020 more explicitly requires that a member will not “effect any transaction in, or induce

---

<sup>40</sup> Merritt B. Fox, Lawrence R. Glosten & Gabriel V. Rauterberg, *Stock Market Manipulation and Its Regulation*, 35 YALE J. REG. 67 (2018).

<sup>41</sup> Daniel R. Fischel & David Ross, *Should the Law Prohibit “Manipulation” in Financial Markets?*, 105 HARV. L. REV. 503 (1991) (arguing against the usefulness of the concept of manipulation).

<sup>42</sup> See, e.g., *GFL Advantage Fund, Ltd. v. Colkitt*, 272 F.3d 189 (3d Cir. 2001) (“[r]egardless of whether market manipulation is achieved through deceptive trading activities or deceptive statements as to the issuing corporation’s value, it is clear that the essential element of the claim is that *inaccurate* information is being injected into the marketplace.”)

<sup>43</sup> *United States v. Mulheren*, 938 F.2d 364, 368 (2d Cir. 1991) (“although we have misgivings about the government’s view of the law, we will assume, without deciding . . . that an investor may lawfully be convicted under Rule 10b-5 where the purpose of his transaction is solely to affect the price of a security.”). Steve Thel was the first to show the broad significance of *Mulheren* for manipulation scholarship. Steve Thel, *\$850,000 in Six Minutes – The Mechanics of Securities Manipulation*, 79 CORNELL L. REV. 219, 240-47 (1994) (discussing how a manipulator may profit by trading so as to alter others’ expectations)

<sup>44</sup> In 2016, the United States Supreme Court was given the opportunity to resolve the circuit split. An opinion of the D.C. Circuit reflected its longstanding position that mere trading could constitute a manipulative act. *Koch v. S.E.C.*, 793 F.3d 147, 155 (D.C. Cir. 2015). The accused manipulator sought Supreme Court review, petitioning for a writ of certiorari based on the split among the circuits and the confusion and disparities of outcome it creates for market participants. *Koch v. SEC, U.S.*, No. 15-781, 3/28/16. The Supreme Court denied the petition, however. *Koch v. S.E.C.*, 793 F.3d 147, 155 (D.C. Cir. 2015), *cert. denied*, 136 S. Ct. 1492, 194 L. Ed. 2d 586 (2016).

<sup>45</sup> The Securities Exchange Act both authorizes and requires FINRA to enforce its member’s compliance with the statute itself, the regulations thereunder, and FINRA’s own rules. See Securities Exchange Act, Section 15(A), 15 U.S.C. 78o-3(b)(2). The remedies available to FINRA include censure, mandates to take remedial actions, restrictions on a member’s activities, fines, and banning members from the securities industry. See 15 U.S.C. 78o-3(b)(7).

<sup>46</sup> FINRA Rule 2010, Standards of Commercial Honor and Principles of Trade.

the purchase or sale of, any security by means of any manipulative, deceptive or other fraudulent device or contrivance.”<sup>47</sup> As will be discussed later,<sup>48</sup> FINRA can promulgate guidance pursuant to these rules clarifying that particular forms of conduct are manipulative.

## ***B. Ex Ante Strategies***

The range of ex ante legal devices available to regulate manipulation is much broader than is typically appreciated. Because of the pressures created by algorithmic manipulation on the efficacy of ex post enforcement, policymakers should make fuller and more aggressive use of this diverse set of tools.

### ***1. Reforming Market Structure***

Traditionally, regulators have relied heavily on surveillance and ex post enforcement to deter manipulation. As those tools become more difficult to deploy, however, regulators should consider reforms to the structure of markets that reduce manipulation indirectly. Appropriate changes could make the market structure itself more hostile to algorithmic strategies that harm market efficiency. Some ideas for reform, while advocated for other reasons, are also worth considering to reduce manipulation.

Two proposals for reforming contemporary equity market structure are “speed bumps” and frequent batched auctions.<sup>49</sup> Both structures alter the current market structure’s emphasis on speed and have generated both substantial opposition and support.<sup>50</sup> Conceptually, a speed bump is an intentional delay imposed by a trading venue on the processing of some or all orders directed to it. The goal of a speed bump is to reduce the incentive to implement ever-faster trading strategies, and to level the playing field for trade execution. The trading platform IEX, which began as an alternative trading system and transitioned to a stock exchange, brought speed bumps to national prominence.<sup>51</sup> IEX’s speed bump was a 350-microsecond delay that was imposed on all incoming trading instructions, before they could arrive at the exchange’s matching engine and be processed.<sup>52</sup> IEX’s speed bump is a “symmetric” speed bump that

---

<sup>47</sup> FINRA Rule 2020. Use of Manipulative, Deceptive or Other Fraudulent Devices.

<sup>48</sup> See *infra* notes \_\_\_ and accompanying text.

<sup>49</sup> For fuller discussion of the effects of extant frequent batched auctions on the equity market, see Gabriel V. Rauterberg, *Innovation in the Stock Market and Alternative Trading Systems*, forthcoming in FINANCIAL MARKET INFRASTRUCTURE (Jens-Hinrich Binder & Paolo Saguato eds. Oxford University Press 2021).

<sup>50</sup> See, e.g., John Nagel, Managing Director, Citadel Securities LLC, Comment Letter on Investors’ Exchange LLC Notice of Filing of Application for Registration as a National Securities Exchange Under Section 6 of the Securities Act of 1934, (Nov. 30, 2015), <https://www.sec.gov/comments/10-222/10222-28.pdf> (letter opposing IEX’s application to be a national stock exchange); BIDISHA CHAKRABARTY ET AL., EFFECTS OF A SPEED BUMP ON MARKET QUALITY AND EXCHANGE COMPETITION 6 n.6 (2019) (“NYSE MKT petitioned to create an access delay on its Pillar platform to add latency of 350 microseconds.”); Alexander Osipovich, *More Exchanges Add ‘Speed Bumps,’ Defying High-Frequency Traders*, WALL STREET JOURNAL (July 30, 2019), <https://www.wsj.com/articles/more-exchanges-add-speed-bumps-defying-high-frequency-traders-11564401611>.

<sup>51</sup> See IEX Form ATS, <https://iextrading.com/docs/IEX+Form+ATS+July+24.pdf>.

<sup>52</sup> IEX’s delay also applied to all outgoing communications regarding events in the order book. See Edwin Hu, *Intentional Access Delays, Market Quality, and Price Discovery: Evidence from IEX Becoming an Exchange* (Working paper 2019) (discussing features of IEX). The speed bump aimed to neutralize the ability of certain traders with speed advantages over most market participants. For instance, without a speed bump, a high-frequency trader can in principle cancel or adjust its quotes at the New York Stock Exchange based on a transaction at

applies in the same way to all market participants and order types.<sup>53</sup> Other exchanges have proposed “asymmetric” speed bumps that would selectively apply an intentional delay to only certain order types, such as marketable orders.<sup>54</sup>

A frequent batched auction is a market structure in which orders are matched and transactions occur at discrete, periodic intervals, rather than in continuous time.<sup>55</sup> Presently, all national stock exchanges offer continuous trading. Under continuous trading, incoming orders are processed serially as they arrive and as quickly as possible, which typically means in millionths of a second or less.<sup>56</sup> The result, critics argue, is to bake into the stock market a socially wasteful arms race for speed. Fast moving market participants can profit from being the first to trade on a new piece of information, such as the demand for a stock by other traders, even when that information becomes publicly available to many market participants at the same time.<sup>57</sup> Frequent batched auctions (for example, every tenth of a second), could end the arms race for speed, increase market efficiency, and promote fairness, without impairing the market’s ability to trade on economic information.<sup>58</sup>

Speed bumps and frequent batched auctions have been advocated for reasons other than deterring unlawful trading conduct, but they would also affect traders’ ability to engage in certain manipulative trading strategies that depend on speed. There are several manipulative strategies that depend on small speed differentials, and which could be eliminated (or rendered more difficult) by appropriate batched auctions or speed bumps. For example, high-frequency traders arguably engage in manipulation by making small, unprofitable orders to detect when

---

Nasdaq, even when a trader sent orders simultaneously to both exchanges. *See* Fox, Glosten & Rauterberg, *supra* note \_\_, at \_\_.

<sup>53</sup> IEX filed its application to be a stock exchange on September 15, 2015. SEC Release No. 34-75925, <https://www.sec.gov/rules/other/2015/investors-exchange-form-1.htm>. The application was approved on June 16, 2016. In the Matter of the Application of Investors’ Exchange, LLC for registration as a National Securities Exchange, SEC Release No. 34-78101.

<sup>54</sup> John McCrank, *Wall Street Braces for Rough Ride as Exchanges Seek More Speed Bumps*, REUTERS (Apr. 4, 2017, 7:13 AM) (NYSE, Nasdaq, and the Chicago Stock Exchange sought SEC approval for proposed speed bumps in 2016-17).

<sup>55</sup> *See, e.g.*, Eric Budish, Peter Cramton & John Shim, *The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response*, 130 Q.J. ECON. 1547, 1548 (2015); Farmer J. Doyne & Skouras Spyros, *Review of the Benefits of a Continuous Market vs. Randomised Stop Auctions and of Alternative Priority Rules* (Policy Options 7 and 12), UK GOVERNMENT’S FORESIGHT PROJECT, THE FUTURE OF COMPUTER TRADING IN FINANCIAL MARKETS, ECONOMIC IMPACT ASSESSMENT EIA 11 (2012); Elaine Wah & Michael Wellman, *Latency Arbitrage, Market Fragmentation, and Efficiency: A Two-Market Model*, Proceedings of the Fourteenth ACM Conference: Electronic Commerce (2013); Michael Wellman, “Countering High-Frequency Trading,” July 30, 2009, [http://mblog.lib.umich.edu/strategic/archives/2009/07/countering\\_high.html](http://mblog.lib.umich.edu/strategic/archives/2009/07/countering_high.html).

<sup>56</sup> In the market structure of certain alternative trading systems, batched auctions occur as frequently as every 900 microseconds (a microsecond is a millionth of a second). *See, e.g.*, IntelligentCross, Form ATS-N 11(a) (describing discrete auction matching mechanism).

<sup>57</sup> Budish, Peter Cramton & John Shim, *Implementation Details for Frequent Batch Auctions: Slowing Down Markets to the Blink of an Eye*, 104 AM. ECON. REV. 418 (2015); *see also* Eric Budish, Robin S. Lee & John J. Shim, *Will the Market Fix Market? A Theory of Stock Exchange Competition and Innovation* (Working Paper 2019).

<sup>58</sup> Elaine Wah, Michael Barr, Uday Rajan & Michael Wellman (2013), Public Comment in response to the U.S. Commodity Futures Trading Commission Concept Release on Risk Controls and System Safeguards for Automated Trading Environments. <http://comments.cftc.gov/PublicComments/ViewComment.aspx?id=59450>; *see also* Paul G. Mahoney, *Equity Market Structure Regulation: Time to Start Over*, forthcoming, MICH. BUS. & ENTR. L. REV. \_\_ (2020); Hester Peirce, *Meeting Market Structure Challenges Where They Are*, 43 J. CORP. L. 335, 357 (2018).

prices can be more (or less) easily moved and then use that information to raise the prices at which they trade with incoming order flow.<sup>59</sup> A second example involves traders who based on latency intervals between when orders arrive at different stock exchanges are able to raise the prices at which those orders transact by more rapidly altering quotes at those exchanges.<sup>60</sup> Of course, the importance of a structural reform's effects on manipulation is just one factor that weighs into a larger calculus about the desirability of such measures. But the reduction of certain manipulative strategies strengthens the relative case for these interventions.

## 2. *Imposing a Specialized Financial Transactions Tax*

A different kind of structural response is to adopt taxes specifically tailored to deter forms of manipulation. The idea of using taxes to affect financial behavior has a long history,<sup>61</sup> and one specific proposal—a financial transactions tax, prominently advocated by James Tobin—is the subject of a large popular and scholarly literature.<sup>62</sup> Proposals to reform trading market structure through tailored taxes have also enjoyed recent popularity.<sup>63</sup>

To illustrate the potential use of taxes to reduce manipulation, turn again to spoofing, where a market is manipulated by a trader's submission of non-bona fide orders designed to affect other participants' views, quotes, and trades. A tax on specific types of quoting and trading behavior could significantly reduce the incentives to spoof. For example, the tax could kick in whenever a market participant posted a large number of high-volume orders and where the ratio of cancellations of large orders to executions of small orders exceeded a certain threshold. By taxing excess canceled trades, the financial transaction tax could make it much more expensive to engage in spoofing, without significantly affecting market liquidity or burdening ordinary investors. A different type of tax could also be structured to require that the trader pay an extremely high capital gains rate on any small orders that execute within some period (e.g., one second) after very large contra-side orders are canceled.<sup>64</sup> There are some contexts, such as spread trading, where one might anticipate a large number of canceled trades, so appropriate exceptions might need to be developed for a canceled trade tax, perhaps focused on symmetric order cancellations or on indicia of market making.

---

<sup>59</sup> Adam D. Clark-Joseph, *Exploratory Trading*, Working Paper, Jan. 13, 2013, <http://www.nanex.net/-aqck2/4136/exploratorytrading.pdf>.

<sup>60</sup> MICHAEL LEWIS, *FLASH BOYS: A WALL STREET REVOLT* 200-01 (2014).

<sup>61</sup> See, e.g., James Tobin, *A Proposal for International Monetary Reform*, 4 EAST. ECON. J. 153 (1978).

<sup>62</sup> See, e.g., Laurence H. Summers & Victoria P. Summers, *When Financial Markets Work Too Well: A Cautious Case for a Securities Transactions Tax*, 3 J. FINANCIAL SERV. RES. 261 (1989); Joseph E. Stiglitz, *Tapping the Brakes: Are Less Active Markets Safer and Better for the Economy?*, Presented at the Federal Reserve Bank of Atlanta 2014 Financial Markets Conference, 2014.

<sup>63</sup> See, e.g., Bruno Biais, Thierry Foucault & Sophie Moinas, *Equilibrium Fast Trading*, 116 J. FIN. ECON. 292, 294 (2015) (arguing that a Pigovian tax could effectively reduce the negative externalities of high-speed trading); Budish, Crampton & Shim, *supra* note \_\_, at 1608-09 (discussing tax reforms as solutions to socially wasteful competition among high-frequency traders for speed to snipe stale quotes); Jiading Gai, Chen Yao & Mao Ye, Ye, Mao and Yao, Chen and Gai, Jiading, *The Externalities of High Frequency Trading*, WBS Finance Group Research Paper No. 180 (2013) (arguing that a Pigovian tax is one solution to a positional arms race among high-frequency traders for speed that has negative externalities for other traders).

<sup>64</sup> Such an idea is not completely foreign to the law. For instance, Section 16(b) requires an insider to return to her employer profits made from opening and closing a position in the firm's stock within six months. [Much as the short-swing profit rule 16(b)].



### 3. *Self-Regulation*

A much more moderate ex ante reform would focus on self-regulation. Regulators could promote norms for trading that depended on self-monitoring by traders. If adopted, firms would voluntarily surveil algorithms' trading strategies for potentially manipulative effects. One illustration of how to enhance traders' awareness of the potential ethical and legal drawbacks of certain trading strategies would be to incorporate it into materials for required certifications. There are a number of licensure certifications that it is common for securities traders to obtain (such as the Securities Industry Essentials Examination, Series 57 (Securities Trader Exam), or Series 24 (General Securities Principal Examination)). These certifications could encourage the individuals who develop and/or deploy trading algorithms to monitor them for their effects on market quality. They could also encourage managers to conduct periodic "audits" that certain commentators on algorithmic development have urged as a best practice.<sup>65</sup>

#### *C. Ex Post Strategies*

Ex ante strategies can eliminate or reduce the incidence of many forms of manipulation, but given the difficulty of enacting ex ante reforms, as well as the diversity of potential manipulative strategies, there will remain a need for ex post actions. How should these ex post actions work?

Core features of machine learning algorithms militate in favor of lower sanction, lower burden of proof proceedings. They also favor choosing a regulator who enjoys expertise in analyzing algorithmic trading. Some of the core actors in trading's regulatory ecosystem bring legal claims that are both high-sanction and high burden of proof. When the Department of Justice brings a criminal claim it has to satisfy the standard of proof beyond a reasonable doubt before a court can impose the severe sanction of a prison sentence. When an algorithmic manipulation is both knowing and socially destructive, then federal criminal prosecution may be precisely the right route. In general, however, enforcing algorithmic manipulation only through the criminal law will leave many harmful cases of market manipulation undeterred, especially those cases involving advanced algorithmic trading. Thus, we need to look to a spectrum of ex post enforcement by a range of regulatory actors.

There is a spectrum between those regulators with the highest burdens of proof, severest sanctions, and lengthiest timetables to enforcement, and those who can act most rapidly to impose weaker sanctions with relatively low burdens of proof. If criminal prosecution stands at one pole, many instances of algorithmic manipulation seem to favor regulators at the other pole. Who stands there? A plausible candidate may be broker-dealers.<sup>66</sup>

A broker-dealer is a regulatory category for a market participant that engages in either or both of the brokering and dealing functions in securities markets. Brokering involves facilitating trading interest by acting as the agent for other parties and executing orders on their

---

<sup>65</sup> Nicol Turner Lee, Paul Resnick & Genie Barton, *Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms*, Brookings Report.

<sup>66</sup> MICHAEL S. BARR, HOWELL E. JACKSON & MARGARET E. TAHYAR, FINANCIAL REGULATION: LAW AND POLICY 427-493, 519-549 (2016) (providing an overview of broker-dealers' roles in retail and exchange settings); Merritt Fox & Gabriel Rauterberg, *Stock Market Futurism*, 42 J. CORP. L. 793, 800 (2017) (addressing the role of broker-dealers).

behalf, while dealing involves facilitating counterparties' trading interest by acting as a principal and trading with the other party directly.<sup>67</sup> Broker-dealers hold a gatekeeping function in the equity market with several different regulatory facets; they are subject to a web of legal requirements, based in federal statutory law, but also including administrative rules, self-regulatory rules, and state law.<sup>68</sup> Aspects of their governing framework could be plausibly repurposed to monitor for manipulation by machine learning-trained algorithms. One rule in particular seems promising and offers a case study in administrative rules that other financial markets should consider adopting.

The directive in question, called the Market Access Rule, requires that a “broker or dealer with market access, or that provides . . . any other person with access . . . shall establish, document, and maintain a system of risk management controls and supervisory procedures reasonably designed to manage the financial, regulatory, and other risks of this business activity.”<sup>69</sup> For instance, Yesha Yadav emphasizes the value of the market access rule in providing a negligence-type cause of action, which has usefully authorized enforcement actions against firms with technological malfunctions and risky trading behavior.<sup>70</sup> Importantly, the risk management controls and procedures mandated by the Market Access Rule should be reasonably designed to “[p]revent the entry of orders unless there has been compliance with all regulatory requirements that must be satisfied on a pre-order entry basis.”<sup>71</sup> The rule establishes both pre- and post-trade supervisory requirements.<sup>72</sup> While it sounds technical (and tedious), the Market Access Rule provides perhaps the most robust existing legal foundation for imposing on a market actor surveillance duties to detect undesirable algorithmic strategies.

Several enforcement actions under the Market Access Rule illustrate its potential reach. Consider a FINRA proceeding against the high-frequency trading firm Two Sigma, based in part on violations of the Market Access Rule.<sup>73</sup> FINRA suggested that its review was prompted by “potentially unusual trading . . . [and] potentially manipulative behavior and/or erroneous/duplicative order entry.”<sup>74</sup> Upon review, FINRA concluded that Two Sigma “had inadequate risk management controls and supervisory procedures pertaining to certain aspects

---

<sup>67</sup> More precisely, a broker-dealer is an individual or institution that engages in one or all of three intermediation activities: acting as an agent executing orders on behalf of customers (i.e., a broker); transacting with customers as a principal for the entity's own account (i.e., a dealer); or operating an equity trading venue that is not a stock exchange. See MERRITT B. FOX, LAWRENCE R. GLOSTEN & GABRIEL V. RAUTERBERG, *THE NEW STOCK MARKET: LAW, ECONOMICS & POLICY* 261 (2019).

<sup>68</sup> BARR, JACKSON & TAHYAR, *supra* note 66, at 519.

<sup>69</sup> Exchange Act Rule 15c3-5, 17 C.F.R. § 240.15c3-5 (2015). See also Risk Management Controls for Brokers or Dealers with Market Access, 75 Fed. Reg. 69792 (Nov. 15, 2011). “Market access” is defined as access to trading on an exchange or alternative trading system. 17 C.F.R. § 240.15c3-5(b).

<sup>70</sup> Yadav, *supra* note **Error! Bookmark not defined.**, at 1057-1061; see also Lindsey C. Crump, *Regulating to Achieve Stability in the Domain of High-Frequency Trading*, 22 MICH. TELECOMM. & TECH. L. REV. 161, 171 (2015) (discussing the market access rule in the context of deterring market manipulation). The rationale for the Market Access Rule also mirrors in large part that offered by Scopino for greater use of the failure-to-supervise claim authorized by CFTC Regulation 166.3. See Scopino, *supra* note **Error! Bookmark not defined.**, at 235.

<sup>71</sup> 17 C.F.R. § 240.15c3-5(c)(2)(i).

<sup>72</sup> For a discussion of the rules, see Regulatory Developments 2011, 67 Bus. Law. 739, 783-87 (2012).

<sup>73</sup> Financial Industry Regulatory Authority Letter of Acceptance, Waiver and Consent No. 20130391658-04, Two Sigma Securities LLC, Respondent, [https://www.finra.org/sites/default/files/fda\\_documents/2013039165804\\_FDA\\_JG412297%20%282019-1563209360457%29.pdf](https://www.finra.org/sites/default/files/fda_documents/2013039165804_FDA_JG412297%20%282019-1563209360457%29.pdf).

<sup>74</sup> *Id.*

of market access, contrary to the requirements of [the market access rule], NASD Rule 3010 and FINRA Rule-2010.”<sup>75</sup> In particular, FINRA found that Two Sigma’s “controls and procedures for complying with regulatory requirements pursuant to [the Market Access Rule] were inadequate” because “the firm lacked any specific controls or surveillance to detect and prevent potentially manipulative activity in the form of spoofing, layering, and algorithmic gaming activity, and had insufficient surveillance for potential marking of the close activity during the 2012-2014 review period.”<sup>76</sup>

Enforcement actions have also been brought against broker-dealers for their failure to adequately monitor and review improper customer trading. For instance, the BATS EDGX exchange settled with JP Morgan in a proceeding based, inter alia, on JP Morgan using a third-party surveillance system to monitor for spoofing that set thresholds for generating spoofing alerts at “levels that were unreasonable to detect activity that might be indicative of layering and spoofing activity.”<sup>77</sup> FINRA has made clear that algorithmic trading is an enforcement priority for it.<sup>78</sup>

In between enforcement by FINRA, oversight by broker dealers and exchanges, and criminal prosecutions by the Department of Justice, lies civil enforcement actions by the SEC and CFTC. To make these agency enforcement actions effective in preserving market integrity, we argue that we must reconceptualize the law of manipulation to focus on trading behavior, not causality or intent. It is to that reconceptualization we turn to next.

#### ***D. Reconceptualizing Manipulation***

In the last part of this article, we suggest that manipulation law should, in addition to imposing criminal and civil sanctions for intentional misconduct, also, at least in part, impose liability based exclusively on *trading behavior* and making no reference to either causality or intent.<sup>79</sup> We have explained why it will be difficult to prosecute certain machine learning-

---

<sup>75</sup> *Id.*

<sup>76</sup> *Id.*

<sup>77</sup> BATS EDGX Exchange Inc. Letter of Acceptance, Waiver and Consent No. 2012034896-05, at ¶ 28. *Id.* at ¶ 29 (“As a result of the above, JPMS failed to adequately supervise certain of its customers’ trading, . . . and failed to detect potentially violative, spoofing activity that occurred on several days on the Exchange between August 12, 2015 and December 2, 2015.”).

<sup>78</sup> In its 2012 Regulatory and Examination Priorities Letter, FINRA noted that “even when there may not be a manipulative intent behind the trading, the velocity of HFT can result in unintended consequences in terms of quote generation and other activities. [This] requires firms using HFT strategies and other trading algorithms to be vigilant when testing these strategies pre- and post-launch to ensure that the strategies do not result in abusive trading and/or unintended consequences.” *Id.* (“Consistent with the Market Access Rule and other supervisory obligations, FINRA will assess whether firms have adequate testing and controls related to HFT and other algorithmic trading strategies. FINRA’s evaluation of firms’ controls may take the form of examinations and targeted investigations. Potential areas of review will include, among other things, the development, testing, deployment and maintenance of algorithmic codes; the adequacy of controls and follow- up regarding message rates; and procedures and controls to detect potential trading abuses such as, without limitation, wash sales and momentum ignition strategies.”).

<sup>79</sup> See Fletcher, *supra* note \_\_, at \_\_. Fletcher similarly argues for abandoning a solely intent-based manipulation standard and adopting one focused on harmful consequences. Fletcher focuses her analysis on the tension between the scienter requirement central to manipulation law and the conceptual and practical difficulties inherent in proving an algorithm’s mental state. She proposes abandoning the scienter requirement and refocusing manipulation law on the harm caused by a trading strategy. But her focus on harm is quite different

trained strategies under current law. In Section A, we sketch a few more reasons favoring a kind of behavioral reformulation of certain manipulative prohibitions. In Section B, we offer a few examples of how this process can occur.

### ***1. Reasons to Reconceptualize Manipulation***

What justifies a prohibition on manipulation, whether imposed criminally, civilly, or by a self-regulatory organization? There are many distinct theories offered to justify and guide manipulation law.<sup>80</sup> They are broadly unified in one respect, however, which is that they justify sanctioning manipulators on the basis of manipulation's negative effects on the quality of markets.<sup>81</sup> Unlike many crimes and torts, the ban on manipulative trading is not closely tethered to social interests in vindicating lost status or rights of victims or in necessary retribution against a malefactor. In fact, quite the opposite. A recurrent criticism of manipulation law, and a feature conceded even by its defenders, is that the line between manipulation and non-manipulative trading is often extremely difficult to draw.<sup>82</sup>

The different forms of manipulation vary considerably in the kind of mental states and conduct involved. Misstatement manipulation, for example, is quite close to fraud and requires an individual to knowingly make a falsehood.<sup>83</sup> This kind of manipulation may seem to fall in the heartland of conduct that the law should be comfortable punishing with fairly severe sanctions. But this is not true of other important forms of manipulation, including those most relevant to our analysis of potential manipulation by machine learning-trained algorithms.

There is serious ambiguity as to two salient features of open market manipulation, both of which are much remarked upon, but worth recalling here. First, it is a strikingly difficult task to construct a principled definition of manipulation that refers purely to a trader's trading behavior and does not make reference to its consequences on market quality. Second, it is

---

than the one FINRA suggests. Fletcher would create a rebuttable presumption that inefficient trading is unlawful. She also proposes lowering the intent standard to one of recklessness.

<sup>80</sup> See, e.g., Albert Kyle & Viswanathan, *How to Define Illegal Price Manipulation*, 98 AM. ECON. REV. 274, 274 (2008) (arguing that a trading strategy should only be viewed as manipulative if it reduces both price accuracy and liquidity); LAWRENCE E. HARRIS, *TRADING AND EXCHANGES* 266 (2002) (offering a nuanced account of manipulation based on its detrimental effects on market functioning); HAZEN, *THE LAW OF SECURITIES REGULATION*, 2 Law Sec. Reg. § 12.1 (6th ed. 2010) ("The purpose of the various statutes and rules prohibiting market manipulation is to prevent activities that rig the market and to thereby facilitate operation of the 'natural law' of supply and demand. . . . manipulation consists of any intentional interference with supply and demand."); Jonathan R. Macey & Maureen O'Hara, *From Markets to Venues: Securities Regulation in an Evolving World*, 58 STAN. L. REV. 563, 565 (2005) (emphasizing manipulation's negative effects on liquidity).

<sup>81</sup> For a discussion of the social functions of markets in general, see THIERRY FOUCAULT, MARCO PAGANO & AILSA RÖELL, *MARKET LIQUIDITY: THEORY, EVIDENCE, AND POLICY* 31 (2013) ("The two main roles of a securities market are to provide trading services for investors who wish to alter their portfolios, and to determine prices that can guide the allocation of capital by investors and firms. . . . [A] market is efficient if it enables investors to trade quickly and cheaply (i.e., if it is liquid) and if it incorporates new information quickly and accurately into prices.").

<sup>82</sup> See, e.g., Robert C. Lower, *Disruptions of the Futures Market: A Comment on Dealing with Market Manipulation*, 8 YALE J. REG. 391, 392 (1991) ("Manipulation is difficult to define . . . . [D]rawing a line between healthy economic behavior and that which is offensive has proved to be too subjective and imprecise to produce an effective regulatory tool."); *In re Henner*, 30 Agric. Dec. 1151 (U.S.D.A. 1971) ("Manipulation' is a vague term used in a wide and inclusive manner, possessing varying shades of meaning") (citation omitted).

<sup>83</sup> See Fox, Glosten & Rauterberg, *supra* note \_\_, at \_\_ (distinguishing among different forms of manipulation).

ambiguous as to whether the kind of morally culpable mental state we associate with a crime is present for certain forms of open market manipulation.

Consider some examples of open-market manipulations.<sup>84</sup> This manipulation exploits an institutional feature of equity market structure. The strategy is somewhat complex, but bears explanation because it illuminates the distinctive character of open market manipulation. Transactions in the stock market occur at multiple different kinds of trading venues, including stock exchanges, alternative trading systems, and internalization.<sup>85</sup> Informed traders—traders with a more accurate view of an asset’s value than what is reflected in its current market trading price—tend to trade on the exchanges, while uninformed traders, who possess no informational edge over other market participants, constitute a larger percentage of the trading population in internalization and on alternative trading systems.<sup>86</sup> Because this fact is widely appreciated, market participants react more to transactions at exchanges than other platforms because those trades are more likely to reflect informed traders’ new information about the value of stocks. This difference in the market’s reaction to trades at different venues is only possible because transactions on exchanges and transactions on other platforms are reported separately through distinct data feeds.<sup>87</sup>

With this background in place we can now see a profitable trading strategy. Suppose the stock for company Manne Inc. is trading at \$50. The same market participant opens two distinct brokerage accounts, one of which she uses to trade at stock exchanges and the other through internalization or alternative trading systems. The manipulator first purchases 1000 shares at exchanges. The price impact of these transactions will be relatively significant and the price may eventually move from \$50 to \$60, with an average purchasing price of \$55. The manipulator has purchased 1000 shares for \$55,000. The manipulator then unwinds her position by selling 1000 shares through internalization. The price impact of these transactions is less significant, resulting in the price of Manne Inc. only falling from \$60 to \$54. The average selling price was \$57, and so the manipulator sold her 1000 shares for \$57,000. As a result, the manipulator sold for \$57,000 the 1000 shares she bought for \$55,000, making a \$2,000 profit.

The SEC believes this trading strategy constitutes manipulation, and we agree. But we also think it is a reasonable view that this is simply a clever trading strategy that exploits heuristics used by other market participants. This trading strategy illustrates both of the difficulties of manipulation noted above. Defining manipulation with sufficient granularity so as to capture this strategy is challenging, and it is unclear whether a party pursuing this strategy has the kind of culpable intent we associate with criminal behavior.<sup>88</sup>

## 2. *Manipulation without Intent*

In this section, we illustrate how a purely behavioral or functional definition of manipulation might work. Jettisoning intent for criminal liability would be an overreaction, but

---

<sup>84</sup> See *infra* note \_\_\_ and accompanying text. *SEC v. Chen*, 1:19-cv-12127 in the United States District Court, District of Massachusetts; *SEC v. Taub and Shmalo*, 2:16-cv-09130 in the United States District Court, District of New Jersey.

<sup>85</sup> See Fox & Rauterberg, *supra* note 65, at \_\_\_.

<sup>86</sup> See Fox, Glostén & Rauterberg, *supra* note 68, at \_\_\_.

<sup>87</sup> *Id.*

<sup>88</sup> In an important sense the first problem, while distinct from the second, is a prerequisite for it. If traders were clearly on notice of what behavior was improper, then their failure to respect the legal directive itself would inculcate them.

alternative approaches for civil liability make sense.<sup>89</sup> A functional definition could be useful in capturing conduct that is detrimental to the market and which should be deterred through civil liability, regardless of trader scienter.

Intriguingly, FINRA has already begun to move toward offering definitions of manipulative activity that omit reference to individual mental states. Consider FINRA's Regulatory Notice 17-22, which illustrates how to prohibit manipulative activity without depending on scienter.<sup>90</sup> The rule changes sought to prohibit two "disruptive" types of trading and quoting activity. Most pertinently, one of the two rule changes effectively sought to prohibit spoofing, but in a purely functional or behavioral way that makes no reference to a mental state.

Under federal law, spoofing is "bidding or offering with the intent to cancel the bid or offer before execution."<sup>91</sup> This definition foregrounds scienter. It makes innocuous and routine behavior—placing bids or offers (and canceling them before execution)—unlawful if done with a specific intent. In contrast, FINRA Rule 5210.03 prohibits a "frequent pattern" in which "a party enters multiple limit orders on one side of the market at various price levels," "the level of supply and demand for the security changes" after those orders are entered, "the party enters one or more orders on the opposite side of the market" that subsequently execute, and the party afterward cancels the initial limit orders.<sup>92</sup> In guidance accompanying promulgation of the rule, FINRA is explicit that the rule "does not include an express intent element."<sup>93</sup> The rule also omits any express requirement of causation. By its terms, it does not require the party's limit orders to have demonstrably caused subsequent changes in the level of a security's supply and demand.

Several other plausible classes of manipulation might also be ripe for this kind of purely functional re-description and prohibition. The lack of intent and causation requirements also make clear why severe sanctions would be inappropriate. Instead, reconceptualizing spoofing from an intent-based crime, sanctionable with prison-time, to disruptive quoting behavior, sanctionable by a permanent cease and desist trading order, with fines and liability, shows how the law could adapt to deterring algorithmic manipulation. These changes suggest that

---

<sup>89</sup> See Fletcher, *supra* note \_\_\_, at \_\_\_.

<sup>90</sup> Regulatory Notice 17-22, FINRA Adopts Rules on Disruptive Quoting and Trading Activity and Expedited Proceedings Effective Date: December 15, 2016, <https://www.finra.org/rules-guidance/notices/17-22>. For analysis of the new FINRA rule, see Stanislav Dolgoplov, *The Doctrinal Quandary of Manipulative Practices in Securities Markets: Artificial Pricing, Price Discovery, and Liquidity Provision*, 45 J. CORP. L. 1, 37-39 (2019) (discussing how the new FINRA rule dispenses with a requirement of intent); Gideon Mark, *Spoofing and Layering*, 45 J. CORP. L. 399, 447 (2020) (noting that "No evidence of improper intent is required to establish a violation."). The stock exchange BATS also has a similar rule. See Order Approving BATS Proposed Rule Changes Relating to Disruptive Quoting and Trading Activity, Exchange Act Release No. 77171, 81 Fed. Reg. 9017 (Feb. 18, 2016); Rules of BATS BZX Exchange, Inc., BATS, at Rule 12.15; see also Michael Morelli, *Implementing High Frequency Trading Regulation: A Critical Analysis of Current Reforms*, 6 MICH. BUS. & ENTREPRENEURIAL L. REV. 201, 218 (2017) (discussing BATS rule).

<sup>91</sup> 7 U.S.C. § 6c(a)(5), adopted as part of the Dodd-Frank Wall Street Reform and Consumer Protection Act, Pub. L. No. 111-203, 124 Stat. 1376 (2010). *Id.* It is also defined as trading that is "of the character of, or is commonly known to the trade as, 'spoofing.'"

<sup>92</sup> See FINRA Rule 5210.03. More broadly, Rule 5210.03 prohibits FINRA's members from engaging in or facilitating "disruptive quoting and trading activity," and disruptive quoting and trading activity is defined to include the behavior above.

<sup>93</sup> See Regulatory Notice 17-22, *supra* note 88; see also FINRA, Proposed Rule Change to Provide a Process for an Expedited Proceeding and Adopt a Rule to Prohibit Disruptive Quoting and Trading Activity, at 15, <https://www.finra.org/rules-guidance/notices/17-22>.

regulators can move toward reconceptualizing manipulation without reference to individual scienter, and have taken the first steps in doing so. As a plausible next step, FINRA should consider adopting a purely functional definition of open market manipulation with an external interest, as in the benchmark example we discussed earlier. Congress could step in to provide the SEC and CFTC with enforcement authority based on this functional approach.

## **CONCLUSION**

Algorithmic trading poses challenges to the law of manipulation. We have shown that these challenges are particularly acute for one type of algorithmic trading, involving reinforcement learning with deep neural networks. Further research using simulated markets will help regulators better understand the market microstructure that is developing with deployment of deep reinforcement learning. While advances in developing simulated markets to improve detection of potential manipulation may assist regulators, we argue that legal reforms are also required to ensure market integrity. The Supreme Court should clarify that open market manipulation is unlawful. Congress and the regulators should consider ex ante strategies to reduce the space for manipulation, including speed bumps or frequent batch auctions, and a specialized financial transaction tax focused on canceled trades. Enforcement reforms are also required, including greater use of a wider spectrum of enforcement actions, especially those that can be deployed rapidly, with lower burdens of proof (and concomitant lesser liability). Lastly, we suggest a reconceptualization of manipulation law, to focus on actual trading behavior, without regard to causality and intent. The sum of these reforms would, ex ante, reduce the incentives to manipulate the market, and ex post, significantly increase the number and scope of enforcement actions. While a full analysis of the costs and benefits of these reforms is beyond the scope of this article, we believe they are directionally correct, and worthy of further study.